



- (51) International Patent Classification:  
H04N 19/36 (2014.01)
- (21) International Application Number:  
PCT/IB2014/065314
- (22) International Filing Date:  
14 October 2014 (14.10.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
61/891,004 15 October 2013 (15.10.2013) US
- (71) Applicant: NOKIA TECHNOLOGIES OY [FI/FI];  
Karaportti 3, FI-02610 Espoo (FI).
- (71) Applicant (for LC only): NOKIA USA INC. [US/US];  
200 South Mathilda Ave., Sunnyvale, California 94086 (US).
- (72) Inventor: HANNUKSELA, Miska Matias; Rusthollinrinne 2, FI-36110 Tampere (FI).
- (74) Agent: NOKIA TECHNOLOGIES OY; c/o Nokia USA Inc., Attn: IP Docketing, 200 South Mathilda Avenue, Sunnyvale, CA 94086 (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.

(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:**

— without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: VIDEO ENCODING AND DECODING

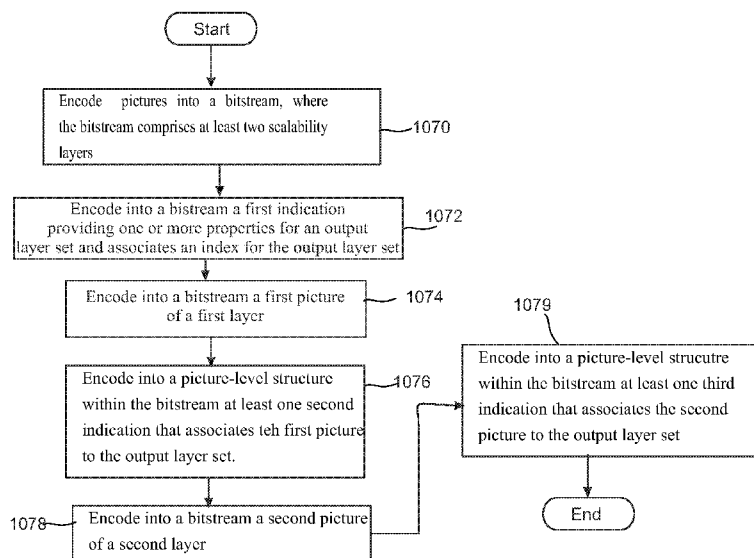


Fig. 7

(57) Abstract: There are disclosed various methods, apparatuses and computer program products for video encoding. In some embodiments pictures are encoded into a bitstream. The bitstream comprises at least two scalability layers and pictures being associated with access units. A first indication and a second indication are encoded into the bitstream. The first indication is configured to indicate an output layer. And the second indication is configured to indicate at least one alternative output layer. A first picture of said at least one alternative output layer is output by a decoding process of the bitstream when no picture of the output layer is in an access unit containing said first picture of said at least one alternative output layer.

WO 2015/056182 A2

## VIDEO ENCODING AND DECODING

### TECHNICAL FIELD

[0001] The present application relates generally to encoding and decoding of digital video  
5 material. In particular, the present application relates to scalable coding.

### BACKGROUND

[0002] This section is intended to provide a background or context to the invention that is recited in the  
claims. The description herein may include concepts that could be pursued, but are not necessarily ones that  
10 have been previously conceived or pursued. Therefore, unless otherwise indicated herein, what is described  
in this section is not prior art to the description and claims in this application and is not admitted to be prior  
art by inclusion in this section.

[0003] A video coding system may comprise an encoder that transforms an input video into a  
compressed representation suited for storage/transmission and a decoder that can uncompress the  
15 compressed video representation back into a viewable form. The encoder may discard some information in  
the original video sequence in order to represent the video in a more compact form, for example, to enable  
the storage/transmission of the video information at a lower bitrate than otherwise might be needed.

### SUMMARY

20 [0004] Some embodiments provide a method for encoding and decoding video information. In some  
embodiments there is provided methods, apparatuses and computer program products for video coding.

[0005] Various aspects of examples of the invention are provided in the detailed description.

[0006] According to a first aspect, there is provided a method comprising:

- encoding pictures into a bitstream, the bitstream comprising at least two scalability layers and  
25 pictures being associated with access units;
- encoding, into the bitstream, a first indication, which indicates an output layer; and
- encoding, into the bitstream, a second indication, which indicates at least one alternative output  
layer
- wherein the first indication and the second indication jointly have an impact that a first picture of  
30 the at least one alternative output layer is output by a decoding process of the bitstream when no  
picture of the output layer is in an access unit containing the first picture.

[0007] According to an embodiment, the method further comprises associating layers with one or more  
scalability dimensions; and associating layers with different value sets of the one or more scalability  
dimensions; wherein the second indication is independent of the one or more scalability dimensions.

[0008] According to a second aspect, there is provided an apparatus comprising at least one processor and at least one memory, said at least one memory stored with code thereon, which when executed by said at least one processor, causes an apparatus to perform at least the following:

- encode pictures into a bitstream, the bitstream comprising at least two scalability layers and pictures being associated with access units;
- encode, into the bitstream, a first indication, which indicates an output layer; and
- encode, into the bitstream, a second indication, which indicates at least one alternative output layer
- wherein the first indication and the second indication jointly have an impact that a first picture of the at least one alternative output layer is output by a decoding process of the bitstream when no picture of the output layer is in an access unit containing the first picture.

[0009] According to an embodiment of the apparatus, said at least one memory stored with code thereon, which when executed by said at least one processor, causes the apparatus to associate layers with one or more scalability dimensions; and associate layers with different value sets of the one or more scalability dimensions; wherein the second indication is independent of the one or more scalability dimensions.

[0010] According to a third aspect, there is provided a computer program product embodied on a non-transitory computer readable medium, comprising computer program code configured to, when executed on at least one processor, cause an apparatus or a system to:

- encode pictures into a bitstream, the bitstream comprising at least two scalability layers and pictures being associated with access units;
- encode, into the bitstream, a first indication, which indicates an output layer; and
- encode, into the bitstream, a second indication, which indicates at least one alternative output layer
- wherein the first indication and the second indication jointly have an impact that a first picture of the at least one alternative output layer is output by a decoding process of the bitstream when no picture of the output layer is in an access unit containing the first picture.

[0011] According to an embodiment of the computer program product, said at least one memory stored with code thereon, which when executed by said at least one processor, causes the apparatus to associate layers with one or more scalability dimensions; and associate layers with different value sets of the one or more scalability dimensions; wherein the second indication is independent of the one or more scalability dimensions.

[0012] According to a fourth aspect, there is provided a method comprising

- decoding pictures from a bitstream, the bitstream comprising at least two scalability layers and pictures being associated with access units;

- decoding, from the bitstream, a first indication, which indicates an output layer;
- decoding, from the bitstream, a second indication, which indicates at least one alternative output layer; and
- outputting a first picture of the at least one alternative output layer when no picture of the output layer is in a first access unit containing the first picture.

5 [0013] According to an embodiment, the method further comprises outputting a second picture of the output layer, the second picture being associated with a second access unit.

[0014] According to an embodiment, the method further comprises:

- associating layers with one or more scalability dimensions; and
- 10 – associating layers with different value sets of the one or more scalability dimensions;
- wherein the second indication is independent of the one or more scalability dimensions.

[0015] According to a fifth aspect, there is provided an apparatus comprising at least one processor and at least one memory, said at least one memory stored with code thereon, which when executed by said at least one processor, causes an apparatus to perform at least the following:

- 15 – decode pictures from a bitstream, the bitstream comprising at least two scalability layers and pictures being associated with access units;
- decode, from the bitstream, a first indication, which indicates an output layer;
- decode, from the bitstream, a second indication, which indicates at least one alternative output layer; and
- 20 – output a first picture of the at least one alternative output layer when no picture of the output layer is in a first access unit containing the first picture.

[0016] According to an embodiment of the apparatus, said at least one memory stored with code thereon, which when executed by said at least one processor, causes the apparatus to output a second picture of the output layer, the second picture being associated with a second access unit.

25 [0017] According to an embodiment of the apparatus, said at least one memory stored with code thereon, which when executed by said at least one processor, causes the apparatus to associate layers with one or more scalability dimensions; and associate layers with different value sets of the one or more scalability dimensions; wherein second indication is independent of the one or more scalability dimensions.

[0018] According to a sixth aspect, there is provided a computer program product embodied on a non-transitory computer readable medium, comprising computer program code configured to, when executed on at least one processor, cause an apparatus or a system to:

- decode pictures from a bitstream, the bitstream comprising at least two scalability layers and pictures being associated with access units;

- decode, from the bitstream, a first indication, which indicates an output layer;
- decode, from the bitstream, a second indication, which indicates at least one alternative output layer; and
- output a first picture of the at least one alternative output layer when no picture of the output layer is in a first access unit containing the first picture.

5 [0019] According to an embodiment of the computer program product, said at least one memory stored with code thereon, which when executed by said at least one processor, causes the apparatus to output a second picture of the output layer, the second picture being associated with a second access unit.

10 [0020] According to an embodiment of the computer program product, said at least one memory stored with code thereon, which when executed by said at least one processor, causes the apparatus to associate layers with one or more scalability dimensions; and associate layers with different value sets of the one or more scalability dimensions; wherein second indication is independent of the one or more scalability dimensions.

[0021] According to a seventh aspect there is provided a method comprising:

- 15 – encoding pictures into a bitstream, the bitstream comprising at least two scalability layers;
- encoding, into the bitstream, at least one first indication that provides one or more properties for an output layer set and associates an index for the output layer set;
- encoding, into the bitstream, a first picture of a first layer;
- encoding, into a picture-level structure within the bitstream, at least one second indication that
- 20 associates the first picture to the output layer set;
- encoding, into the bitstream, a second picture of a second layer;
- encoding, into a picture-level structure within the bitstream, at least one third indication that associates the second picture to the output layer set;
- wherein the first, the second and the third indications jointly have an impact that the pictures being
- 25 associated with the output layer set are output by a decoding process of the bitstream wherein the output layer set has been selected.

[0022] According to an eighth aspect, there is provided an apparatus comprising at least one processor and at least one memory, said at least one memory stored with code thereon, which when executed by said at least one processor, causes an apparatus to perform at least the following:

- 30 – encode pictures into a bitstream, the bitstream comprising at least two scalability layers;
- encode, into the bitstream, at least one first indication that provides one or more properties for an output layer set and associates an index for the output layer set;
- encode, into the bitstream, a first picture of a first layer;

- encode, into a picture-level structure within the bitstream, at least one second indication that associates the first picture to the output layer set;
- encode, into the bitstream, a second picture of a second layer;
- encode, into a picture-level structure within the bitstream, at least one third indication that  
5 associates the second picture to the output layer set;
- wherein the first, the second and the third indications jointly have an impact that the pictures being associated with the output layer set are output by a decoding process of the bitstream wherein the output layer set has been selected.

[0023] According to a ninth aspect, there is provided a computer program product embodied on a non-  
10 transitory computer readable medium, comprising computer program code configured to, when executed on at least one processor, cause an apparatus or a system to:

- encode pictures into a bitstream, the bitstream comprising at least two scalability layers;
- encode, into the bitstream, at least one first indication that provides one or more properties for an output layer set and associates an index for the output layer set;
- 15 – encode, into the bitstream, a first picture of a first layer;
- encode, into a picture-level structure within the bitstream, at least one second indication that associates the first picture to the output layer set;
- encode, into the bitstream, a second picture of a second layer;
- encode, into a picture-level structure within the bitstream, at least one third indication that  
20 associates the second picture to the output layer set;
- wherein the first, the second and the third indications jointly have an impact that the pictures being associated with the output layer set are output by a decoding process of the bitstream wherein the output layer set has been selected.

[0024] According to a tenth aspect, there is provided a method comprising:  
25 – decoding pictures from a bitstream, the bitstream comprising at least two scalability layers;

- decoding, from the bitstream, at least one first indication that provides one or more properties for an output layer set and associates an index for the output layer set;
- determining that the output layer set is used in outputting from the decoder;
- decoding, from the bitstream, a first picture of a first layer;
- 30 – decoding, from a picture-level structure within the bitstream, at least one second indication that associates the first picture to the output layer set;
- decoding, from the bitstream, a second picture of a second layer;

- decoding, from a picture-level structure within the bitstream, at least one third indication that associates the second picture to the output layer set;
- outputting the first picture and the second picture, as they are pictures that are associated with the output layer set.

5 [0025] According to a eleventh aspect, there is provided an apparatus comprising at least one processor and at least one memory, said at least one memory stored with code thereon, which when executed by said at least one processor, causes an apparatus to perform at least the following:

- decode pictures from a bitstream, the bitstream comprising at least two scalability layers;
- decode, from the bitstream, at least one first indication that provides one or more properties for  
10 an output layer set and associates an index for the output layer set;
- determine that the output layer set is used in outputting from the decoder;
- decode, from the bitstream, a first picture of a first layer;
- decode, from a picture-level structure within the bitstream, at least one second indication that associates the first picture to the output layer set;
- 15 – decode, from the bitstream, a second picture of a second layer;
- decode, from a picture-level structure within the bitstream, at least one third indication that associates the second picture to the output layer set;
- output the first picture and the second picture, as they are pictures that are associated with the output layer set.

20 [0026] According to a twelfth example, there is provided a computer program product embodied on a non-transitory computer readable medium, comprising computer program code configured to, when executed on at least one processor, cause an apparatus or a system to:

- decode pictures from a bitstream, the bitstream comprising at least two scalability layers;
- decode, from the bitstream, at least one first indication that provides one or more properties for  
25 an output layer set and associates an index for the output layer set;
- determine that the output layer set is used in outputting from the decoder;
- decode, from the bitstream, a first picture of a first layer;
- decode, from a picture-level structure within the bitstream, at least one second indication that associates the first picture to the output layer set;
- 30 – decode, from the bitstream, a second picture of a second layer;
- decode, from a picture-level structure within the bitstream, at least one third indication that associates the second picture to the output layer set;

- output the first picture and the second picture, as they are pictures that are associated with the output layer set.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

- 5 [0027] For a more complete understanding of example embodiments of the present invention, reference is now made to the following descriptions taken in connection with the accompanying drawings in which:
- [0028] Figure 1 shows a block diagram of a video coding system according to an embodiment;
- [0029] Figure 2 shows an apparatus for video coding according to an embodiment;
- [0030] Figure 3 shows an arrangement for video coding comprising a plurality of apparatuses;
- 10 [0031] Figure 4 shows an example of scalable video coding;
- [0032] Figure 5 shows a high level flow chart of an encoding method according to an embodiment;
- [0033] Figure 6 shows a high level flow chart of a decoding method according to an embodiment
- [0034] Figure 7 shows a high level flow chart of an encoding method according to an embodiment; and
- [0035] Figure 8 shows a high level flow chart of a decoding method according to an embodiment.

15

### **DETAILED DESCRIPTION OF SOME EXAMPLE EMBODIMENTS**

- [0036] In the following, several embodiments of the invention will be described in the context of one video coding arrangement. It is to be noted, however, that the invention is not limited to this particular arrangement. In fact, the different embodiments have applications widely in any environment where
- 20 improvement of scalable and/or multiview video coding is required. For example, the invention may be applicable to video coding systems like streaming systems, DVD players, digital television receivers, personal video recorders, systems and computer programs on personal computers, handheld computers and communication devices, as well as network elements such as transcoders and cloud computing arrangements where video data is handled.
- 25 [0037] Figure 1 shows a block diagram of a video coding system according to an example embodiment as a schematic block diagram of an exemplary apparatus or electronic device 50, which may incorporate a codec according to an embodiment of the invention. Figure 2 shows a layout of an apparatus according to an example embodiment. The elements of Figs. 1 and 2 will be explained next.
- [0038] The electronic device 50 may for example be a mobile terminal or user equipment of a wireless
- 30 communication system. However, it would be appreciated that embodiments of the invention may be implemented within any electronic device or apparatus which may require encoding and decoding or encoding or decoding video images.
- [0039] The apparatus 50 may comprise a housing 30 for incorporating and protecting the device. The apparatus 50 further may comprise a display 32 in the form of a liquid crystal display. In other

embodiments of the invention the display may be any suitable display technology suitable to display an image or video. The apparatus 50 may further comprise a keypad 34. In other embodiments of the invention any suitable data or user interface mechanism may be employed. For example the user interface may be implemented as a virtual keyboard or data entry system as part of a touch-sensitive display. The apparatus  
5 may comprise a microphone 36 or any suitable audio input which may be a digital or analogue signal input. The apparatus 50 may further comprise an audio output device which in embodiments of the invention may be any one of: an earpiece 38, speaker, or an analogue audio or digital audio output connection. The apparatus 50 may also comprise a battery 40 (or in other embodiments of the invention the device may be powered by any suitable mobile energy device such as solar cell, fuel cell or clockwork generator). The  
10 apparatus may further comprise a camera 42 capable of recording or capturing images and/or video. In some embodiments the apparatus 50 may further comprise an infrared port for short range line of sight communication to other devices. In other embodiments the apparatus 50 may further comprise any suitable short range communication solution such as for example a Bluetooth wireless connection or a USB/firewire wired connection.

15 [0040] The apparatus 50 may comprise a controller 56 or processor for controlling the apparatus 50. The controller 56 may be connected to memory 58 which in embodiments of the invention may store both data in the form of image and audio data and/or may also store instructions for implementation on the controller 56. The controller 56 may further be connected to codec circuitry 54 suitable for carrying out coding and decoding of audio and/or video data or assisting in coding and decoding carried out  
20 by the controller.

[0041] The apparatus 50 may further comprise a card reader 48 and a smart card 46, for example a UICC and UICC reader for providing user information and being suitable for providing authentication information for authentication and authorization of the user at a network.

[0042] The apparatus 50 may comprise radio interface circuitry 52 connected to the controller and  
25 suitable for generating wireless communication signals for example for communication with a cellular communications network, a wireless communications system or a wireless local area network. The apparatus 50 may further comprise an antenna 44 connected to the radio interface circuitry 52 for transmitting radio frequency signals generated at the radio interface circuitry 52 to other apparatus(es) and for receiving radio frequency signals from other apparatus(es).

30 [0043] In some embodiments of the invention, the apparatus 50 comprises a camera capable of recording or detecting individual frames which are then passed to the codec 54 or controller for processing. In some embodiments of the invention, the apparatus may receive the video image data for processing from another device prior to transmission and/or storage. In some embodiments of the

invention, the apparatus 50 may receive either wirelessly or by a wired connection the image for coding/decoding.

[0044] Fig. 3 shows an arrangement for video coding comprising a plurality of apparatuses, networks and network elements according to an example embodiment. With respect to Figure 3, an example of a system within which embodiments of the present invention can be utilized is shown. The system 10 comprises multiple communication devices which can communicate through one or more networks. The system 10 may comprise any combination of wired or wireless networks including, but not limited to a wireless cellular telephone network (such as a GSM, UMTS, CDMA network etc), a wireless local area network (WLAN) such as defined by any of the IEEE 802.x standards, a Bluetooth personal area network, an Ethernet local area network, a token ring local area network, a wide area network, and the Internet.

[0045] The system 10 may include both wired and wireless communication devices or apparatus 50 suitable for implementing embodiments of the invention. For example, the system shown in Figure 3 shows a mobile telephone network 11 and a representation of the internet 28. Connectivity to the internet 28 may include, but is not limited to, long range wireless connections, short range wireless connections, and various wired connections including, but not limited to, telephone lines, cable lines, power lines, and similar communication pathways.

[0046] The example communication devices shown in the system 10 may include, but are not limited to, an electronic device or apparatus 50, a combination of a personal digital assistant (PDA) and a mobile telephone 14, a PDA 16, an integrated messaging device (IMD) 18, a desktop computer 20, a notebook computer 22. The apparatus 50 may be stationary or mobile when carried by an individual who is moving. The apparatus 50 may also be located in a mode of transport including, but not limited to, a car, a truck, a taxi, a bus, a train, a boat, an airplane, a bicycle, a motorcycle or any similar suitable mode of transport.

[0047] Some or further apparatuses may send and receive calls and messages and communicate with service providers through a wireless connection 25 to a base station 24. The base station 24 may be connected to a network server 26 that allows communication between the mobile telephone network 11 and the internet 28. The system may include additional communication devices and communication devices of various types. The communication devices may communicate using various transmission technologies including, but not limited to, code division multiple access (CDMA), global systems for mobile communications (GSM), universal mobile telecommunications system (UMTS), time divisional multiple access (TDMA), frequency division multiple access (FDMA), transmission control protocol-internet protocol (TCP-IP), short messaging service (SMS), multimedia messaging service (MMS), email, instant messaging service (IMS), Bluetooth, IEEE 802.11 and any similar wireless communication technology. A communications device involved in implementing various embodiments

of the present invention may communicate using various media including, but not limited to, radio, infrared, laser, cable connections, and any suitable connection.

[0048] The H.264/AVC standard was developed by the Joint Video Team (JVT) of the Video Coding Experts Group (VCEG) of the Telecommunications Standardization Sector of International  
5 Telecommunication Union (ITU-T) and the Moving Picture Experts Group (MPEG) of International Organisation for Standardization (ISO) / International Electrotechnical Commission (IEC). The H.264/AVC standard is published by both parent standardization organizations, and it is referred to as ITU-T Recommendation H.264 and ISO/IEC International Standard 14496-10, also known as MPEG-4 Part 10  
10 Advanced Video Coding (AVC). There have been multiple versions of the H.264/AVC standard, integrating new extensions or features to the specification. These extensions include Scalable Video Coding (SVC) and Multiview Video Coding (MVC).

[0049] The H.265/HEVC standard was developed by the Joint Collaborative Team on Video Coding (JCT-VC) of VCEG and MPEG. The H.265/HEVC standard will be published by both parent  
15 standardization organizations, and is referred to as ITU-T Recommendation H.265 and ISO/IEC International Standard 23008-2, also known as MPEG-H Part 2 High Efficiency Video Coding (HEVC). There are currently ongoing standardization projects to develop extensions to H.265/HEVC, including scalable, multiview, three-dimensional, and fidelity range extensions.

[0050] A scalable video codec for quality scalability (also known as Signal-to-Noise or SNR) and/or  
20 spatial scalability may be implemented as follows. For a base layer, a conventional non-scalable video encoder and decoder is used. The reconstructed/decoded pictures of the base layer are included in the reference picture buffer for an enhancement layer. In H.264/AVC, HEVC, and similar codecs using reference picture list(s) for inter prediction, the base layer decoded pictures may be inserted into a reference picture list(s) for coding/decoding of an enhancement layer pictures similarly to the decoded reference pictures of the enhancement layer. Consequently, the encoder may choose a base-layer reference picture as  
25 inter prediction reference and may indicate its use e.g. with a reference picture index in the coded bitstream. The decoder decodes from the bitstream, for example from a reference picture index, that a base-layer picture is used as an inter prediction reference for the enhancement layer. When a decoded base-layer picture is used as a prediction reference for an enhancement layer, it is referred to as an inter-layer reference picture.

30 [0051] Various technologies for providing three-dimensional (3D) video content are currently investigated and developed. Especially, intense studies have been focused on various multiview applications wherein a viewer is able to see only one pair of stereo video from a specific viewpoint and another pair of stereo video from a different viewpoint. One of the most feasible approaches for such multiview applications has turned out to be such wherein only a limited number of views, e.g. a mono or a

stereo video plus some supplementary data, is provided to a decoder side and all required views are then rendered (i.e. synthesized) locally by the decoder to be displayed on a display.

[0052] Some key definitions, bitstream and coding structures, and concepts of H.264/AVC and HEVC are described in this section as an example of a video encoder, decoder, encoding method, decoding method, and a bitstream structure, wherein the embodiments may be implemented. Some of the key definitions, bitstream and coding structures, and concepts of H.264/AVC are the same as in HEVC – hence, they are described below jointly. The aspects of the invention are not limited to H.264/AVC or HEVC, but rather the description is given for one possible basis on top of which the invention may be partly or fully realized.

[0053] When describing H.264/AVC and HEVC as well as in example embodiments, common notation for arithmetic operators, logical operators, relation operators, bit-wise operators, assignment operators, and range notation e.g. as specified in H.264/AVC or HEVC may be used. Furthermore, common mathematical functions e.g. as specified in H.264/AVC or HEVC may be used and a common order or precedence and execution order (from left to right or from right to left) of operators e.g. as specified in H.264/AVC or HEVC may be used.

[0054] When describing H.264/AVC and HEVC as well as in example embodiments, the following description may be used to specify the parsing process of each syntax element.

- b(8): byte having any pattern of bit string (8 bits).
- se(v): signed integer Exp-Golomb-coded syntax element with the left bit first.
- u(n): unsigned integer using n bits. When n is “v” in the syntax table, the number of bits varies in a manner dependent on the value of other syntax elements. The parsing process for this descriptor is specified by n next bits from the bitstream interpreted as a binary representation of an unsigned integer with the most significant bit written first.
- ue(v): unsigned integer Exp-Golomb-coded syntax element with the left bit first.

[0055] An Exp-Golomb bit string may be converted to a code number (codeNum) for example using the following table:

Bit string	codeNum
1	0
0 1 0	1
0 1 1	2
0 0 1 0 0	3
0 0 1 0 1	4
0 0 1 1 0	5
0 0 1 1 1	6
0 0 0 1 0 0 0	7
0 0 0 1 0 0 1	8
0 0 0 1 0 1 0	9
...	...

[0056] A code number corresponding to an Exp-Golomb bit string may be converted to  $se(v)$  for example using the following table:

codeNum	syntax element value
0	0
1	1
2	-1
3	2
4	-2
5	3
6	-3
...	...

5 [0057] When describing H.264/AVC and HEVC as well as in example embodiments, syntax structures, semantics of syntax elements, and decoding process may be specified as follows. Syntax elements in the bitstream are represented in **bold** type. Each syntax elements is described by its name (all lower case letters with underscore characteristics), optionally its one or two syntax categories, and one or two descriptors for its method of coded representation. The decoding process behaves according to the value of the syntax  
 10 element and to the values of previously decoded syntax elements. When a value of a syntax element is used in the syntax tables or the text, named by a mixture of lower case and upper case letter and without any underscore characters. Variables starting with an upper case letter are derived for the decoding of the current syntax structure and all depending syntax structures. Variables starting with an upper case letter may be used in the decoding process for later syntax structures without mentioning the originating syntax structure  
 15 of the variable. Variables starting with a lower case letter are only used within the context in which they are derived. In some cases, "mnemonic" names for syntax element values or variable values are used interchangeably with their numerical values. Sometimes "mnemonic" names are used without any associated numerical values. The association of values and names is specified in the text. The names are constructed from one or more groups of letters separated by an underscore character. Each group starts with  
 20 an upper case letter and may contain more upper case letters.

[0058] When describing H.264/AVC and HEVC as well as in example embodiments, a syntax structure may be specified using the following. A group of statements enclosed in curly brackets is a compound statement and is treated functionally as a single statement. A "while" structure specifies a test of whether a condition is true, and if true, specifies evaluation of a statement (or compound statement) repeatedly until  
 25 the condition is no longer true. A "do ... while" structure specifies evaluation of a statement once, followed by a test of whether a condition is true, and if true, specifies repeated evaluation of the statement until the condition is no longer true. An "if ... else" structure specifies a test of whether a condition is true, and if the condition is true, specifies evaluation of a primary statement, otherwise, specifies evaluation of an

alternative statement. The "else" part of the structure and the associated alternative statement is omitted if no alternative statement evaluation is needed. A "for" structure specifies evaluation of an initial statement, followed by a test of a condition, and if the condition is true, specifies repeated evaluation of a primary statement followed by a subsequent statement until the condition is no longer true.

5 [0059] In the description of existing standards as well as in the description of example embodiments, a phrase "by external means" or "through external means" may be used. For example, an entity, such as a syntax structure or a value of a variable used in the decoding process, may be provided "by external means" to the decoding process. The phrase "by external means" may indicate that the entity is not included in the bitstream created by the encoder, but rather conveyed externally from the bitstream for example using a control protocol. It may alternatively or additionally mean that the entity is not created by the encoder, but  
10 may be created for example in the player or decoding control logic or alike that is using the decoder. The decoder may have an interface for inputting the external means, such as variable values.

[0060] Similarly to many earlier video coding standards, the bitstream syntax and semantics as well as the decoding process for error-free bitstreams are specified in H.264/AVC and HEVC. The encoding  
15 process is not specified, but encoders must generate conforming bitstreams. Bitstream and decoded conformance can be verified with the Hypothetical Reference Decoder losses, but the use of the tools in encoding is optional and no decoding process has been specified for erroneous bitstreams.

[0061] The elementary unit for the input to an H.264/AVC or HEVC encoder and the output of an H.264/AVC or HEVC decoder, respectively, is a picture. A picture may either be a frame or a field. A frame  
20 comprises a matrix of luma samples and corresponding chroma samples. A field is a set of alternate sample rows of a frame and may be used as encoder input, when the source signal is interlaced. Chroma pictures may be subsampled when compared to luma pictures. For example, in the 4:2:0 sampling pattern the spatial resolution of chroma pictures is half of that of the luma picture along both coordinate axes.

[0062] In H.264/AVC, a macroblock is a 16x16 block of luma samples and the corresponding blocks  
25 of chroma samples. For example, in the 4:2:0 sampling pattern, a macroblock contains one 8x8 block of chroma samples per each chroma component. In H.264/AVC, a picture is partitioned to one or more slice groups, and a slice group contains one or more slices. In H.264/AVC, a slice consists of an integer number of macroblocks ordered consecutively in the raster scan within a particular slice group.

[0063] In HEVC, video pictures are divided into coding units (CU) covering the area of the picture. A  
30 CU consists of one or more prediction units (PU) defining the prediction process for the samples within the CU and one or more transform units (TU) defining the prediction error coding process for the samples in the CU. Typically, a CU consists of a square block of samples with a size selectable from a predefined set of possible CU sizes. A CU with the maximum allowed size is typically named as CTU (coding tree unit) and the video picture is divided into non-overlapping CTUs. An CTU can be further split into a combination

of smaller CUs, e.g. by recursively splitting the CTU and resultant CUs. Each resulting CU typically has at least one PU and at least one TU associated with it. Each PU and TU can further be split into smaller PUs and TUs in order to increase granularity of the prediction and prediction error coding processes, respectively. The PU splitting can be realized by splitting the CU into four equal size square PUs or splitting  
5 the CU into two rectangle PUs vertically or horizontally in a symmetric or asymmetric way. The division of the image into CUs, and division of CUs into PUs and TUs is typically signaled in the bitstream allowing the decoder to reproduce the intended structure of these units

[0064] In HEVC, a picture can be partitioned in tiles, which are rectangular and contain an integer number of CTUs. In HEVC, the partitioning to tiles forms a regular grid, where heights and widths of tiles  
10 differ from each other by one CTU at the maximum.

[0065] In HEVC, a slice is defined to be an integer number of coding tree units contained in one independent slice segment and all subsequent dependent slice segments (if any) that precede the next independent slice segment (if any) within the same access unit. In HEVC, a slice segment is defined to be an integer number of coding tree units ordered consecutively in the tile scan and contained in a single NAL  
15 unit. The division of each picture into slice segments is a partitioning. In HEVC, an independent slice segment is defined to be a slice segment for which the values of the syntax elements of the slice segment header are not inferred from the values for a preceding slice segment, and a dependent slice segment is defined to be a slice segment for which the values of some syntax elements of the slice segment header are inferred from the values for the preceding independent slice segment in decoding order. In HEVC, a slice  
20 header is defined to be the slice segment header of the independent slice segment that is a current slice segment or is the independent slice segment that precedes a current dependent slice segment, and a slice segment header is defined to be a part of a coded slice segment containing the data elements pertaining to the first or all coding tree units represented in the slice segment. The CUs are scanned in the raster scan order of LCUs within tiles or within a picture, if tiles are not in use. Within an LCU, the CUs have a specific  
25 scan order.

[0066] In a Working Draft (WD) 5 of HEVC, some key definitions and concepts for picture partitioning are defined as follows. A partitioning is defined as the division of a set into subsets such that each element of the set is in exactly one of the subsets.

[0067] A basic coding unit in a HEVC WD5 is a treeblock. A treeblock is an NxN block of luma samples and two corresponding blocks of chroma samples of a picture that has three sample arrays, or an  
30 NxN block of samples of a monochrome picture or a picture that is coded using three separate colour planes. A treeblock may be partitioned for different coding and decoding processes. A treeblock partition is a block of luma samples and two corresponding blocks of chroma samples resulting from a partitioning of a treeblock for a picture that has three sample arrays or a block of luma samples resulting from a partitioning

of a treeblock for a monochrome picture or a picture that is coded using three separate colour planes. Each treeblock is assigned a partition signalling to identify the block sizes for intra or inter prediction and for transform coding. The partitioning is a recursive quadtree partitioning. The root of the quadtree is associated with the treeblock. The quadtree is split until a leaf is reached, which is referred to as the coding node. The coding node is the root node of two trees, the prediction tree and the transform tree. The prediction tree specifies the position and size of prediction blocks. The prediction tree and associated prediction data are referred to as a prediction unit. The transform tree specifies the position and size of transform blocks. The transform tree and associated transform data are referred to as a transform unit. The splitting information for luma and chroma is identical for the prediction tree and may or may not be identical for the transform tree. The coding node and the associated prediction and transform units form together a coding unit.

[0068] In a HEVC WD5, pictures are divided into slices and tiles. A slice may be a sequence of treeblocks but (when referring to a so-called fine granular slice) may also have its boundary within a treeblock at a location where a transform unit and prediction unit coincide. Treeblocks within a slice are coded and decoded in a raster scan order. For the primary coded picture, the division of each picture into slices is a partitioning.

[0069] In a HEVC WD5, a tile is defined as an integer number of treeblocks co-occurring in one column and one row, ordered consecutively in the raster scan within the tile. For the primary coded picture, the division of each picture into tiles is a partitioning. Tiles are ordered consecutively in the raster scan within the picture. Although a slice contains treeblocks that are consecutive in the raster scan within a tile, these treeblocks are not necessarily consecutive in the raster scan within the picture. Slices and tiles need not contain the same sequence of treeblocks. A tile may comprise treeblocks contained in more than one slice. Similarly, a slice may comprise treeblocks contained in several tiles.

[0070] In H.264/AVC and HEVC, in-picture prediction may be disabled across slice boundaries. Thus, slices can be regarded as a way to split a coded picture into independently decodable pieces, and slices are therefore regarded as elementary units for transmission. In many cases, encoders may indicate in the bitstream which types of in-picture prediction are turned off across slice boundaries, and the decoder operation takes this information into account for example when concluding which prediction sources are available. For example, samples from a neighboring macroblock or CU may be regarded as unavailable for intra prediction, if the neighboring macroblock or CU resides in a different slice.

[0071] A syntax element may be defined as an element of data represented in the bitstream. A syntax structure may be defined as zero or more syntax elements present together in the bitstream in a specified order.

[0072] The elementary unit for the output of an H.264/AVC or HEVC encoder and the input of an H.264/AVC or HEVC decoder respectively, is a Network Abstraction Layer (NAL) unit. For transport over

packet-oriented networks or storage into structured files, NAL units may be encapsulated into packets or similar structures. A bytestream format has been specified in H.264/AVC and HEVC for transmission or storage environments that do not provide framing structures. The bytestream format separates NAL units from each other by attaching a start code in front of each NAL unit. To avoid false detection of NAL unit boundaries, encoders may run a byte-oriented start code emulation prevention algorithm, which adds an emulation prevention byte to the NAL unit payload if a start code would have occurred otherwise. In order to enable straightforward gateway operation between packet- and stream-oriented systems, start code emulation prevention may always be performed regardless of whether the bytestream format is in use or not.

5 [0073] NAL units consist of a header and payload. In H.264/AVC, the NAL unit header indicates the type of the NAL unit and whether a coded slice contained in the NAL unit is a part of a reference picture or a non-reference picture. H.264/AVC includes a 2-bit `nal_ref_idc` syntax element, which when equal to 0 indicates that a coded slice contained in the NAL unit is a part of a non-reference picture and when greater than 0 indicates that a coded slice contained in the NAL unit is a part of a reference picture. The header for SVC and MVC NAL units may be additionally contain various indications related to scalability and multiview hierarchy.

[0074] In HEVC, a two-byte NAL unit header is used for all specified NAL unit types. The NAL unit header contains one reserved bit, a six-bit NAL unit type indication, a six-bit reserved field (called `nuh_layer_id`) and a three-bit `temporal_id_plus1` indication for temporal level. The `temporal_id_plus1` syntax element may be regarded as a temporal identifier for the NAL unit, and a zero-based `TemporalId` variable may be derived as follows:  $\text{TemporalId} = \text{temporal\_id\_plus1} - 1$ . `TemporalId` equal to 0 corresponds to the lowest temporal level. The value of `temporal_id_plus1` is required to be non-zero in order to avoid start code emulation involving the two NAL unit header bytes. The bitstream created by excluding all VCL NAL units having a `TemporalId` greater than or equal to a selected value and including all other VCL NAL units remains conforming. Consequently, a picture having `TemporalId` equal to `TID` does not use any picture having a `TemporalId` greater than `TID` as inter prediction reference. A sub-layer or a temporal sub-layer may be defined to be a temporal scalable layer of a temporal scalable bitstream, consisting of VCL NAL units with a particular value of the `TemporalId` variable and the associated non-VCL NAL units. Without loss of generality, in some example embodiments a variable `LayerId` is derived from the value of `nuh_layer_id` for example as follows:  $\text{LayerId} = \text{nuh\_layer\_id}$ . In the following, `LayerId`, `nuh_layer_id` and `layer_id` are used interchangeably unless otherwise indicated.

20 [0075] It is expected that `nuh_layer_id` and/or similar syntax elements in NAL unit header would carry information on the scalability hierarchy. For example, the `LayerId` value may be mapped to values of variables or syntax elements describing different scalability dimensions, such as `quality_id` or similar,

dependency\_id or similar, any other type of layer identifier, view order index or similar, view identifier, an indication whether the NAL unit concerns depth or texture i.e. depth\_flag or similar, or an identifier similar to priority\_id of SVC indicating a valid sub-bitstream extraction if all NAL units greater than a specific identifier value are removed from the bitstream. nuh\_layer\_id and/or similar syntax elements may be partitioned into one or more syntax elements indicating scalability properties. For example, a certain number of bits among nuh\_layer\_id and/or similar syntax elements may be used for dependency\_id or similar, while another certain number of bits among nuh\_layer\_id and/or similar syntax elements may be used for quality\_id or similar. Alternatively, a mapping of LayerId values or similar to values of variables or syntax elements describing different scalability dimensions may be provided for example in a Video Parameter Set, a Sequence Parameter Set or another syntax structure.

[0076] NAL units can be categorized into Video Coding Layer (VCL) NAL units and non-VCL NAL units. VCL NAL units are typically coded slice NAL units. In H.264/AVC, coded slice NAL units contain syntax elements representing one or more coded macroblocks, each of which corresponds to a block of samples in the uncompressed picture. In HEVC, coded slice NAL units contain syntax elements representing one or more CU.

[0077] In H.264/AVC, a coded slice NAL unit can be indicated to be a coded slice in an Instantaneous Decoding Refresh (IDR) picture or coded slice in a non-IDR picture.

[0078] In HEVC, a coded slice NAL unit can be indicated to be one of the following types:

<b>nal_unit_type</b>	<b>Name of nal_unit_type</b>	<b>Content of NAL unit and RBSP syntax structure</b>
0, 1	TRAIL_N, TRAIL_R	Coded slice segment of a non-TSA, non-STSA trailing picture slice_segment_layer_rbsp()
2, 3	TSA_N, TSA_R	Coded slice segment of a TSA picture slice_segment_layer_rbsp()
4, 5	STSA_N, STSA_R	Coded slice segment of an STSA picture slice_layer_rbsp()
6, 7	RADL_N, RADL_R	Coded slice segment of a RADL picture slice_layer_rbsp()

8, 9	RASL_N, RASL_R,	Coded slice segment of a RASL picture slice_layer_rbsp( )
10, 12, 14	RSV_VCL_N10 RSV_VCL_N12 RSV_VCL_N14	Reserved // reserved non-RAP non-reference VCL NAL unit types
11, 13, 15	RSV_VCL_R11 RSV_VCL_R13 RSV_VCL_R15	Reserved // reserved non-RAP reference VCL NAL unit types
16, 17, 18	BLA_W_LP BLA_W_DLP (a.k.a. BLA_W_RADL) BLA_N_LP	Coded slice segment of a BLA picture slice_segment_layer_rbsp( )
19, 20	IDR_W_DLP (a.k.a. IDR_W_RADL) IDR_N_LP	Coded slice segment of an IDR picture slice_segment_layer_rbsp( )
21	CRA_NUT	Coded slice segment of a CRA picture slice_segment_layer_rbsp( )
22, 23	RSV_RAP_VCL22.. RSV_RAP_VCL23	Reserved // reserved RAP VCL NAL unit types
24..31	RSV_VCL24.. RSV_VCL31	Reserved // reserved non-RAP VCL NAL unit types

[0079] In a draft HEVC standard, abbreviations for picture types may be defined as follows: trailing (TRAIL) picture, Temporal Sub-layer Access (TSA), Step-wise Temporal Sub-layer Access (STSA), Random Access Decodable Leading (RADL) picture, Random Access Skipped Leading (RASL) picture, Broken Link Access (BLA) picture, Instantaneous Decoding Refresh (IDR) picture, Clean Random Access (CRA) picture.

[0080] A Random Access Point (RAP) picture, which may also or alternatively be referred to as intra random access point (IRAP) picture, is a picture where each slice or slice segment has nal\_unit\_type in the range of 16 to 23, inclusive. A RAP picture contains only intra-coded slices, and may be a BLA picture, a CRA picture or an IDR picture. The first picture in the bitstream is a RAP picture. Provided the necessary parameter sets are available when they need to be activated, the RAP picture and all subsequent non-RASL pictures in decoding order can be correctly decoded without performing the decoding process of any

pictures that precede the RAP picture in decoding order. There may be pictures in a bitstream that contain only intra-coded slices that are not RAP pictures.

[0081] In HEVC a CRA picture may be the first picture in the bitstream in decoding order, or may appear later in the bitstream. CRA pictures in HEVC allow so-called leading pictures that follow the CRA picture in decoding order but precede it in output order. Some of the leading pictures, so-called RASL pictures, may use pictures decoded before the CRA picture as a reference. Pictures that follow a CRA picture in both decoding and output order are decodable if random access is performed at the CRA picture, and hence clean random access is achieved similarly to the clean random access functionality of an IDR picture.

5 [0082] A CRA picture may have associated RADL or RASL pictures. When a CRA picture is the first picture in the bitstream in decoding order, the CRA picture is the first picture of a coded video sequence in decoding order, and any associated RASL pictures are not output by the decoder and may not be decodable, as they may contain references to pictures that are not present in the bitstream.

[0083] A leading picture is a picture that precedes the associated RAP picture in output order. The associated RAP picture is the previous RAP picture in decoding order (if present). A leading picture may either be a RADL picture or a RASL picture.

[0084] All RASL pictures are leading pictures of an associated BLA or CRA picture. When the associated RAP picture is a BLA picture or is the first coded picture in the bitstream, the RASL picture is not output and may not be correctly decodable, as the RASL picture may contain references to pictures that are not present in the bitstream. However, a RASL picture can be correctly decoded if the decoding had started from a RAP picture before the associated RAP picture of the RASL picture. RASL pictures are not used as reference pictures for the decoding process of non-RASL pictures. When present, all RASL pictures precede, in decoding order, all trailing pictures of the same associated RAP picture. In some drafts of the HEVC standard, a RASL picture was referred to a Tagged for Discard (TFD) picture.

20 [0085] All RADL pictures are leading pictures. RADL pictures are not used as reference pictures for the decoding process of trailing pictures of the same associated RAP picture. When present, all RADL pictures precede, in decoding order, all trailing pictures of the same associated RAP picture. RADL pictures do not refer to any picture preceding the associated RAP picture in decoding order and can therefore be correctly decoded when the decoding starts from the associated RAP picture. In some drafts of the HEVC standard, a RADL picture was referred to a Decodable Leading Picture (DLP).

25 [0086] Decodable leading pictures may be such that can be correctly decoded when the decoding is started from the CRA picture. In other words, decodable leading pictures use only the initial CRA picture or subsequent pictures in decoding order as reference in inter prediction. Non-decodable leading pictures are such that cannot be correctly decoded when the decoding is started from the initial CRA picture. In

other words, non-decodable leading pictures use pictures prior, in decoding order, to the initial CRA picture as references in inter prediction.

[0087] Concatenation of coded video data, which may also be referred to as splicing, may occur for example coded video sequences are concatenated into a bitstream that is broadcast or streamed or stored in a mass memory. For example, coded video sequences representing commercials or advertisements may be concatenated with movies or other “primary” content. A spliced bitstream may be defined as the bitstream that is added after a first bitstream and the concatenation of the first bitstream and the spliced bitstream form a new bitstream. If a spliced bitstream contains only one coded video sequence, it may be referred to as the spliced coded video sequence. Similarly, the first coded video sequence of the spliced bitstream may be referred to as the spliced coded video sequence. An entity performing the concatenation of the bitstreams may be referred to as a splicer.

[0088] When a part of a bitstream starting from a CRA picture is included in another bitstream, the RASL pictures associated with the CRA picture might not be correctly decodable, because some of their reference pictures might not be present in the combined bitstream. To make such a splicing operation straightforward, the NAL unit type of the CRA picture can be changed to indicate that it is a BLA picture. The RASL pictures associated with a BLA picture may not be correctly decodable hence are not be output/displayed. Furthermore, the RASL pictures associated with a BLA picture may be omitted from decoding.

[0089] A BLA picture may be the first picture in the bitstream in decoding order, or may appear later in the bitstream. Each BLA picture begins a new coded video sequence, and has similar effect on the decoding process as an IDR picture. However, a BLA picture contains syntax elements that specify a non-empty reference picture set. When a BLA picture has nal\_unit\_type equal to BLA\_W\_LP, it may have associated RASL pictures, which are not output by the decoder and may not be decodable, as they may contain references to pictures that are not present in the bitstream. When a BLA picture has nal\_unit\_type equal to BLA\_W\_LP, it may also have associated RADL pictures, which are specified to be decoded. When a BLA picture has nal\_unit\_type equal to BLA\_W\_DLP, it does not have associated RASL pictures but may have associated RADL pictures, which are specified to be decoded. BLA\_W\_DLP may also be referred to as BLA\_W\_RADL. When a BLA picture has nal\_unit\_type equal to BLA\_N\_LP, it does not have any associated leading pictures.

[0090] An IDR picture having nal\_unit\_type equal to IDR\_N\_LP does not have associated leading pictures present in the bitstream. An IDR picture having nal\_unit\_type equal to IDR\_W\_DLP does not have associated RASL pictures present in the bitstream, but may have associated RADL pictures in the bitstream. IDR\_W\_DLP may also be referred to as IDR\_W\_RADL.

[0091] When the value of `nal_unit_type` is equal to `TRAIL_N`, `TSA_N`, `STSA_N`, `RADL_N`, `RASL_N`, `RSV_VCL_N10`, `RSV_VCL_N12`, or `RSV_VCL_N14`, the decoded picture is not used as a reference for any other picture of the same temporal sub-layer. That is, in HEVC, when the value of `nal_unit_type` is equal to `TRAIL_N`, `TSA_N`, `STSA_N`, `RADL_N`, `RASL_N`, `RSV_VCL_N10`, `RSV_VCL_N12`, or  
5 `RSV_VCL_N14`, the decoded picture is not included in any of `RefPicSetStCurrBefore`, `RefPicSetStCurrAfter` and `RefPicSetLtCurr` of any picture with the same value of `TemporalId`. A coded picture with `nal_unit_type` equal to `TRAIL_N`, `TSA_N`, `STSA_N`, `RADL_N`, `RASL_N`, `RSV_VCL_N10`, `RSV_VCL_N12`, or `RSV_VCL_N14` may be discarded without affecting the decodability of other pictures with the same value of `TemporalId`.

10 [0092] A trailing picture may be defined as a picture that follows the associated RAP picture in output order. Any picture that is a trailing picture does not have `nal_unit_type` equal to `RADL_N`, `RADL_R`, `RASL_N` or `RASL_R`. Any picture that is a leading picture may be constrained to precede, in decoding order, all trailing pictures that are associated with the same RAP picture. No RASL pictures are present in the bitstream that are associated with a BLA picture having `nal_unit_type` equal to `BLA_W_DLP` or  
15 `BLA_N_LP`. No RADL pictures are present in the bitstream that are associated with a BLA picture having `nal_unit_type` equal to `BLA_N_LP` or that are associated with an IDR picture having `nal_unit_type` equal to `IDR_N_LP`. Any RASL picture associated with a CRA or BLA picture may be constrained to precede any RADL picture associated with the CRA or BLA picture in output order. Any RASL picture associated with a CRA picture may be constrained to follow, in output order, any other RAP picture that precedes the  
20 CRA picture in decoding order.

[0093] In HEVC, there are two picture types, the TSA and STSA picture types, that can be used to indicate temporal sub-layer switching points. If temporal sub-layers with `TemporalId` up to `N` had been decoded until the TSA or STSA picture (exclusive) and the TSA or STSA picture has `TemporalId` equal to `N+1`, the TSA or STSA picture enables decoding of all subsequent pictures (in decoding order) having  
25 `TemporalId` equal to `N+1`. The TSA picture type may impose restrictions on the TSA picture itself and all pictures in the same sub-layer that follow the TSA picture in decoding order. None of these pictures is allowed to use inter prediction from any picture in the same sub-layer that precedes the TSA picture in decoding order. The TSA definition may further impose restrictions on the pictures in higher sub-layers that follow the TSA picture in decoding order. None of these pictures is allowed to refer a picture that  
30 precedes the TSA picture in decoding order if that picture belongs to the same or higher sub-layer as the TSA picture. TSA pictures have `TemporalId` greater than 0. The STSA is similar to the TSA picture but does not impose restrictions on the pictures in higher sub-layers that follow the STSA picture in decoding order and hence enable up-switching only onto the sub-layer where the STSA picture resides.

[0094] A non-VCL NAL unit may be for example one of the following types: a sequence parameter set, a picture parameter set, a supplemental enhancement information (SEI) NAL unit, an access unit delimiter, an end of sequence NAL unit, an end of stream NAL unit, or a filler data NAL unit. Parameter sets may be needed for the reconstruction of decoded pictures, whereas many of the other non-VCL NAL units are not necessary for the reconstruction of decoded sample values.

[0095] Parameters that remain unchanged through a coded video sequence may be included in a sequence parameter set (SPS). In addition to the parameters that may be essential to the decoding process, the sequence parameter set may optionally contain video usability information (VUI), which includes parameters that may be important for buffering, picture output timing, rendering and resource reservation.

There are three NAL units specified in H.264/AVC to carry sequence parameter sets: the sequence parameter set NAL unit containing all the data for H.264/AVC VCL NAL units in the sequence, the sequence parameter set extension NAL unit containing the data for auxiliary coded pictures, and the subset sequence parameter set for MVC and SVC VCL NAL units. A picture parameter set (PPS) contains such parameters that are likely to be unchanged in several coded pictures.

[0096] Parameter set syntax structures may have extensions mechanisms, which may for example be used to include parameters that are specific to extensions of a coding standard. An example syntax of an extension mechanism is provided in the following for SPS:

	<b>Descriptor</b>
seq_parameter_set_rbsp( ) {	
...	
<b>sps_extension_flag</b>	u(1)
if( sps_extension_flag )	
while( more_rbsp_data( ) )	
<b>sps_extension_data_flag</b>	u(1)
rsbsp_trailing_bits( )	
}	

[0097] Decoders of particular version(s) of a coding standard or a coding scheme may ignore sps\_extension\_data\_flag, while in another version of the coding standard or the coding scheme, an extension syntax structure may be specified and may appear within the sps\_extension\_data\_flag bits. Similar extensions mechanisms may be specified also for other types of parameter sets.

[0098] In a draft version of HEVC, there is also a third type of parameter sets, here referred to as Adaptation Parameter Set (APS), which includes parameters that are likely to be unchanged in several coded slices. In a draft version of HEVC, the APS syntax structure includes parameters or syntax elements related to context-based adaptive binary arithmetic coding (CABAC), adaptive sample offset, adaptive loop filtering, and deblocking filtering. In a draft version of HEVC, an APS is a NAL unit and coded without

reference or prediction from any other NAL unit. An identifier, referred to as `aps_id` syntax element, is included in APS NAL unit, and included and used in the slice header to refer to a particular APS. However, APS was not included in the final H.265/HEVC standard.

[0099] H.265/HEVC also includes another type of a parameter set, called a video parameter set (VPS).

5 A video parameter set RBSP may include parameters that can be referred to by one or more sequence parameter set RBSPs.

[0100] The relationship and hierarchy between VPS, SPS, and PPS may be described as follows. VPS resides one level above SPS in the parameter set hierarchy and in the context of scalability and/or 3DV. VPS may include parameters that are common for all slices across all (scalability or view) layers in the entire coded video sequence. SPS includes the parameters that are common for all slices in a particular (scalability or view) layer in the entire coded video sequence, and may be shared by multiple (scalability or view) layers. PPS includes the parameters that are common for all slices in a particular layer representation (the representation of one scalability or view layer in one access unit) and are likely to be shared by all slices in multiple layer representations.

15 [0101] VPS may provide information about the dependency relationships of the layers in a bitstream, as well as many other information that are applicable to all slices across all (scalability or view) layers in the entire coded video sequence. In a scalable extension of HEVC, VPS may for example include a mapping of the `LayerId` value derived from the NAL unit header to one or more scalability dimension values, for example correspond to `dependency_id`, `quality_id`, `view_id`, and `depth_flag` for the layer defined similarly to SVC and MVC. VPS may include profile and level information for one or more layers as well as the profile and/or level for one or more temporal sub-layers (consisting of VCL NAL units at and below certain `TemporalId` values) of a layer representation. VPS may also provide the maximum number of layers present in the bitstream. For example, the syntax element `vps_max_layers_minus1` may be included in the syntax and `vps_max_layer_minus1 + 1` may indicate the maximum number of layers present in the bitstream. The actual number of layers in the bitstream may be smaller than or equal to `vps_max_layer_minus1 + 1`.

[0102] An example syntax of a VPS extension intended to be a part of the VPS is provided in the following. The presented VPS extension provides the dependency relationships among other things. It should be understood that the VPS extension syntax is provided as an example and other similar and/or extended syntax structures may be equivalently applied with different embodiments.

<code>vps_extension( ) {</code>	Descriptor
<code>while( !byte_aligned( ) )</code>	
<code>    <b>vps_extension_byte_alignment_reserved_one_bit</b></code>	<code>u(1)</code>
<code>for( i = 0, NumScalabilityTypes = 0; i &lt; 16; i++ ) {</code>	
<code>    <b>scalability_mask[ i ]</b></code>	<code>u(1)</code>
<code>    NumScalabilityTypes += scalability_mask[ i ]</code>	

}	
for( j = 0; j < NumScalabilityTypes; j++ )	
<b>dimension_id_len_minus1[ j ]</b>	u(3)
<b>vps_nuh_layer_id_present_flag</b>	u(1)
for( i = 0; i <= vps_max_layers_minus1; i++ ) {	
if( vps_nuh_layer_id_present_flag && i > 0 )	
<b>layer_id_in_nuh[ i ]</b>	u(6)
for( j = 0; j < NumScalabilityTypes; j++ )	
<b>dimension_id[ i ][ j ]</b>	u(v)
}	
for( i = 1; i <= vps_max_layers_minus1; i++ )	
for( j = 0; j < i; j++ )	
<b>direct_dependency_flag[ i ][ j ]</b>	u(1)
<b>direct_dep_type_len_minus2</b>	ue(v)
for( i = 1; i <= vps_max_layers_minus1; i++ )	
for( j = 0; j < i; j++ )	
if( direct_dependency_flag[ i ][ j ] )	
<b>direct_dependency_type[ i ][ j ]</b>	u(v)
}	

[0103] The semantics of the presented VPS extension may be specified as described in the following paragraphs.

[0104] vps\_extension\_byte\_alignment\_reserved\_one\_bit is equal to 1 and is used to achieve alignment of the next syntax element to a byte boundary.

[0105] scalability\_mask[i] equal to 1 indicates that dimension\_id syntax elements corresponding to the i-th scalability dimension are present. scalability\_mask[i] equal to 0 indicates that dimension\_id syntax elements corresponding to the i-th scalability dimension are not present. The scalability dimensions corresponding to each value of of i in scalability\_mask[i] may be specified for example to include the following or any subset thereof along with other scalability dimensions.

scalability_mask index	Scalability dimension	ScalabilityId mapping
0	multiview	ViewId
1	spatial or quality scalability	DependencyId

[0106] dimension\_id\_len\_minus1[ j ] plus 1 specifies the length, in bits, of the dimension\_id[i][ j ] syntax element. vps\_nuh\_layer\_id\_present\_flag specifies whether the layer\_id\_in\_nuh[i] syntax is present.

layer\_id\_in\_nuh[i] specifies the value of the nuh\_layer\_id syntax element in VCL NAL units of the i-th layer. When not present, the value of layer\_id\_in\_nuh[i] is inferred to be equal to i. layer\_id\_in\_nuh[i] is greater than layer\_id\_in\_nuh[i - 1]. The variable LayerIdxInVps[layer\_id\_in\_nuh[i]] is set equal to i. dimension\_id[i][j] specifies the identifier of the j-th scalability dimension type of the i-th layer. When not present, the value of dimension\_id[i][j] is inferred to be equal to 0. The number of bits used for the representation of dimension\_id[i][j] is dimension\_id\_len\_minus1[j] + 1 bits.

[0107] direct\_dependency\_flag[i][j] equal to 0 specifies that the layer with index j is not a direct reference layer for the layer with index i. direct\_dependency\_flag[i][j] equal to 1 specifies that the layer with index j may be a direct reference layer for the layer with index i. When direct\_dependency\_flag[i][j] is not present for i and j in the range of 0 to vps\_max\_num\_layers\_minus1, it is inferred to be equal to 0.

[0108] The variables NumDirectRefLayers[i] and RefLayerId[i][j] may be derived as follows:

```

for( i = 1; i <= vps_max_layers_minus1; i++ )
  for( j = 0, NumDirectRefLayers[ i ] = 0; j < i; j++ )
    if( direct_dependency_flag[ i ][ j ] == 1 )
      RefLayerId[ i ][ NumDirectRefLayers[ i ]++ ] = layer_id_in_nuh[ j ]

```

[0109] direct\_dep\_type\_len\_minus2 plus 2 specifies the number of bits of the direct\_dependency\_type[i][j] syntax element. direct\_dependency\_type[i][j] equal to 0 indicates that sample prediction may be used and motion prediction is not used for layer identified by i from layer identified by j. direct\_dependency\_type[i][j] equal to 1 indicates that motion prediction may be used and sample prediction is not used for layer identified by i from layer identified by j. direct\_dependency\_type[i][j] equal to 2 indicates that both sample and motion prediction may be used for layer identified by i from layer identified by j.

[0110] The variables NumSamplePredRefLayers[i], NumMotionPredRefLayers[i], SamplePredEnabledFlag[i][j], MotionPredEnabledFlag[i][j], NumDirectRefLayers[i], RefLayerId[i][j], MotionPredRefLayerId[i][j], and SamplePredRefLayerId[i][j] may be derived as follows:

```

for( i = 0; i < 64; i++ ) {
  NumSamplePredRefLayers[ i ] = 0
  NumMotionPredRefLayers[ i ] = 0
  NumDirectRefLayers[ i ] = 0
  for( j = 0; j < 64; j++ ) {
    SamplePredEnabledFlag[ i ][ j ] = 0
    MotionPredEnabledFlag[ i ][ j ] = 0
    RefLayerId[ i ][ j ] = 0
    SamplePredRefLayerId[ i ][ j ] = 0
    MotionPredRefLayerId[ i ][ j ] = 0
  }
}

```

```

for( i = 1; i <= vps_max_layers_minus1; i++ ) {
  iNuhLid = layer_id_in_nuh[ i ]
  for( j = 0; j < i; j++ )
    if( direct_dependency_flag[ i ][ j ] ) {
5      RefLayerId[ iNuhLid ][ NumDirectRefLayers[ iNuhLid ]++ ] =
      layer_id_in_nuh[ j ]
      SamplePredEnabledFlag[ iNuhLid ][ j ] = ( ( direct_dependency_type[ i ][ j ] + 1 ) & 1 )
      NumSamplePredRefLayers[ iNuhLid ] += SamplePredEnabledFlag[ iNuhLid ][ j ]
      MotionPredEnabledFlag[ iNuhLid ][ j ] = ( ( ( direct_dependency_type[ i ][ j ] + 1 )
10      & 2 ) >> 1 )
      NumMotionPredRefLayers[ iNuhLid ] += MotionPredEnabledFlag[ iNuhLid ][ j ]
    }
  }

for( i = 1, mIdx = 0, sIdx = 0; i <= vps_max_layers_minus1; i++ ) {
15  iNuhLid = layer_id_in_nuh[ i ]
  for( j = 0, j < i; j++ ) {
    if( MotionPredEnabledFlag[ iNuhLid ][ j ] )
      MotionPredRefLayerId[ iNuhLid ][ mIdx++ ] = layer_id_in_nuh[ j ]
    if( SamplePredEnabledFlag[ iNuhLid ][ j ] )
20      SamplePredRefLayerId[ iNuhLid ][ sIdx++ ] = layer_id_in_nuh[ j ]
  }
}

```

[0111] H.264/AVC and HEVC syntax allows many instances of parameter sets, and each instance is  
25 identified with a unique identifier. In H.264/AVC, each slice header includes the identifier of the picture  
parameter set that is active for the decoding of the picture that contains the slice, and each picture parameter  
set contains the identifier of the active sequence parameter set. Consequently, the transmission of picture  
and sequence parameter set does not have to be accurately synchronized with the transmission of slices.  
Instead, it is sufficient that the active sequence and picture parameter sets are received at any moment before  
30 they are referenced, which allows transmission of parameter sets “out-of-band” using a more reliable  
transmission mechanism compared to the protocols used for the slice data. For example, parameter sets can  
be included as a parameter in the session description for Real-time Transport Protocol (RTP) sessions. If  
parameter sets are transmitted in-band, they can be repeated to improve error robustness.

[0112] A SEI NAL unit may contain one or more SEI message, which are not required for the decoding  
35 of output pictures but assist in related processes, such as picture output timing, rendering, error detection,  
error concealment, and resource reservation. Several SEI messages are specified in H.264/AVC and HEVC,  
and the user data SEI messages enable organizations and companies to specify SEI messages for their own  
use. H.264/AVC and HEVC contain the syntax and semantics for the specified SEI messages but no process  
for handling the messages in the recipient is defined. Consequently, encoders are required to follow the  
40 H.264/AVC standard or the HEVC standard when they create SEI messages, and decoders conforming to  
the H.264/AVC standard or the HEVC standard, respectively, are not required to process SEI messages for

output order conformance. One of the reasons to include the syntax and semantics of SEI messages in H.264/AVC and HEVC is to allow different system specifications to interpret the supplemental information identically and hence interoperate. It is intended that system specifications can require the use of particular SEI messages both in the encoding end and in the decoding end, and additionally the process for handling particular SEI messages in the recipient can be specified.

[0113] Several nesting SEI messages have been specified in the AVC and HEVC standards or proposed otherwise. The idea of nesting SEI messages is to contain one or more SEI messages within a nesting SEI message and provide a mechanism for associating the contained SEI messages with a subsets of the bitstream and/or a subset of decoded data. It may be required that a nesting SEI message contains one or more SEI messages that are not nesting SEI messages themselves. An SEI message contained in a nesting SEI message may be referred to as a nested SEI message. An SEI message not contained in a nesting SEI message may be referred to as a non-nested SEI message. The scalable nesting SEI message of HEVC enables to identify either a bitstream subset (resulting from a sub-bitstream extraction process) or a set of layers to which the nested SEI messages apply. A bitstream subset may also be referred to as a sub-bitstream.

[0114] A coded picture is a coded representation of a picture. A coded picture in H.264/AVC comprises the VCL NAL units that are required for the decoding of the picture. In H.264/AVC, a coded picture can be a primary coded picture or a redundant coded picture. A primary coded picture is used in the decoding process of valid bitstreams, whereas a redundant coded picture is a redundant representation that should only be decoded when the primary coded picture cannot be successfully decoded. In HEVC, no redundant coded picture has been specified.

[0115] In H.264/AVC and HEVC, an access unit comprises a primary coded picture and those NAL units that are associated with it. In HEVC, an access unit is defined as a set of NAL units that are associated with each other according to a specified classification rule, are consecutive in decoding order, and contain exactly one coded picture. In H.264/AVC, the appearance order of NAL units within an access unit is constrained as follows. An optional access unit delimiter NAL unit may indicate the start of an access unit. It is followed by zero or more SEI NAL units. The coded slices of the primary coded picture appear next. In H.264/AVC, the coded slice of the primary coded picture may be followed by coded slices for zero or more redundant coded pictures. A redundant coded picture is a coded representation of a picture or a part of a picture. A redundant coded picture may be decoded if the primary coded picture is not received by the decoder for example due to the loss in transmission or a corruption in physical storage medium.

[0116] In H.264/AVC, an access unit may also include an auxiliary coded picture, which is a picture that supplements the primary coded picture and may be used for example in the display process. An auxiliary coded picture may for example be used as an alpha channel or alpha plane specifying the

transparency level of the samples in the decoded pictures. An alpha channel or plane may be used in a layered composition or rendering system, where the output picture is formed by overlaying pictures being at least partly transparent on top of each other. An auxiliary coded picture has the same syntactic and semantic restrictions as a monochrome redundant coded picture. In H.264/AVC, an auxiliary coded picture  
5 contains the same number of macroblocks as the primary coded picture.

[0117] In HEVC, an access unit may be defined as a set of NAL units that are associated with each other according to a specified classification rule, are consecutive in decoding order, and contain exactly one coded picture. In addition to containing the VCL NAL units of the coded picture, an access unit may also contain non-VCL-NAL units. In HEVC, the decoding of an access unit results in a decoded picture.

10 [0118] In H.264/AVC, a coded video sequence is defined to be a sequence of consecutive access units in decoding order from an IDR access unit, inclusive, to the next IDR access unit, exclusive, or to the end of the bitstream, whichever appears earlier. In HEVC, a coded video sequence is defined to be a sequence of access units that consists, in decoding order, of a CRA (Clean Random Access) access unit that is the first access unit in the bitstream, and IDR access unit or a BLA (Broken Link Access) access unit, followed  
15 by zero or more non-IDR and non-BLA access units including all subsequent access units up to but no including any subsequent IDR or BLA access unit.

[0119] A group of pictures (GOP) and its characteristics may be defined as follows. A GOP can be decoded regardless of whether any previous pictures were decoded. An open GOP is such a group of pictures in which pictures preceding the initial intra picture in output order might not be correctly decodable  
20 when the decoding starts from the initial intra picture of the open GOP. In other words, pictures of an open GOP may refer (in inter prediction) to pictures belonging to a previous GOP. An H.264/AVC decoder can recognize an intra picture starting an open GOP from the recovery point SEI message in an H.264/AVC bitstream. An HEVC decoder can recognize an intra picture starting an open GOP, because a specific NAL unit type, CDR NAL unit type, is used for its coded slices. A closed GOP is such a group of pictures in  
25 which all pictures can be correctly decoded when the decoding starts from the initial intra picture of the closed GOP. In other words, no picture in a closed GOP refers to any pictures in previous GOPs. In H.264/AVC and HEVC, a closed GOP starts from an IDR access unit. In HEVC a closed GOP may also start from BLA W DLP or a BLA N LP picture. As a result, closed GOP structure has more error resilience potential in comparison to the open GOP structure, however at the cost of possible reduction in the  
30 compression efficiency. Open GOP coding structure is potentially more efficient in the compression, due to a larger flexibility in selection of reference pictures.

[0120] The bitstream syntax of H.264/AVC and HEVC indicates whether a particular picture is a reference picture for inter prediction of any other picture. Pictures of any coding type (I, P, B) can be reference pictures or non-reference pictures in H.264/AVC and HEVC. The NAL unit header indicates the

type of the NAL unit and whether a coded slice contained in the NAL unit is a part of a reference picture or a non-reference picture.

[0121] Many hybrid video codecs, including H.264/AVC and HEVC, encode video information in two phases. In the first phase, predictive coding is applied for example as so-called sample prediction and/or so-called syntax prediction.

[0122] In the sample prediction, pixel or sample values in a certain picture area or “block” are predicted. These pixel or sample values can be predicted, for example, using one or more of the following ways:

- Motion compensation mechanisms (which may also be referred to as temporal prediction or motion-compensated temporal prediction or motion-compensated prediction or MCP), which involve finding and indicating an area in one of the previously encoded video frames that corresponds closely to the block being coded.
- Inter-view prediction, which involves finding and indicating an area in one of the previously encoded view components that corresponds closely to the block being coded.
- View synthesis prediction, which involves synthesizing a prediction block or image area where a prediction block is derived on the basis of reconstructed/decoded ranging information.
- Inter-layer prediction using reconstructed/decoded samples, such as the so-called IntraBL (base layer) mode of SVC.
- Inter-layer residual prediction, in which for example the coded residual of a reference layer or a derived residual from a difference of a reconstructed/decoded reference layer picture and a corresponding reconstructed/decoded enhancement layer picture may be used for predicting a residual block of the current enhancement layer block. A residual block may be added for example to a motion-compensated prediction block to obtain a final prediction block for the current enhancement layer block.
- Intra prediction, where pixel or sample values can be predicted by spatial mechanisms which involve finding and indicating a spatial region relationship.

[0123] In the syntax prediction, which may also be referred to as parameter prediction, syntax elements and/or syntax element values and/or variables derived from syntax elements are predicted from syntax elements (de)coded earlier and/or variables derived earlier. Non-limiting examples of syntax prediction are provided below:

- In motion vector prediction, motion vectors e.g. for inter and/or inter-view prediction may be coded differentially with respect to a block-specific predicted motion vector. In many video codecs, the predicted motion vectors are created in a predefined way, for example by calculating the median of the encoded or decoded motion vectors of the adjacent blocks. Another way to create motion vector predictions, sometimes referred to as advanced motion vector prediction (AMVP), is to generate a

list of candidate predictions from adjacent blocks and/or co-located blocks in temporal reference pictures and signalling the chosen candidate as the motion vector predictor. In addition to predicting the motion vector values, the reference index of previously coded/decoded picture can be predicted. The reference index is typically predicted from adjacent blocks and/or co-located blocks in temporal reference picture. Differential coding of motion vectors is typically disabled across slice boundaries.

- The block partitioning, e.g. from CTU to CUs and down to PUs, may be predicted.
- In filter parameter prediction, the filtering parameters e.g. for sample adaptive offset may be predicted.

10 [0124] Prediction approaches using image information from a previously coded image can also be called as inter prediction methods which may also be referred to as temporal prediction and motion compensation. Prediction approaches using image information within the same image can also be called as intra prediction methods.

15 [0125] The second phase is one of coding the error between the predicted block of pixels or samples and the original block of pixels or samples. This may be accomplished by transforming the difference in pixel or sample values using a specified transform. This transform may be e.g. a Discrete Cosine Transform (DCT) or a variant thereof. After transforming the difference, the transformed difference is quantized and entropy coded.

20 [0126] By varying the fidelity of the quantization process, the encoder can control the balance between the accuracy of the pixel or sample representation (i.e. the visual quality of the picture) and the size of the resulting encoded video representation (i.e. the file size or transmission bit rate).

25 [0127] The decoder reconstructs the output video by applying a prediction mechanism similar to that used by the encoder in order to form a predicted representation of the pixel or sample blocks (using the motion or spatial information created by the encoder and included in the compressed representation of the image) and prediction error decoding (the inverse operation of the prediction error coding to recover the quantized prediction error signal in the spatial domain).

[0128] After applying pixel or sample prediction and error decoding processes the decoder combines the prediction and the prediction error signals (the pixel or sample values) to form the output video frame.

30 [0129] The decoder (and encoder) may also apply additional filtering processes in order to improve the quality of the output video before passing it for display and/or storing as a prediction reference for the forthcoming pictures in the video sequence.

[0130] In many video codecs, including H.264/AVC and HEVC, motion information is indicated by motion vectors associated with each motion compensated image block. Each of these motion vectors represents the displacement of the image block in the picture to be coded (in the encoder) or decoded (at

the decoder) and the prediction source block in one of the previously coded or decoded images (or picture). H.264/AVC and HEVC, as many other video compression standards, divide a picture into a mesh of rectangles, for each of which a similar block in one of the reference pictures is indicated for inter prediction. The location of the prediction block is coded as a motion vector that indicates the position of the prediction block relative to the block being coded.

[0131] H.264/AVC and HEVC include a concept of picture order count (POC). A value of POC is derived for each picture and is non-decreasing with increasing picture position in output order. POC therefore indicates the output order of pictures. POC may be used in the decoding process for example for implicit scaling of motion vectors in the temporal direct mode of bi-predictive slices, for implicitly derived weights in weighted prediction, and for reference picture list initialization. Furthermore, POC may be used in the verification of output order conformance. In H.264/AVC, POC is specified relative to the previous IDR picture or a picture containing a memory management control operation marking all pictures as “unused for reference”.

[0132] Inter prediction process may be characterized using one or more of the following factors.

[0133] The accuracy of motion vector representation. For example, motion vectors may be of quarter-pixel accuracy, and sample values in fractional-pixel positions may be obtained using a finite impulse response (FIR) filter.

[0134] Block partitioning for inter prediction. Many coding standards, including H.264/AVC and HEVC, allow selection of the size and shape of the block for which a motion vector is applied for motion-compensated prediction in the encoder, and indicating the selected size and shape in the bitstream so that decoders can reproduce the motion-compensated prediction done in the encoder.

[0135] Number of reference pictures for inter prediction. The sources of inter prediction are previously decoded pictures. Many coding standards, including H.264/AVC and HEVC, enable storage of multiple reference pictures for inter prediction and selection of the used reference picture on a block basis. For example, reference pictures may be selected on macroblock or macroblock partition basis in H.264/AVC and on PU or CU basis in HEVC. Many coding standards, such as H.264/AVC and HEVC, include syntax structures in the bitstream that enable decoders to create one or more reference picture lists. A reference picture index to a reference picture list may be used to indicate which one of the multiple reference pictures is used for inter prediction for a particular block. A reference picture index may be coded by an encoder into the bitstream in some inter coding modes or it may be derived (by an encoder and a decoder) for example using neighboring blocks in some other inter coding modes.

[0136] Motion vector prediction. In order to represent motion vectors efficiently in bitstreams, motion vectors may be coded differentially with respect to a block-specific predicted motion vector. In many video codecs, the predicted motion vectors are created in a predefined way, for example by calculating the median

of the encoded or decoded motion vectors of the adjacent blocks. Another way to create motion vector predictions, sometimes referred to as advanced motion vector prediction (AMVP), is to generate a list of candidate predictions from adjacent blocks and/or co-located blocks in temporal reference pictures and signalling the chosen candidate as the motion vector predictor. In addition to predicting the motion vector values, the reference index of previously coded/decoded picture can be predicted. The reference index may be predicted e.g. from adjacent blocks and/or co-located blocks in temporal reference picture. Differential coding of motion vectors may be disabled across slice boundaries.

[0137] Multi-hypothesis motion-compensated prediction. H.264/AVC and HEVC enable the use of a single prediction block in P slices (herein referred to as uni-predictive slices) or a linear combination of two motion-compensated prediction blocks for bi-predictive slices, which are also referred to as B slices. Individual blocks in B slices may be bi-predicted, uni-predicted, or intra-predicted, and individual blocks in P slices may be uni-predicted or intra-predicted. The reference pictures for a bi-predictive picture may not be limited to be the subsequent picture and the previous picture in output order, but rather any reference pictures may be used. In many coding standards, such as H.264/AVC and HEVC, one reference picture list, referred to as reference picture list 0, is constructed for P slices, and two reference picture lists, list 0 and list 1, are constructed for B slices. For B slices, when prediction in forward direction may refer to prediction from a reference picture in reference picture list 0, and prediction in backward direction may refer to prediction from a reference picture in reference picture list 1, even though the reference pictures for prediction may have any decoding or output order relation to each other or to the current picture.

[0138] Weighted prediction. Many coding standards use a prediction weight of 1 for prediction blocks of inter (P) pictures and 0.5 for each prediction block of a B picture (resulting into averaging). H.264/AVC allows weighted prediction for both P and B slices. In implicit weighted prediction, the weights are proportional to picture order counts (POC), while in explicit weighted prediction, prediction weights are explicitly indicated.

[0139] In many video codecs, the prediction residual after motion compensation is first transformed with a transform kernel (like DCT) and then coded. The reason for this is that often there still exists some correlation among the residual and transform can in many cases help reduce this correlation and provide more efficient coding.

[0140] In HEVC, each PU has prediction information associated with it defining what kind of a prediction is to be applied for the pixels within that PU (e.g. motion vector information for inter predicted PUs and intra prediction directionality information for intra predicted PUs). Similarly each TU is associated with information describing the prediction error decoding process for the samples within the TU (including e.g. DCT coefficient information). It may be signaled at CU level whether prediction error coding is applied

or not for each CU. In the case there is no prediction error residual associated with the CU, it can be considered there are no TUs for the CU.

[0141] In some coding formats and codecs, a distinction is made between so-called short-term and long-term reference pictures. This distinction may affect some decoding processes such as motion vector scaling in the temporal direct mode or implicit weighted prediction. If both of the reference pictures used for the temporal direct mode are short-term reference pictures, the motion vector used in the prediction may be scaled according to the picture order count difference between the current picture and each of the reference pictures. However, if at least one reference picture for the temporal direct mode is a long-term reference picture, default scaling of the motion vector may be used, for example scaling the motion to half may be used. Similarly, if a short-term reference picture is used for implicit weighted prediction, the prediction weight may be scaled according to the POC difference between the POC of the current picture and the POC of the reference picture. However, if a long-term reference picture is used for implicit weighted prediction, a default prediction weight may be used, such as 0.5 in implicit weighted prediction for bi-predicted blocks.

[0142] Some video coding formats, such as H.264/AVC, include the frame\_num syntax element, which is used for various decoding processes related to multiple reference pictures. In H.264/AVC, the value of frame\_num for IDR pictures is 0. The value of frame\_num for non-IDR pictures is equal to the frame\_num of the previous reference picture in decoding order incremented by 1 (in modulo arithmetic, i.e., the value of frame\_num wrap over to 0 after a maximum value of frame\_num).

[0143] A syntax structure for (decoded) reference picture marking may exist in a video coding system. For example, when the decoding of the picture has been completed, the decoded reference picture marking syntax structure, if present, may be used to adaptively mark pictures as “unused for reference” or “used for long-term reference”. If the decoded reference picture marking syntax structure is not present and the number of pictures marked as “used for reference” can no longer increase, a sliding window reference picture marking may be used, which basically marks the earliest (in decoding order) decoded reference picture as unused for reference.

[0144] H.264/AVC specifies the process for decoded reference picture marking in order to control the memory consumption in the decoder. The maximum number of reference pictures used for inter prediction, referred to as M, is determined in the sequence parameter set. When a reference picture is decoded, it is marked as “used for reference”. If the decoding of the reference picture caused more than M pictures marked as “used for reference”, at least one picture is marked as “unused for reference”. There are two types of operation for decoded reference picture marking: adaptive memory control and sliding window. The operation mode for decoded reference picture marking is selected on picture basis. The adaptive memory control enables explicit signaling which pictures are marked as “unused for reference” and may also assign long-term indices to short-term reference pictures. The adaptive memory control may require the presence

of memory management control operation (MMCO) parameters in the bitstream. MMCO parameters may be included in a decoded reference picture marking syntax structure. If the sliding window operation mode is in use and there are M pictures marked as “used for reference”, the short-term reference picture that was the first decoded picture among those short-term reference pictures that are marked as “used for reference” is marked as “unused for reference”. In other words, the sliding window operation mode results into first-in-first-out buffering operation among short-term reference pictures.

[0145] One of the memory management control operations in H.264/AVC causes all reference pictures except for the current picture to be marked as “unused for reference”. An instantaneous decoding refresh (IDR) picture contains only intra-coded slices and causes a similar “reset” of reference pictures.

[0146] In HEVC, reference picture marking syntax structures and related decoding processes have been replaced with a reference picture set (RPS) syntax structure and decoding process are used instead for a similar purpose. A reference picture set valid or active for a picture includes all the reference pictures used as reference for the picture and all the reference pictures that are kept marked as “used for reference” for any subsequent pictures in decoding order. There are six subsets of the reference picture set, which are referred to as RefPicSetStCurr0 (a.k.a. RefPicSetStCurrBefore), RefPicSetStCurr1 (a.k.a. RefPicSetStCurrAfter), RefPicSetStFoll0, RefPicSetStFoll1, RefPicSetLtCurr, and RefPicSetLtFoll. RefPicSetStFoll0 and RefPicSetStFoll1 may also be considered to form jointly one subset RefPicSetStFoll. The notation of the six subsets is as follows. “Curr” refers to the reference pictures that are included in the reference picture lists of the current picture and hence may be used as inter prediction reference for the current picture. “Foll” refers to reference pictures that are not included in the reference picture lists of the current picture but may be used in subsequent pictures in decoding order as reference pictures. “St” refers to short-term reference pictures, which may generally be identified through a certain number of least significant bits of their POC value. “Lt” refers to long-term reference pictures, which are specifically identified and generally have a greater difference of POC values relative to the current picture than what can be represented by the mentioned certain number of least significant bits. “0” refers to those reference pictures that have a smaller POC value than that of the current picture. “1” refers to those reference pictures that have a greater POC value than that of the current picture. RefPicSetStCurr0, RefPicSetStCurr1, RefPicSetStFoll0 and RefPicSetStFoll1 are collectively referred to as the short-term subset of the reference picture set. RefPicSetLtCurr and RefPicSetLtFoll are collectively referred to as the long-term subset of the reference picture set.

[0147] In HEVC, a reference picture set may be specified in a picture parameter set and taken into use in the slice header through an index to the reference picture set. A reference picture set may also be specified in a slice header. A long-term subset of a reference picture set is generally specified only in a slice header, while the short-term subsets of the same reference picture set may be specified in the picture parameter set

or slice header. A reference picture set may be coded independently or may be predicted from another reference picture set (known as inter-RPS prediction). When a reference picture set is independently coded, the syntax structure includes up to three loops iterating over different types of reference pictures; short-term reference pictures with lower POC value than the current picture, short-term reference pictures with higher POC value than the current picture, and long-term reference pictures. Each loop entry specifies a picture to be marked as “used for reference”. In general, the picture is specified with a differential POC value. The inter-RPS prediction exploits the fact that the reference picture set of the current picture can be predicted from the reference picture set of a previously decoded picture. This is because all the reference pictures of the current picture are either reference pictures of the previous picture or the previously decoded picture itself. It is only necessary to indicate which of these pictures should be reference pictures and be used for the prediction of the current picture. In both types of reference picture set coding, a flag (used\_by\_curr\_pic\_X\_flag) is additionally sent for each reference picture indicating whether the reference picture is used for reference by the current picture (included in a \*Curr list) or not (included in a \*Foll list). Pictures that are included in the reference picture set used by the current slice are marked as “used for reference”, and pictures that are not in the reference picture set used by the current slice are marked as “unused for reference”. If the current picture is an IDR picture, RefPicSetStCurr0, RefPicSetStCurr1, RefPicSetStFoll0, RefPicSetStFoll1, RefPicSetLtCurr, and RefPicSetLtFoll are all set to empty.

[0148] A Decoded Picture Buffer (DPB) may be used in the encoder and/or in the decoder. There are two reasons to buffer decoded pictures, for references in inter prediction and for reordering decoded pictures into output order. As H.264/AVC and HEVC provide a great deal of flexibility for both reference picture marking and output reordering, separate buffers for reference picture buffering and output picture buffering may waste memory resources. Hence, the DPB may include a unified decoded picture buffering process for reference pictures and output reordering. A decoded picture may be removed from the DPB when it is no longer used as a reference and is not needed for output.

[0149] In many coding modes of H.264/AVC and HEVC, the reference picture for inter prediction is indicated with an index to a reference picture list. The index may be coded with CABAC or variable length coding. In general, the smaller the index is, the shorter the corresponding syntax element may become. In H.264/AVC and HEVC, two reference picture lists (reference picture list 0 and reference picture list 1) are generated for each bi-predictive (B) slice, and one reference picture list (reference picture list 0) is formed for each inter-coded (P) slice. In addition, for a B slice in a draft version of the HEVC standard, a combined list (List C) may be constructed after the final reference picture lists (List 0 and List 1) have been constructed. The combined list may be used for uni-prediction (also known as uni-directional prediction) within B slices. However, in the final H.265/HEVC standard, no combined list is constructed.

[0150] A reference picture list, such as the reference picture list 0 and the reference picture list 1, may be constructed in two steps: First, an initial reference picture list is generated. The initial reference picture list may be generated for example on the basis of frame\_num, POC, temporal\_id, or information on the prediction hierarchy such as a GOP structure, or any combination thereof. Second, the initial reference picture list may be reordered by reference picture list reordering (RPLR) commands, also known as reference picture list modification syntax structure, which may be contained in slice headers. The RPLR commands indicate the pictures that are ordered to the beginning of the respective reference picture list. This second step may also be referred to as the reference picture list modification process, and the RPLR commands may be included in a reference picture list modification syntax structure. If reference picture sets are used, the reference picture list 0 may be initialized to contain RefPicSetStCurr0 first, followed by RefPicSetStCurr1, followed by RefPicSetLtCurr. Reference picture list 1 may be initialized to contain RefPicSetStCurr1 first, followed by RefPicSetStCurr0. The initial reference picture lists may be modified through the reference picture list modification syntax structure, where pictures in the initial reference picture lists may be identified through an entry index to the list.

[0151] Since multiview video provides encoders and decoders the possibility to utilize inter-view redundancy, decoded inter-view frames may be included in the reference picture list(s) as well.

[0152] Examples of motion vector prediction schemes and related coding modes are provided in the next paragraphs.

[0153] In addition to the motion-compensated macroblock modes for which a differential motion vector is coded, a P macroblock may also be coded in the so-called P\_Skip type in H.264/AVC. For this coding type, no differential motion vector, reference index, or quantized prediction error signal is coded into the bitstream. The reference picture of a macroblock coded with the P\_Skip type has index 0 in reference picture list 0. The motion vector used for reconstructing the P\_Skip macroblock is obtained using median motion vector prediction for the macroblock without any differential motion vector being added. P\_Skip may be beneficial for compression efficiency particularly in areas where the motion field is smooth.

[0154] In B slices of H.264/AVC, four different types of inter prediction are supported: uni-predictive from reference picture list 0, uni-directional from reference picture list 1, bi-predictive, direct prediction, and B\_skip. The type of inter prediction can be selected separately for each macroblock partition. B slices utilize a similar macroblock partitioning as P slices. For a bi-predictive macroblock partition, the prediction signal is formed by a weighted average of motion-compensated list 0 and list 1 prediction signals. Reference indices, motion vector differences, as well as quantized prediction error signal may be coded for uni-predictive and bi-predictive B macroblock partitions.

[0155] Two direct modes are included in H.264/AVC, temporal direct and spatial direct, and one of them can be selected into use for a slice in a slice header, although their use may be constrained further for

example in profiles or alike. In the temporal direct mode, the reference index for reference picture list 1 is set to 0 and the reference index for reference picture list 0 is set to point to the reference picture that is used in the co-located block (compared to the current block cb) of the reference picture having index 0 in the reference picture list 1 if that reference picture is available, or set to 0 if that reference picture is not available. The motion vector predictor for cb is essentially derived by considering the motion information within a co-located block of the reference picture having index 0 in reference picture list 1. Motion vector predictors for a temporal direct block are derived by scaling a motion vector from the co-located block. The scaling is proportional to picture order count differences between the current picture and the reference pictures associated with the inferred reference indexes in list 0 and list 1, and by selecting the sign for the motion vector predictor depending on which reference picture list it is using.

[0156] In spatial direct mode of H.264/AVC, motion information of spatially adjacent blocks is exploited. Motion vector prediction in spatial direct mode can be divided into three steps: reference index determination, determination of uni- or bi-prediction, and motion vector prediction. In the first step, the reference picture with the minimum non-negative reference index (i.e., non-intra block) is selected from each of reference picture list 0 and reference picture list 1 of the neighboring blocks A, B, and C (where A is the adjacent block on the left of the current block, B is the adjacent block above the current block and C is the adjacent block on the top-right side of the current block). If no non-negative reference index exists in reference picture list 0 of the neighboring blocks A, B, and C, and likewise no non-negative reference index exists in reference picture list 1 of the neighboring blocks A, B, and C, reference index 0 is selected for both reference picture lists.

[0157] The use of uni- or bi-prediction for H.264/AVC spatial direct mode is determined as follows: If a minimum non-negative reference index for both reference picture lists was found in the reference index determination step, bi-prediction is used. If a minimum non-negative reference index for either but not both of reference picture list 0 or reference picture list 1 was found in the reference index determination step, uni-prediction from either reference picture list 0 or reference picture list 1, respectively, is used.

[0158] In the motion vector prediction for H.264/AVC spatial direct mode, certain conditions, such as whether a negative reference index was concluded in the first step, are checked and, if fulfilled, a zero motion vector is determined. Otherwise, the motion vector predictor is derived similarly to the motion vector predictor of P blocks using the motion vectors of spatially adjacent blocks A, B, and C.

[0159] No motion vector differences or reference indices are present in the bitstream for a direct mode block in H.264/AVC, while quantized prediction error signal may be coded and present therefore present in the bitstream. A B\_skip macroblock mode in H.264/AVC is similar to the direct mode but no prediction error signal is coded and included in the bitstream.

[0160] H.265/HEVC includes two motion vector prediction schemes, namely the advanced motion vector prediction (AMVP) and the merge mode. In the AMVP or the merge mode, a list of motion vector candidates is derived for a PU. There are two kinds of candidates: spatial candidates and temporal candidates, where temporal candidates may also be referred to as TMVP candidates.

5 [0161] A candidate list derivation may be performed for example as follows, while it should be understood that other possibilities exist for candidate list derivation. If the occupancy of the candidate list is not at maximum, the spatial candidates are included in the candidate list first if they are available and not already exist in the candidate list. After that, if occupancy of the candidate list is not yet at maximum, a temporal candidate is included in the candidate list. If the number of candidates still does not reach the  
10 maximum allowed number, the combined bi-predictive candidates (for B slices) and a zero motion vector are added in. After the candidate list has been constructed, the encoder decides the final motion information from candidates for example based on a rate-distortion optimization (RDO) decision and encodes the index of the selected candidate into the bitstream. Likewise, the decoder decodes the index of the selected candidate from the bitstream, constructs the candidate list, and uses the decoded index to select a motion  
15 vector predictor from the candidate list.

[0162] In H.265/HEVC, AMVP and the merge mode may be characterized as follows. In AMVP, the encoder indicates whether uni-prediction or bi-prediction is used and which reference pictures are used as well as encodes a motion vector difference. In the merge mode, only the chosen candidate from the candidate list is encoded into the bitstream indicating the current prediction unit has the same motion  
20 information as that of the indicated predictor. Thus, the merge mode creates regions composed of neighboring prediction blocks sharing identical motion information, which is only signaled once for each region. Another difference between AMVP and the merge mode in H.265/HEVC is that the maximum number of candidates of AMVP is 2 while that of the merge mode is 5.

[0163] The advanced motion vector prediction may operate for example as follows, while other similar  
25 realizations of advanced motion vector prediction are also possible for example with different candidate position sets and candidate locations with candidate position sets. Two spatial motion vector predictors (MVPs) may be derived and a temporal motion vector predictor (TMVP) may be derived. They may be selected among the positions: three spatial motion vector predictor candidate positions located above the current prediction block (B0, B1, B2) and two on the left (A0, A1). The first motion vector predictor that  
30 is available (e.g. resides in the same slice, is inter-coded, etc.) in a pre-defined order of each candidate position set, (B0, B1, B2) or (A0, A1), may be selected to represent that prediction direction (up or left) in the motion vector competition. A reference index for the temporal motion vector predictor may be indicated by the encoder in the slice header (e.g. as a `collocated_ref_idx` syntax element). The motion vector obtained from the co-located picture may be scaled according to the proportions of the picture order count differences

of the reference picture of the temporal motion vector predictor, the co-located picture, and the current picture. Moreover, a redundancy check may be performed among the candidates to remove identical candidates, which can lead to the inclusion of a zero motion vector in the candidate list. The motion vector predictor may be indicated in the bitstream for example by indicating the direction of the spatial motion vector predictor (up or left) or the selection of the temporal motion vector predictor candidate.

[0164] The merging/merge mode/process/mechanism may operate for example as follows, while other similar realizations of the merge mode are also possible for example with different candidate position sets and candidate locations with candidate position sets.

[0165] In the merging/merge mode/process/mechanism, where all the motion information of a block/PU is predicted and used without any modification/correction. The aforementioned motion information for a PU may comprise one or more of the following: 1) The information whether 'the PU is uni-predicted using only reference picture list0' or 'the PU is uni-predicted using only reference picture list1' or 'the PU is bi-predicted using both reference picture list0 and list1'; 2) Motion vector value corresponding to the reference picture list0, which may comprise a horizontal and vertical motion vector component; 3) Reference picture index in the reference picture list0 and/or an identifier of a reference picture pointed to by the Motion vector corresponding to reference picture list 0, where the identifier of a reference picture may be for example a picture order count value, a layer identifier value (for inter-layer prediction), or a pair of a picture order count value and a layer identifier value; 4) Information of the reference picture marking of the reference picture, e.g. information whether the reference picture was marked as "used for short-term reference" or "used for long-term reference"; 5) - 7) The same as 2) - 4), respectively, but for reference picture list1.

[0166] Similarly, predicting the motion information is carried out using the motion information of adjacent blocks and/or co-located blocks in temporal reference pictures. A list, often called as a merge list, may be constructed by including motion prediction candidates associated with available adjacent/co-located blocks and the index of selected motion prediction candidate in the list is signalled and the motion information of the selected candidate is copied to the motion information of the current PU. When the merge mechanism is employed for a whole CU and the prediction signal for the CU is used as the reconstruction signal, i.e. prediction residual is not processed, this type of coding/decoding the CU is typically named as skip mode or merge based skip mode. In addition to the skip mode, the merge mechanism may also be employed for individual PUs (not necessarily the whole CU as in skip mode) and in this case, prediction residual may be utilized to improve prediction quality. This type of prediction mode is typically named as an inter-merge mode.

[0167] One of the candidates in the merge list and/or the candidate list for AMVP or any similar motion vector candidate list may be a TMVP candidate or alike, which may be derived from the collocated block

within an indicated or inferred reference picture, such as the reference picture indicated for example in the slice header. In HEVC, the reference picture list to be used for obtaining a collocated partition is chosen according to the collocated\_from\_l0\_flag syntax element in the slice header. When the flag is equal to 1, it specifies that the picture that contains the collocated partition is derived from list 0, otherwise the picture is derived from list 1. When collocated\_from\_l0\_flag is not present, it is inferred to be equal to 1. The collocated\_ref\_idx in the slice header specifies the reference index of the picture that contains the collocated partition. When the current slice is a P slice, collocated\_ref\_idx refers to a picture in list 0. When the current slice is a B slice, collocated\_ref\_idx refers to a picture in list 0 if collocated\_from\_l0 is 1, otherwise it refers to a picture in list 1. collocated\_ref\_idx always refers to a valid list entry, and the resulting picture is the same for all slices of a coded picture. When collocated\_ref\_idx is not present, it is inferred to be equal to 0.

[0168] In HEVC the so-called target reference index for temporal motion vector prediction in the merge list is set as 0 when the motion coding mode is the merge mode. When the motion coding mode in HEVC utilizing the temporal motion vector prediction is the advanced motion vector prediction mode, the target reference index values are explicitly indicated (e.g. per each PU).

[0169] In HEVC, the availability of a candidate predicted motion vector (PMV) may be determined as follows (both for spatial and temporal candidates) (SRTP = short-term reference picture, LTRP = long-term reference picture):

reference picture for target reference index	reference picture for candidate PMV	candidate PMV availability
STRP	STRP	“available” (and scaled)
STRP	LTRP	“unavailable”
LTRP	STRP	“unavailable”
LTRP	LTRP	“available” (but not scaled)

[0170] In HEVC, when the target reference index value has been determined, the motion vector value of the temporal motion vector prediction may be derived as follows: The motion vector PMV at the block that is collocated with the bottom-right neighbor of the current prediction unit is obtained. The picture where the collocated block resides may be e.g. determined according to the signalled reference index in the slice header as described above. If the PMV at location C0 is not available, the motion vector PMV at location C1 of the collocated picture is obtained. The determined available motion vector PMV at the collocated block is scaled with respect to the ratio of a first picture order count difference and a second picture order count difference. The first picture order count difference is derived between the picture containing the co-located block and the reference picture of the motion vector of the co-located block. The second picture order count difference is derived between the current picture and the target reference picture. If one

but not both of the target reference picture and the reference picture of the motion vector of the collocated block is a long-term reference picture (while the other is a short-term reference picture), the TMVP candidate may be considered unavailable. If both of the target reference picture and the reference picture of the motion vector of the collocated block are long-term reference pictures, no POC-based motion vector scaling may be applied.

[0171] Motion parameter types or motion information may include but are not limited to one or more of the following types:

- an indication of a prediction type (e.g. intra prediction, uni-prediction, bi-prediction) and/or a number of reference pictures;
- an indication of a prediction direction, such as inter (a.k.a. temporal) prediction, inter-layer prediction, inter-view prediction, view synthesis prediction (VSP), and inter-component prediction (which may be indicated per reference picture and/or per prediction type and where in some embodiments inter-view and view-synthesis prediction may be jointly considered as one prediction direction) and/or
- an indication of a reference picture type, such as a short-term reference picture and/or a long-term reference picture and/or an inter-layer reference picture (which may be indicated e.g. per reference picture)
- a reference index to a reference picture list and/or any other identifier of a reference picture (which may be indicated e.g. per reference picture and the type of which may depend on the prediction direction and/or the reference picture type and which may be accompanied by other relevant pieces of information, such as the reference picture list or alike to which reference index applies);
- a horizontal motion vector component (which may be indicated e.g. per prediction block or per reference index or alike);
- a vertical motion vector component (which may be indicated e.g. per prediction block or per reference index or alike);
- one or more parameters, such as picture order count difference and/or a relative camera separation between the picture containing or associated with the motion parameters and its reference picture, which may be used for scaling of the horizontal motion vector component and/or the vertical motion vector component in one or more motion vector prediction processes (where said one or more parameters may be indicated e.g. per each reference picture or each reference index or alike);
- coordinates of a block to which the motion parameters and/or motion information applies, e.g. coordinates of the top-left sample of the block in luma sample units;
- extents (e.g. a width and a height) of a block to which the motion parameters and/or motion information applies.

[0172] In general, motion vector prediction mechanisms, such as those motion vector prediction mechanisms presented above as examples, may include prediction or inheritance of certain pre-defined or indicated motion parameters.

[0173] A motion field associated with a picture may be considered to comprise of a set of motion information produced for every coded block of the picture. A motion field may be accessible by coordinates of a block, for example. A motion field may be used for example in TMVP or any other motion prediction mechanism where a source or a reference for prediction other than the current (de)coded picture is used.

[0174] Different spatial granularity or units may be applied to represent and/or store a motion field. For example, a regular grid of spatial units may be used. For example, a picture may be divided into rectangular blocks of certain size (with the possible exception of blocks at the edges of the picture, such as on the right edge and the bottom edge). For example, the size of the spatial unit may be equal to the smallest size for which a distinct motion can be indicated by the encoder in the bitstream, such as a 4x4 block in luma sample units. For example, a so-called compressed motion field may be used, where the spatial unit may be equal to a pre-defined or indicated size, such as a 16x16 block in luma sample units, which size may be greater than the smallest size for indicating distinct motion. For example, an HEVC encoder and/or decoder may be implemented in a manner that a motion data storage reduction (MDSR) or motion field compression is performed for each decoded motion field (prior to using the motion field for any prediction between pictures). In an HEVC implementation, MDSR may reduce the granularity of motion data to 16x16 blocks in luma sample units by keeping the motion applicable to the top-left sample of the 16x16 block in the compressed motion field. The encoder may encode indication(s) related to the spatial unit of the compressed motion field as one or more syntax elements and/or syntax element values for example in a sequence-level syntax structure, such as a video parameter set or a sequence parameter set. In some (de)coding methods and/or devices, a motion field may be represented and/or stored according to the block partitioning of the motion prediction (e.g. according to prediction units of the HEVC standard). In some (de)coding methods and/or devices, a combination of a regular grid and block partitioning may be applied so that motion associated with partitions greater than a pre-defined or indicated spatial unit size is represented and/or stored associated with those partitions, whereas motion associated with partitions smaller than or unaligned with a pre-defined or indicated spatial unit size or grid is represented and/or stored for the pre-defined or indicated units.

[0175] In HEVC, several improvements have been made to enable the codec to better utilize parallelism, i.e. parallel processing of encoding and/or decoding tasks, thus more efficiently utilizing modern multi-core processor architectures. While slices in principle can be used to parallelize the decoder, employing slices for parallelism typically results in relatively poor coding efficiency. The concept of wavefront processing has been introduced to HEVC to improve the utilization of parallelism.

[0176] To enable wavefront processing, the encoder and/or the decoder uses the CABAC state of the second CTU of the previous CTU row as the initial CABAC state of the current CTU row. Hence, the processing of the current CTU row can be started when the processing of the second CTU of the previous CTU has been finished. Thanks to this property, CTU rows can be processed in a parallel fashion. In general, it may be pre-defined e.g. in a coding standard which CTU is used for transferring the entropy (de)coding state of the previous row of CTUs or it may be determined and indicated in the bitstream by the encoder and/or decoded from the bitstream by the decoder.

[0177] The wavefront processing in HEVC may be used in two parallelization approaches, Wavefront Parallel Processing (WPP) and Overlapped Wavefront (OWF). WPP allows creating picture partitions that can be processed in parallel without incurring high coding losses.

[0178] WPP processes rows of coding tree units (CTU) in parallel while preserving all coding dependencies. In WPP, entropy coding, predictive coding as well as in-loop filtering can be applied in a single processing step, which makes the implementations of WPP rather straightforward. OWF, in turn, enables to overlap the execution of consecutive pictures. When the processing of a coding tree unit row in the current picture has been finished and no more rows are available, the processing of the next picture can be started instead of waiting for the current picture to finish.

[0179] When a coded picture has been constrained for wavefront processing or when tiles have been used, CTU rows or tiles (respectively) may be byte-aligned in the bitstream and may be preceded by a start code. Additionally, entry points may be provided in the bitstream (e.g. in the slice header) and/or externally (e.g. in a container file). An entry point is a byte pointer or a byte count or a similar straightforward reference mechanism to the start of a CTU row (for wavefront-enabled coded pictures) or a tile. In HEVC, entry points may be specified using `entry_point_offset_minus1[ i ]` of the slice header. In the HEVC file format (ISO/IEC 14496-15), the sub-sample information box may provide the information of entry points. In some scenarios, the use of dependent slice segments may be useful instead of or in addition to entry points. A dependent slice segment may be formed for example for a CTU row when a coded picture is constrained for wavefront processing and consequently the start of the dependent slice segment NAL unit may be used to determine CTU row boundaries.

[0180] Many video coding standards specify buffering models and buffering parameters for bitstreams. Such buffering models may be called Hypothetical Reference Decoder (HRD) or Video Buffer Verifier (VBV). A standard compliant bitstream complies with the buffering model with a set of buffering parameters specified in the corresponding standard. Such buffering parameters for a bitstream may be explicitly or implicitly signaled. 'Implicitly signaled' means for example that the default buffering parameter values according to the profile and level apply. The HRD/VBV parameters are used, among other things, to impose constraints on the bit rate variations of compliant bitstreams.

[0181] Video coding standards use variable-bitrate coding, which is caused for example by the flexibility of the encoder to select adaptively between intra and inter coding techniques for compressing video frames. To handle fluctuation in the bitrate variation of the compressed video, buffering may be used at the encoder and decoder side. Hypothetical Reference Decoder (HRD) may be regarded as a hypothetical decoder model that specifies constraints on the variability within conforming bitstreams, conforming NAL unit streams or conforming byte streams that an encoding process may produce.

[0182] A bitstream is compliant if it can be decoded by the HRD without buffer overflow or, in some cases, underflow. Buffer overflow happens if more bits are to be placed into the buffer when it is full. Buffer underflow happens if some bits are not in the buffer when said bits are to be fetched from the buffer for decoding/playback.

[0183] An HRD may be a part of an encoder or operationally connected to the output of the encoder. The buffering occupancy and possibly other information of the HRD may be used to control the encoding process. For example, if a coded data buffer in the HRD is about to overflow, the encoding bitrate may be reduced for example by increasing a quantizer step size.

[0184] The operation of the HRD may be controlled by HRD parameters, such as buffer size(s) and initial delay(s). The HRD parameter values may be created as part of the HRD process included or operationally connected to encoding. Alternatively, HRD parameters may be generated separately from encoding, for example in an HRD verifier that processes the input bitstream with the specified HRD process and generates such HRD parameter values according to which the bitstream is conforming. Another use for an HRD verifier is to verify that a given bitstream and given HRD parameters actually result into a conforming HRD operation and output.

[0185] HRD conformance checking may concern for example the following two types of bitstreams: The first such type of bitstream, called Type I bitstream, is a NAL unit stream containing only the VCL NAL units and filler data NAL units for all access units in the bitstream. The second type of bitstream, called a Type II bitstream, may contain, in addition to the VCL NAL units and filler data NAL units for all access units in the bitstream, additional non-VCL NAL units other than filler data NAL units and/or syntax elements such as `leading_zero_8bits`, `zero_byte`, `start_code_prefix_one_3bytes`, and `trailing_zero_8bits` that form a byte stream from the NAL unit stream.

[0186] Two types of HRD parameters (NAL HRD parameters and VCL HRD parameters) may be used. The HRD parameters may be indicated through video usability information included in the sequence parameter set syntax structure. The HRD parameters may, for example, include buffer size and input bitrate.

[0187] Buffering and picture timing parameters (e.g. included in sequence parameter sets and picture parameter sets referred to in the VCL NAL units and in buffering period and picture timing SEI messages) may be conveyed to the HRD, in a timely manner, either in the bitstream (by non-VCL NAL units), or by

out-of-band means externally from the bitstream e.g. using a signalling mechanism, such as media parameters included in the media line of a session description formatted e.g. according to the Session Description Protocol (SDP). For the purpose of counting bits in the HRD, only the appropriate bits that are actually present in the bitstream may be counted. When the content of a non-VCL NAL unit is conveyed  
5 for the application by some means other than presence within the bitstream, the representation of the content of the non-VCL NAL unit may or may not use the same syntax as would be used if the non-VCL NAL unit were in the bitstream.

[0188] The HRD may contain a coded picture buffer (CPB), an instantaneous decoding process, a decoded picture buffer (DPB), and output cropping.

10 [0189] The CPB may operate on decoding unit basis. A decoding unit may be an access unit or it may be a subset of an access unit, such as an integer number of NAL units. Encoders may determine that decoding units are for example tiles or CTU rows (when encoding constraints enabling wavefront processing have been applied). When a decoding unit is a subset of picture, a lower latency in the encoding and decoding may be achieved. The selection of the decoding unit may be indicated by an encoder in the  
15 bitstream. For example, decoding unit SEI messages may indicate decoding units as follows: The set of NAL units associated with a decoding unit information SEI message consists, in decoding order, of the SEI NAL unit containing the decoding unit information SEI message and all subsequent NAL units in the access unit up to but not including any subsequent SEI NAL unit containing a decoding unit information SEI message. Each decoding unit may be required to include at least one VCL NAL unit. All non-VCL NAL  
20 units associated with a VCL NAL unit may be included in the decoding unit containing the VCL NAL unit.

[0190] The HRD may operate as follows. Data associated with decoding units that flow into the CPB according to a specified arrival schedule may be delivered by the Hypothetical Stream Scheduler (HSS). The arrival schedule may be determined by the encoder and indicated for example through picture timing SEI messages, and/or the arrival schedule may be derived for example based on a bitrate which may be  
25 indicated for example as part of HRD parameters in video usability information (which may be included in the sequence parameter set). The HRD parameters in video usability information may contain many sets of parameters, each for different bitrate or delivery schedule. The data associated with each decoding unit may be removed and decoded instantaneously by the instantaneous decoding process at CPB removal times. A CPB removal time may be determined for example using an initial CPB buffering delay, which may be  
30 determined by the encoder and indicated for example through a buffering period SEI message, and differential removal delays indicated for each picture for example through picture timing SEI messages. The initial arrival time (i.e. the arrival time of the first bit) of the very first decoding unit may be determined to be 0. The initial arrival time of any subsequent decoding unit may be determined to be equal to the final arrival time of the previous decoding unit. Each decoded picture is placed in the DPB. A decoded picture

may be removed from the DPB at the later of the DPB output time or the time that it becomes no longer needed for inter-prediction reference. Thus, the operation of the CPB of the HRD may comprise timing of decoding unit initial arrival (when the first bit of the decoding unit enters the CPB), timing of decoding unit removal and decoding of decoding unit, whereas the operation of the DPB of the HRD may comprise  
5 removal of pictures from the DPB, picture output, and current decoded picture marking and storage.

[0191] The operation of an AU-based coded picture buffering in the HRD can be described in a simplified manner as follows. It is assumed that bits arrive into the CPB at a constant arrival bitrate (when the so-called low-delay mode is not in use). Hence, coded pictures or access units are associated with initial arrival time, which indicates when the first bit of the coded picture or access unit enters the CPB.

10 Furthermore, in the low-delay mode the coded pictures or access units are assumed to be removed instantaneously when the last bit of the coded picture or access unit is inserted into CPB and the respective decoded picture is inserted then to the DPB, thus simulating instantaneous decoding. This time is referred to as the removal time of the coded picture or access unit. The removal time of the first coded picture of the coded video sequence is typically controlled, for example by the Buffering Period Supplemental  
15 Enhancement Information (SEI) message. This so-called initial coded picture removal delay ensures that any variations of the coded bitrate, with respect to the constant bitrate used to fill in the CPB, do not cause starvation or overflow of the CPB. It is to be understood that the operation of the CPB is somewhat more sophisticated than what described here, having for example the low-delay operation mode and the capability to operate at many different constant bitrates. Moreover, the operation of the CPB may be specified  
20 differently in different standards.

[0192] When the bitstream starts at an IRAP picture, for example as a result of accessing a file or stream randomly and starting the decoding from an IRAP picture or tuning into a broadcast, there can be leading pictures (RADL and/or RASL pictures) that follow the IRAP picture in decoding order and precede it in output order. It is possible to discard or omit the decoding of these leading pictures following the RAP  
25 picture without affecting the decoding operation, as these leading pictures have no effect on the decoding process of any other pictures.

[0193] The buffering period SEI message of HEVC supports indicating two sets of initial buffering delay and initial buffering delay offset parameters, which can be signaled for example at an IRAP picture. One set of values specifies the required initial buffering when the leading pictures associated with the IRAP  
30 picture (with which the buffering period SEI message is associated) are present in the bitstream. The other set of values specifies the required initial buffering when leading pictures are not present in the bitstream or are discarded prior to scheduling them with HSS and/or inputting them into the CPB. The HRD operation may be required to be verified with the HRD for both sets of buffering parameters provided in the buffering period SEI message.

[0194] The DPB is used, among other things, to control the required memory resources for decoding of conforming bitstreams. There are two reasons to buffer decoded pictures, for references in prediction and for reordering decoded pictures into output order. As H.264/AVC and HEVC provide a great deal of flexibility for both reference picture marking and output reordering, separate buffers for reference picture buffering and output picture buffering could have been a waste of memory resources. Hence, the DPB includes a unified decoded picture buffering process for reference pictures and output reordering. A decoded picture may be removed from the DPB when it is no longer used as reference and needed for output.

[0195] In output cropping, lines and/or columns of samples may be removed from decoded pictures according to a cropping rectangle to form output pictures. In HEVC, a cropped decoded picture is defined as the result of cropping a decoded picture based on the conformance cropping window specified in the SPS that is referred to by the corresponding coded picture. Conforming decoders are required to produce numerically identical cropped decoded pictures as the decoding process specified in HEVC. The output cropping of HEVC produces cropped decoded pictures.

[0196] The HRD may be used to check conformance of bitstreams and decoders.

[0197] Bitstream conformance requirements of the HRD may comprise for example the following and/or alike. The CPB is required not to overflow (relative to the size which may be indicated for example within HRD parameters of video usability information) or underflow (i.e. the removal time of a decoding unit cannot be smaller than the arrival time of the last bit of that decoding unit). The number of pictures in the DPB may be required to be smaller than or equal to a certain maximum number, which may be indicated for example in the sequence parameter set. All pictures used as prediction references may be required to be present in the DPB. It may be required that the interval for outputting consecutive pictures from the DPB is not smaller than a certain minimum.

[0198] Decoder conformance requirements of the HRD may comprise for example the following and/or alike. A decoder claiming conformance to a specific profile and level may be required to decode successfully all conforming bitstreams specified for decoder conformance provided that all sequence parameter sets and picture parameter sets referred to in the VCL NAL units, and appropriate buffering period and picture timing SEI messages are conveyed to the decoder, in a timely manner, either in the bitstream (by non-VCL NAL units), or by external means. There may be two types of conformance that can be claimed by a decoder: output timing conformance and output order conformance.

[0199] To check conformance of a decoder, test bitstreams conforming to the claimed profile and level may be delivered by a hypothetical stream scheduler (HSS) both to the HRD and to the decoder under test (DUT). All pictures output by the HRD may also be required to be output by the DUT and, for each picture

output by the HRD, the values of all samples that are output by the DUT for the corresponding picture may also be required to be equal to the values of the samples output by the HRD.

[0200] For output timing decoder conformance, the HSS may operate e.g. with delivery schedules selected from those indicated in the HRD parameters of video usability information, or with "interpolated" delivery schedules. The same delivery schedule may be used for both the HRD and DUT. For output timing decoder conformance, the timing (relative to the delivery time of the first bit) of picture output may be required to be the same for both HRD and the DUT up to a fixed delay.

[0201] For output order decoder conformance, the HSS may deliver the bitstream to the DUT "by demand" from the DUT, meaning that the HSS delivers bits (in decoding order) only when the DUT requires more bits to proceed with its processing. The HSS may deliver the bitstream to the HRD by one of the schedules specified in the bitstream such that the bit rate and CPB size are restricted. The order of pictures output may be required to be the same for both HRD and the DUT.

[0202] In scalable video coding, a video signal can be encoded into a base layer and one or more enhancement layers. An enhancement layer may enhance the temporal resolution (i.e., the frame rate), the spatial resolution, or simply the quality of the video content represented by another layer or part thereof. Each layer together with all its dependent layers is one representation of the video signal at a certain spatial resolution, temporal resolution and quality level. In this document, we refer to a scalable layer together with all of its dependent layers as a "scalable layer representation". The portion of a scalable bitstream corresponding to a scalable layer representation can be extracted and decoded to produce a representation of the original signal at certain fidelity.

[0203] In the following, the term layer is used in context of any type of scalability, including view scalability and depth enhancements. An enhancement layer refers to any type of an enhancement, such as SNR, spatial, multiview, depth, bit-depth, chroma format, and/or color gamut enhancement. A base layer also refers to any type of a base operation point, such as a base view, a base layer for SNR/spatial scalability, or a texture base view for depth-enhanced video coding.

[0204] Scalable video (de)coding may be realized with a concept known as single-loop decoding, where decoded reference pictures are reconstructed only for the highest layer being decoded while pictures at lower layers may not be fully decoded or may be discarded after using them for inter-layer prediction. In single-loop decoding, the decoder performs motion compensation and full picture reconstruction only for the scalable layer desired for playback (called the "desired layer" or the "target layer"), thereby reducing decoding complexity when compared to multi-loop decoding. All of the layers other than the desired layer do not need to be fully decoded because all or part of the coded picture data is not needed for reconstruction of the desired layer. However, lower layers (than the target layer) may be used for inter-layer syntax or parameter prediction, such as inter-layer motion prediction. Additionally or alternatively, lower layers may

be used for inter-layer intra prediction and hence intra-coded blocks of lower layers may have to be decoded. Additionally or alternatively, inter-layer residual prediction may be applied, where the residual information of the lower layers may be used for decoding of the target layer and the residual information may need to be decoded or reconstructed. In some coding arrangements, a single decoding loop is needed for decoding  
5 of most pictures, while a second decoding loop may be selectively applied to reconstruct so-called base representations (i.e. decoded base layer pictures), which may be needed as prediction references but not for output or display.

[0205] Some aspects of the SVC extension of the H.264/AVC standard are described next as an example of a scalable video coding standard.

10 [0206] SVC includes support for coarse-grain quality and spatial scalability (CGS), medium-grain quality scalability (MGS) and temporal scalability. In some scalable video coding schemes, data in an enhancement layer can be truncated after a certain location, or even at arbitrary positions, where each truncation position may include additional data representing increasingly enhanced visual quality. Such scalability is referred to as fine-grained (granularity) scalability (FGS). FGS was included in some draft  
15 versions of the SVC standard, but it was eventually excluded from the final SVC standard. FGS is subsequently discussed in the context of some draft versions of the SVC standard. The scalability provided by those enhancement layers that cannot be truncated is referred to as coarse-grained (granularity) scalability (CGS). It collectively includes the traditional quality (SNR) scalability and spatial scalability. The SVC standard supports the so-called medium-grained scalability (MGS), where quality enhancement  
20 pictures are coded similarly to SNR scalable layer pictures but indicated by high-level syntax elements similarly to FGS layer pictures, by having the quality\_id syntax element greater than 0.

[0207] SVC uses an inter-layer prediction mechanism, wherein certain information can be predicted from layers other than the currently reconstructed layer or the next lower layer. Information that could be inter-layer predicted includes intra texture, motion and residual data. Inter-layer motion prediction includes  
25 the prediction of block coding mode, header information, etc., wherein motion from the lower layer may be used for prediction of the higher layer. In case of intra coding, a prediction from surrounding macroblocks or from co-located macroblocks of lower layers is possible. These prediction techniques do not employ information from earlier coded access units and hence, are referred to as intra prediction techniques. Furthermore, residual data from lower layers can also be employed for prediction of the current  
30 layer.

[0208] SVC allows the use of single-loop decoding. It is enabled by using a constrained intra texture prediction mode, whereby the inter-layer intra texture prediction can be applied to macroblocks (MBs) for which the corresponding block of the base layer is located inside intra-MBs. At the same time, those intra-MBs in the base layer use constrained intra-prediction (e.g., having the syntax element

“constrained\_intra\_pred\_flag” equal to 1). In single-loop decoding, the decoder performs motion compensation and full picture reconstruction only for the scalable layer desired for playback (called the “desired layer” or the “target layer”), thereby greatly reducing decoding complexity. All of the layers other than the desired layer do not need to be fully decoded because all or part of the data of the MBs not used for inter-layer prediction (be it inter-layer intra texture prediction, inter-layer motion prediction or inter-layer residual prediction) is not needed for reconstruction of the desired layer. A single decoding loop is needed for decoding of most pictures, while a second decoding loop is selectively applied to reconstruct the base representations, which are needed as prediction references but not for output or display, and are reconstructed only for the so called key pictures (for which “store\_ref\_base\_pic\_flag” is equal to 1).

5 [0209] The scalability structure in the SVC draft is characterized by three syntax elements: “temporal\_id,” “dependency\_id” and “quality\_id.” The syntax element “temporal\_id” is used to indicate the temporal scalability hierarchy or, indirectly, the frame rate. A scalable layer representation comprising pictures of a smaller maximum “temporal\_id” value has a smaller frame rate than a scalable layer representation comprising pictures of a greater maximum “temporal\_id”. A given temporal layer typically depends on the lower temporal layers (i.e., the temporal layers with smaller “temporal\_id” values) but does not depend on any higher temporal layer. The syntax element “dependency\_id” is used to indicate the CGS inter-layer coding dependency hierarchy (which, as mentioned earlier, includes both SNR and spatial scalability). At any temporal level location, a picture of a smaller “dependency\_id” value may be used for inter-layer prediction for coding of a picture with a greater “dependency\_id” value. The syntax element  
15 “quality\_id” is used to indicate the quality level hierarchy of a FGS or MGS layer. At any temporal location, and with an identical “dependency\_id” value, a picture with “quality\_id” equal to QL uses the picture with “quality\_id” equal to QL-1 for inter-layer prediction. A coded slice with “quality\_id” larger than 0 may be coded as either a truncatable FGS slice or a non-truncatable MGS slice.

[0210] For simplicity, all the data units (e.g., Network Abstraction Layer units or NAL units in the SVC context) in one access unit having identical value of “dependency\_id” are referred to as a dependency unit or a dependency representation. Within one dependency unit, all the data units having identical value of “quality\_id” are referred to as a quality unit or layer representation.

[0211] A base representation, also known as a decoded base picture, is a decoded picture resulting from decoding the Video Coding Layer (VCL) NAL units of a dependency unit having “quality\_id” equal to 0 and for which the “store\_ref\_base\_pic\_flag” is set equal to 1. An enhancement representation, also referred to as a decoded picture, results from the regular decoding process in which all the layer representations that are present for the highest dependency representation are decoded.

[0212] As mentioned earlier, CGS includes both spatial scalability and SNR scalability. Spatial scalability is initially designed to support representations of video with different resolutions. For each time

instance, VCL NAL units are coded in the same access unit and these VCL NAL units can correspond to different resolutions. During the decoding, a low resolution VCL NAL unit provides the motion field and residual which can be optionally inherited by the final decoding and reconstruction of the high resolution picture. When compared to older video compression standards, SVC's spatial scalability has been  
5 generalized to enable the base layer to be a cropped and zoomed version of the enhancement layer.

[0213] MGS quality layers are indicated with "quality\_id". In a draft version of SVC, quality\_id was also used for indicating FGS quality layers. For each dependency unit (with the same "dependency\_id"), there is a layer with "quality\_id" equal to 0 and there can be other layers with "quality\_id" greater than 0. These layers with "quality\_id" greater than 0 are MGS layers and in a draft version of SVC they were either  
10 MGS layers or FGS layers, depending on whether the slices are coded as truncatable slices.

[0214] In the basic form of FGS enhancement layers, only inter-layer prediction is used. Therefore, FGS enhancement layers can be truncated freely without causing any error propagation in the decoded sequence. However, the basic form of FGS suffers from low compression efficiency. This issue arises because only low-quality pictures are used for inter prediction references. It has therefore been proposed  
15 that FGS-enhanced pictures be used as inter prediction references. However, this may cause encoding-decoding mismatch, also referred to as drift, when some FGS data are discarded.

[0215] One feature of a draft SVC standard is that the FGS NAL units can be freely dropped or truncated, and a feature of the SVC standard is that MGS NAL units can be freely dropped (but cannot be truncated) without affecting the decodability of the bitstream. As discussed above, when those FGS or MGS  
20 data have been used for inter prediction reference during encoding, dropping or truncation of the data would result in a mismatch between the decoded pictures in the decoder side and in the encoder side. This mismatch is also referred to as drift.

[0216] To control drift due to the dropping or truncation of FGS or MGS data, SVC applied the following solution: In a certain dependency unit, a base representation (by decoding only the CGS picture with "quality\_id" equal to 0 and all the dependent-on lower layer data) is stored in the decoded picture  
25 buffer. When encoding a subsequent dependency unit with the same value of "dependency\_id," all of the NAL units, including FGS or MGS NAL units, use the base representation for inter prediction reference. Consequently, all drift due to dropping or truncation of FGS or MGS NAL units in an earlier access unit is stopped at this access unit. For other dependency units with the same value of "dependency\_id," all of the  
30 NAL units use the decoded pictures for inter prediction reference, for high coding efficiency.

[0217] Each NAL unit includes in the NAL unit header a syntax element "use\_ref\_base\_pic\_flag." When the value of this element is equal to 1, decoding of the NAL unit uses the base representations of the reference pictures during the inter prediction process. The syntax element "store\_ref\_base\_pic\_flag"

specifies whether (when equal to 1) or not (when equal to 0) to store the base representation of the current picture for future pictures to use for inter prediction.

[0218] NAL units with “quality\_id” greater than 0 do not contain syntax elements related to reference picture lists construction and weighted prediction, i.e., the syntax elements “num\_ref\_active\_lx\_minus1” (x=0 or 1), the reference picture list reordering syntax table, and the weighted prediction syntax table are not present. Consequently, the MGS or FGS layers have to inherit these syntax elements from the NAL units with “quality\_id” equal to 0 of the same dependency unit when needed.

[0219] In SVC, a reference picture list consists of either only base representations (when “use\_ref\_base\_pic\_flag” is equal to 1) or only decoded pictures not marked as “base representation” (when “use\_ref\_base\_pic\_flag” is equal to 0), but never both at the same time.

[0220] In an H.264/AVC bit stream, coded pictures in one coded video sequence uses the same sequence parameter set, and at any time instance during the decoding process, only one sequence parameter set is active. In SVC, coded pictures from different scalable layers may use different sequence parameter sets. If different sequence parameter sets are used, then, at any time instant during the decoding process, there may be more than one active sequence picture parameter set. In the SVC specification, the one for the top layer is denoted as the active sequence picture parameter set, while the rest are referred to as layer active sequence picture parameter sets. Any given active sequence parameter set remains unchanged throughout a coded video sequence in the layer in which the active sequence parameter set is referred to.

[0221] As indicated earlier, MVC is an extension of H.264/AVC. Many of the definitions, concepts, syntax structures, semantics, and decoding processes of H.264/AVC apply also to MVC as such or with certain generalizations or constraints. Some definitions, concepts, syntax structures, semantics, and decoding processes of MVC are described in the following.

[0222] An access unit in MVC is defined to be a set of NAL units that are consecutive in decoding order and contain exactly one primary coded picture consisting of one or more view components. In addition to the primary coded picture, an access unit may also contain one or more redundant coded pictures, one auxiliary coded picture, or other NAL units not containing slices or slice data partitions of a coded picture. The decoding of an access unit results in one decoded picture consisting of one or more decoded view components, when decoding errors, bitstream errors or other errors which may affect the decoding do not occur. In other words, an access unit in MVC contains the view components of the views for one output time instance.

[0223] A view component in MVC is referred to as a coded representation of a view in a single access unit.

[0224] Inter-view prediction may be used in MVC and refers to prediction of a view component from decoded samples of different view components of the same access unit. In MVC, inter-view prediction is

realized similarly to inter prediction. For example, inter-view reference pictures are placed in the same reference picture list(s) as reference pictures for inter prediction, and a reference index as well as a motion vector are coded or inferred similarly for inter-view and inter reference pictures.

[0225] An anchor picture is a coded picture in which all slices may reference only slices within the same access unit, i.e., inter-view prediction may be used, but no inter prediction is used, and all following coded pictures in output order do not use inter prediction from any picture prior to the coded picture in decoding order. Inter-view prediction may be used for IDR view components that are part of a non-base view. A base view in MVC is a view that has the minimum value of view order index in a coded video sequence. The base view can be decoded independently of other views and does not use inter-view prediction. The base view can be decoded by H.264/AVC decoders supporting only the single-view profiles, such as the Baseline Profile or the High Profile of H.264/AVC.

[0226] In the MVC standard, many of the sub-processes of the MVC decoding process use the respective sub-processes of the H.264/AVC standard by replacing term "picture", "frame", and "field" in the sub-process specification of the H.264/AVC standard by "view component", "frame view component", and "field view component", respectively. Likewise, terms "picture", "frame", and "field" are often used in the following to mean "view component", "frame view component", and "field view component", respectively.

[0227] In scalable multiview coding, the same bitstream may contain coded view components of multiple views and at least some coded view components may be coded using quality and/or spatial scalability.

[0228] Many video encoders utilize the Lagrangian cost function to find rate-distortion optimal coding modes, for example the desired macroblock mode and associated motion vectors. This type of cost function uses a weighting factor or  $\lambda$  (lambda) to tie together the exact or estimated image distortion due to lossy coding methods and the exact or estimated amount of information required to represent the pixel/sample values in an image area. The Lagrangian cost function may be represented by the equation:

$$[0229] \quad C=D+\lambda R$$

[0230] where C is the Lagrangian cost to be minimized, D is the image distortion (for example, the mean-squared error between the pixel/sample values in original image block and in coded image block) with the mode and motion vectors currently considered,  $\lambda$  is a Lagrangian coefficient and R is the number of bits needed to represent the required data to reconstruct the image block in the decoder (including the amount of data to represent the candidate motion vectors).

[0231] There are ongoing standardization activities to specify a multiview extension of HEVC (which may be referred to as MV-HEVC), a depth-enhanced multiview extension of HEVC (which may be referred

to as 3D-HEVC), and a scalable extension of HEVC (which may be referred to as SHVC). A multi-loop decoding operation has been envisioned to be used in all these specifications.

[0232] In scalable video coding schemes utilizing multi-loop (de)coding, decoded reference pictures for each (de)coded layer may be maintained in a decoded picture buffer (DPB). The memory consumption  
5 for DPB may therefore be significantly higher than that for scalable video coding schemes with single-loop (de)coding operation. However, multi-loop (de)coding may have other advantages, such as relatively few additional parts compared to single-layer coding.

[0233] In scalable video coding with multi-loop decoding, enhanced layers may be predicted from pictures that had been already decoded in the base (reference) layer. Such pictures may be stored in the  
10 DPB of base layer and may be marked as used for reference. In certain circumstances, a picture marked as used for reference may be stored in fast memory, in order to provide fast random access to its samples, and may remain stored after the picture is supposed to be displayed in order to be used as reference for prediction. This imposes requirements on memory organization. In order to relax such memory requirements, a conventional design in multi-loop multilayer video coding schemes (such as MVC)  
15 assumes restricted utilization of inter-layer predictions. Inter-layer/inter-view prediction for enhanced view is allowed from a decoded picture of the base view located at the same access unit, in other word representing the scene at the same time entity. In such designs, the number of reference pictures available for predicting enhanced views is increased by 1 for each reference view.

[0234] It has been proposed that in scalable video coding with multi-loop (de)coding operation pictures  
20 marked as used for reference need not originate from the same access units in all layers. For example, a smaller number of reference pictures may be maintained in an enhancement layer compared to the base layer. In some embodiments a temporal inter-layer prediction, which may also be referred to as a diagonal inter-layer prediction or diagonal prediction, can be used to improve compression efficiency in such coding scenarios. In general, diagonal prediction may refer to any prediction where the prediction crosses more  
25 than one scalability domain or scalability type. For example, diagonal prediction may refer to prediction that takes place from a different component type (e.g. from depth to texture) and from a different time instant (e.g. from a picture of a previous access unit in (de)coding order to a picture in the current access unit).

[0235] A decoding process may be specified with reference to a layer identifier list  
30 TargetDecLayerIdList, which specifies the list of layer identifier values, such as nuh\_layer\_id values. The layer identifier values may be in TargetDecLayerIdList in increasing order of the NAL units to be decoded. TargetDecLayerIdList may include the layer identifiers for layers that are intended to be output by the decoder as well as all the layers on which the output layers depend in the decoding process.

[0236] Work is ongoing to specify scalable and multiview extensions to the HEVC standard. The multiview extension of HEVC, referred to as MV-HEVC, is similar to the MVC extension of H.264/AVC. Similarly to MVC, in MV-HEVC, inter-view reference pictures can be included in the reference picture list(s) of the current picture being coded or decoded. The scalable extension of HEVC, referred to as SHVC, is planned to be specified so that it uses multi-loop decoding operation (unlike the SVC extension of H.264/AVC). SHVC uses a reference index based design, where an inter-layer reference picture can be included in a one or more reference picture lists of the current picture being coded or decoded (as described above). In an earlier draft of SHVC, another design was also investigated, which may be referred to as IntraBL or TextureRL, where a specific coding mode, e.g. in CU level, is used for using decoded/reconstructed sample values of a reference layer picture for prediction in an enhancement layer picture. The SHVC development has concentrated on development of spatial and coarse grain quality scalability.

[0237] Both MV-HEVC and SHVC use reference-index-based scalability and are therefore similar to each other. The high-level syntax, semantics and decoding operation of MV-HEVC and SHVC has been aligned to a great extent. A difference between SHVC and MV-HEVC is that SHVC enables spatial scalability and hence includes upsampling of sample arrays and motion field mapping or upsampling.

[0238] It is possible to use many of the same syntax structures, semantics, and decoding processes for MV-HEVC and reference-index-based SHVC. Furthermore, it is possible to use the same syntax structures, semantics, and decoding processes for depth coding too. Hereafter, term scalable multiview extension of HEVC (SMV-HEVC) is used to refer to a coding process, a decoding process, syntax, and semantics where largely the same (de)coding tools are used regardless of the scalability type and where the reference index based approach without changes in the syntax, semantics, or decoding process below the slice header is used. SMV-HEVC might not be limited to multiview, spatial, and coarse grain quality scalability but may also support other types of scalability, such as depth-enhanced video.

[0239] For the enhancement layer coding, the same concepts and coding tools of HEVC may be used in SHVC, MV-HEVC, and/or SMV-HEVC. However, the additional inter-layer prediction tools, which employ already coded data (including reconstructed picture samples and motion parameters a.k.a motion information) in reference layer for efficiently coding an enhancement layer, may be integrated to SHVC, MV-HEVC, and/or SMV-HEVC codec.

[0240] An access unit in SHVC, MV-HEVC and SMV-HEVC may be defined as a set of NAL units that are associated with each other according to a specified classification rule, are consecutive in decoding order, and contain the VCL NAL units of all coded pictures associated with the same output time and their associated non-VCL NAL units.

- [0241] In MV-HEVC, SMV-HEVC, and reference index based SHVC solution, the block level syntax and decoding process are not changed for supporting inter-layer texture prediction. Only the high-level syntax has been modified (compared to that of HEVC) so that reconstructed pictures (upsampled if necessary) from a reference layer of the same access unit can be used as the reference pictures for coding the current enhancement layer picture. The inter-layer reference pictures as well as the temporal reference pictures are included in the reference picture lists. The signalled reference picture index is used to indicate whether the current Prediction Unit (PU) is predicted from a temporal reference picture or an inter-layer reference picture. The use of this feature may be controlled by the encoder and indicated in the bitstream for example in a video parameter set, a sequence parameter set, a picture parameter, and/or a slice header. The indication(s) may be specific to an enhancement layer, a reference layer, a pair of an enhancement layer and a reference layer, specific TemporalId values, specific picture types (e.g. IRAP pictures), specific slice types (e.g. P and B slices but not I slices), pictures of a specific POC value, and/or specific access units, for example. The scope and/or persistence of the indication(s) may be indicated along with the indication(s) themselves and/or may be inferred.
- [0242] The reference list(s) in MV-HEVC, SMV-HEVC, and a reference index based SHVC solution may be initialized using a specific process in which the inter-layer reference picture(s), if any, may be included in the initial reference picture list(s). are constructed as follows. For example, the temporal references may be firstly added into the reference lists (L0, L1) in the same manner as the reference list construction in HEVC. After that, the inter-layer references may be added after the temporal references. The inter-layer reference pictures may be for example concluded from the layer dependency information, such as the RefLayerId[ i ] variable derived from the VPS extension as described above. The inter-layer reference pictures may be added to the initial reference picture list L0 if the current enhancement-layer slice is a P-Slice, and may be added to both initial reference picture lists L0 and L1 if the current enhancement-layer slice is a B-Slice. The inter-layer reference pictures may be added to the reference picture lists in a specific order, which can but need not be the same for both reference picture lists. For example, an opposite order of adding inter-layer reference pictures into the initial reference picture list 1 may be used compared to that of the initial reference picture list 0. For example, inter-layer reference pictures may be inserted into the initial reference picture 0 in an ascending order of nuh\_layer\_id, while an opposite order may be used to initialize the initial reference picture list 1.
- [0243] In the coding and/or decoding process, the inter-layer reference pictures may be treated as a long-term reference pictures.
- [0244] In SMV-HEVC and a reference index based SHVC solution, inter-layer motion parameter prediction may be performed by setting the inter-layer reference picture as the collocated picture for TMVP derivation. A motion field mapping process between two layers may be performed for example to avoid

block level decoding process modification in TMVP derivation. A motion field mapping could also be performed for multiview coding, but a present draft of MV-HEVC (JCT-3V document JCT3V-E1004) does not include such a process. The use of the motion field mapping feature may be controlled by the encoder and indicated in the bitstream for example in a video parameter set, a sequence parameter set, a picture parameter, and/or a slice header. The indication(s) may be specific to an enhancement layer, a reference layer, a pair of an enhancement layer and a reference layer, specific TemporalId values, specific picture types (e.g. RAP pictures), specific slice types (e.g. P and B slices but not I slices), pictures of a specific POC value, and/or specific access units, for example. The scope and/or persistence of the indication(s) may be indicated along with the indication(s) themselves and/or may be inferred.

5 [0245] In a motion field mapping process for spatial scalability, the motion field of the upsampled inter-layer reference picture is attained based on the motion field of the respective reference layer picture. The motion parameters (which may e.g. include a horizontal and/or vertical motion vector value and a reference index) and/or a prediction mode for each block of the upsampled inter-layer reference picture may be derived from the corresponding motion parameters and/or prediction mode of the collocated block in the reference layer picture. The block size used for the derivation of the motion parameters and/or prediction mode in the upsampled inter-layer reference picture may be for example 16×16. The 16×16 block size is the same as in HEVC TMVP derivation process where compressed motion field of reference picture is used.

15 [0246] In H.265/HEVC, the `sps_temporal_mvp_enabled_flag` indicates whether the TMVP mechanism may be in use (when the flag is equal to 1) or is not in use (when the flag is equal to 0) in the HEVC base layer/view (with `nuh_layer_id` equal to 0). When `sps_temporal_mvp_enabled_flag` is equal to 1, the `slice_temporal_mvp_enabled_flag` is present in the slice header and indicates if the TMVP mechanism is in use for the current picture.

20 [0247] There may be “black box” implementations of scalable extensions of HEVC, where the base layer decoding/coding is implemented with an existing HEVC v1 implementation without changes. Such an implementation of base layer decoding/coding would store motion fields only if `sps_temporal_mvp_enabled_flag` is equal to 1.

25 [0248] Base layer motion fields may be used for either or both of the following two purposes: temporal motion vector prediction between pictures of the base layer and inter-layer motion prediction. If the base layer motion fields are used only for inter-layer motion prediction, the memory used for base layer motion fields could be de-allocated or used for other purposes after decoding of the access unit has been finished (or, more accurately, decoding of all layers within the access unit that may use the abase layer as motion prediction reference has been finished). However, when `sps_temporal_mvp_enabled_flag` is used to control the storage of base layer motion fields, it cannot be used to indicate that base layer motion fields are used only for inter-layer motion prediction and not for temporal motion vector prediction within the base layer.

[0249] In a textureRL based SHVC solution, the inter-layer texture prediction may be performed at CU level for which a new prediction mode, named as textureRL mode, is introduced. The collocated upsampled base layer block is used as the prediction for the enhancement layer CU coded in textureRL mode. For an input CU of the enhancement layer encoder, the CU mode may be determined among intra, inter and textureRL modes, for example. The use of the textureRL feature may be controlled by the encoder and indicated in the bitstream for example in a video parameter set, a sequence parameter set, a picture parameter, and/or a slice header. The indication(s) may be specific to an enhancement layer, a reference layer, a pair of an enhancement layer and a reference layer, specific TemporalId values, specific picture types (e.g. RAP pictures), specific slice types (e.g. P and B slices but not I slices), pictures of a specific POC value, and/or specific access units, for example. The scope and/or persistence of the indication(s) may be indicated along with the indication(s) themselves and/or may be inferred. Furthermore, the textureRL may be selected by the encoder at CU level and may be indicated in the bitstream per each CU for example using a CU level flag (texture\_rl\_flag) which may be entropy-coded e.g. using context adaptive arithmetic coding (e.g. CABAC).

[0250] The residue of textureRL predicted CU may be coded as follows. The transform process of textureRL predicted CU may be the same as that for the intra predicted CU, where a discrete sine transform (DST) is applied to TU of luma component having 4x4 size and a discrete cosine transform (DCT) is applied to the other type of TUs. Transform coefficient coding of a textureRL-predicted CU may be the same to that of inter predicted CU, where no\_residue\_flag may be used to indicate whether the coefficients of the whole CU are skipped.

[0251] In a textureRL based SHVC solution, in addition to spatially and temporally neighboring PUs, the motion parameters of the collocated reference-layer block may also be used to form the merge candidate list. The base layer merge candidate may be derived at a location collocated to the central position of the current PU and may be inserted in a particular location of the merge list, such as as the first candidate in merge list. In the case of spatial scalability, the reference-layer motion vector may be scaled according to the spatial resolution ratio between the two layers. The pruning (duplicated candidates check) may be performed for each spatially neighboring candidate with collocated base layer candidate. For the collocated base layer merge candidate and spatial merge candidate derivation, a certain maximum number of merge candidates may be used; for example four merge candidates may be selected among candidates that are located in six different positions. The temporal merge candidate may be derived in the same manner as done for HEVC merge list. When the number of candidates does not reach to maximum number of merge candidates (which may be determined by the encoder and may be indicated in the bitstream and may be assigned to the variable MaxNumMergeCand), the additional candidates, including combined bi-predictive

candidates and zero merge candidates, may be generated and added at the end of the merge list, similarly or identically to HEVC merge list construction.

[0252] In some coding and/or decoding arrangements, a reference index based scalability and a block-level scalability approach, such a textureRL based approach, may be combined. For example, multiview-  
5 video-plus-depth coding and/or decoding may be performed as follows. A textureRL approach may be used between the components of the same view. For example, a depth view component may be inter-layer predicted using a textureRL approach from a texture view component of the same view. A reference index based approach may be used for inter-view prediction, and in some embodiments inter-view prediction may be applied only between view components of the same component type.

10 [0253] Work is also ongoing to specify depth-enhanced video coding extensions to the HEVC standard, which may be referred to as 3D-HEVC, in which texture views and depth views may be coded into a single bitstream where some of the texture views may be compatible with HEVC. In other words, an HEVC decoder may be able to decode some of the texture views of such a bitstream and can omit the remaining texture views and depth views..

15 [0254] Other types of scalability and scalable video coding include bit-depth scalability, where base layer pictures are coded at lower bit-depth (e.g. 8 bits) per luma and/or chroma sample than enhancement layer pictures (e.g. 10 or 12 bits), chroma format scalability, where enhancement layer pictures provide higher fidelity and/or higher spatial resolution in chroma (e.g. coded in 4:4:4 chroma format) than base layer pictures (e.g. 4:2:0 format), and color gamut scalability, where the enhancement layer pictures have a  
20 richer/broader color representation range than that of the base layer pictures - for example the enhancement layer may have UHD TV (ITU-R BT.2020) color gamut and the base layer may have the ITU-R BT.709 color gamut. Additionally or alternatively, depth enhancement layers or views, providing distance and/or disparity information, and/or layers with alpha pictures, providing transparency information, and/or other types of auxiliary picture layers may be provided as scalable layers. Any number of such other types of  
25 scalability may be realized for example with a reference index based approach or a block-based approach e.g. as described above.

[0255] Another categorization of scalable coding is based on whether the same or different coding standard or technology is used as the basis for the base layer and enhancement layers. Terms hybrid codec scalability or standards scalability may be used to indicate a scenario where one coding standard or system  
30 is used for some layers, while another coding standard or system is used for some other layers. For example, the base layer may be AVC-coded, while one or more enhancement layers may be coded with an HEVC extension, such as SHVC or MV-HEVC. It is possible that more than one layer is of a first coding standard or system, such as AVC or its extensions like MVC, and/or more than one layer is a second coding standard. It is possible that layers represent more than two coding standards. For example, the base layer may be

coded with AVC, an enhancement layer may be coded with MVC and represent a non-base view, and either or both of the previous layers may be enhanced by a spatial or quality scalable layer coded with SHVC.

[0256] In many video communication or transmission systems, transport mechanisms and multimedia container file formats there are mechanisms to transmit or store the base layer separately from the enhancement layer(s). It may be considered that layers are stored in or transmitted through separate logical channels. Examples are provided in the following:

1. ISO Base Media File Format (ISOBMFF, ISO/IEC International Standard 14496-12): Base layer can be stored as a track and each enhancement layer can be stored in another track. Similarly, in a hybrid codec scalability case, a non-HEVC-coded base layer can be stored as a track (e.g. of sample entry type 'avc1'), while the enhancement layer(s) can be stored as another track which is linked to the base-layer track using so-called track references.

2. Real-time Transport Protocol (RTP): either RTP session multiplexing or synchronization source (SSRC) multiplexing can be used to logically separate different layers.

3. MPEG-2 transport stream (TS): Each layer can have a different packet identifier (PID) value.

[0257] Many video communication or transmission systems, transport mechanisms and multimedia container file formats provides means to associate coded data of separate logical channels, such as of different tracks or sessions, with each other. For example, there are mechanisms to associate coded data of the same access unit together. For example, decoding or output times may be provided in the container file format or transport mechanism, and coded data with the same decoding or output time may be considered to form an access unit.

[0258] A way of categorizing different types of prediction is to consider across which domains or scalability types the prediction crosses. This categorization may lead into one or more of the following types of prediction, which may also sometimes be referred to as prediction directions:

- Temporal prediction e.g. of sample values or motion vectors from an earlier picture usually of the same scalability layer, view and component type (texture or depth).
- Inter-view prediction (which may be also referred to as cross-view prediction) referring to prediction taking place between view components usually of the same time instant or access unit and the same component type.
- Inter-layer prediction referring to prediction taking place between layers usually of the same time instant, of the same component type, and of the same view.
- Inter-component prediction may be defined to comprise prediction of syntax element values, sample values, variable values used in the decoding process, or anything alike from a component picture of one type to a component picture of another type. For example, inter-component

prediction may comprise prediction of a texture view component from a depth view component, or vice versa.

[0259] Prediction approaches using image information from a previously coded image can also be called as inter prediction methods. Inter prediction may sometimes be considered to only include motion-compensated temporal prediction, while it may sometimes be considered to include all types of prediction where a reconstructed/decoded block of samples is used as prediction source, therefore including conventional inter-view prediction for example. Inter prediction may be considered to comprise only sample prediction but it may alternatively be considered to comprise both sample and syntax prediction. As a result of syntax and sample prediction, a predicted block of pixels of samples may be obtained.

[0260] If the prediction, such as predicted variable values and/or prediction blocks, is not refined by the encoder using any form of prediction error or residual coding, prediction may be referred to as inheritance. For example, in the merge mode of HEVC, the prediction motion information is not refined e.g. by (de)coding motion vector differences, and hence the merge mode may be considered as an example of motion information inheritance.

[0261] Video coding schemes may utilize a prediction scheme between pictures. As discussed, prediction may be performed in the encoder for example through a process of block partitioning and block matching between a currently coded block (Cb) in the current picture and a reference block (Rb) in the picture which is selected as a reference. Therefore parameters of such prediction can be defined as motion information (MI) comprising for example on or more of the following: spatial coordinates of the Cb (e.g. coordinates of the top-left pixel of the Cb); a reference index refIdx or similar which specifies the picture in the reference picture list which is selected as reference picture; a motion vector (MV) specifying displacement between the spatial coordinated of the Cb and Rb in the reference picture ; and the size and shape of the motion partition (the size and shape of the matching block).

[0262] A motion field associated with a picture may be considered to comprise a set of motion information produced for every coded block of the picture. A motion field may be accessible by coordinates of a block, for example. A motion field may be used for example in Temporal motion vector prediction or any other motion prediction mechanism where a source or a reference for prediction other than the current decoded/coded picture is used.

[0263] Video coding schemes may utilize a temporal motion vector prediction scheme, such as the temporal direct mode in H.264/AVC or the temporal motion vector predictor (TMVP) candidate in the merge and AVMP modes of H.265/HEVC. In a temporal motion vector prediction scheme, at least a subset of the motion information of another picture is used to derive motion information or motion information predictor(s) for the current picture. Temporal motion vector prediction therefore requires storage of motion information of reference pictures.

[0264] In H.265/HEVC, the sequence parameter set includes the `sps_temporal_mvp_enabled_flag` syntax element, which indicates if the slice header includes the `slice_temporal_mvp_enabled` flag. If `sps_temporal_mvp_enabled_flag` is equal to 0, no temporal motion vector predictors are used in the coded video sequence. `slice_temporal_mvp_enabled_flag` specifies whether temporal motion vector predictors  
5 can be used for inter prediction. When `slice_temporal_mvp_enabled_flag` is equal to 1, there are syntax elements in the slice header that identify the collocated picture used to derive the temporal motion vector predictors.

[0265] Temporal motion vector prediction can also be used in scalable video coding when a motion field of an inter-layer reference picture is used to predict or derive motion information of the current picture.

[0266] Motion field mapping may be used for example when an inter-layer reference picture is of  
10 different spatial resolution than the current picture. In a motion field mapping process for spatial scalability, the motion field of the upsampled inter-layer reference picture is attained based on the motion field of the respective reference layer picture. The motion parameters (which may e.g. include a horizontal and/or vertical motion vector value and a reference index) and/or a prediction mode for each block of the  
15 upsampled inter-layer reference picture may be derived from the corresponding motion parameters and/or prediction mode of the collocated block in the reference layer picture.

[0267] The storage of motion information may be performed for example on the basis of the minimum size of a motion partition, e.g. 4x4 (of luma samples) in the case of H.264/AVC. In another example, the spatial granularity of motion information may be pre-defined for example in a coding standard and the  
20 coded motion information may be resampled or converted to that spatial granularity. For example, motion information can be stored for 16x16 blocks (of luma samples) in H.265/HEVC.

[0268] A HRD for a scalable video bitstream may operate similarly to a HRD for a single-layer bitstream. However, some changes may be required or desirable, particularly when it comes to the DPB operation in multi-loop decoding of a scalable bitstream. It is possible to specify DPB operation for multi-  
25 loop decoding of a scalable bitstream in multiple ways. In a layer-wise approach, each layer may have conceptually its own DPB, which may otherwise operate independently but some DPB parameters may be provided jointly for all the layer-wise DPBs and picture output may operate synchronously so that the pictures having the same output time are output at the same time or, in output order conformance checking, pictures from the same access unit are output next to each other. In another approach, referred to as the  
30 resolution-specific approach, layers having the same key properties share the same sub-DPB. The key properties may include one or more of the following: picture width, picture height, chroma format, bitdepth, color format/gamut.

[0269] It may be possible to support both layer-wise and resolution-specific DPB approach with the same DPB model, which may be referred to as the sub-DPB model. The DPB is partitioned into several

sub-DPBs, and each sub-DPB is otherwise managed independently but some DPB parameters may be provided jointly for all the sub-DPBs and picture output may operate synchronously so that the pictures having the same output time are output at the same time or, in output order conformance checking, pictures from the same access unit are output next to each other.

5 [0270] A coding standard may include a sub-bitstream extraction process, and such is specified for example in SVC, MVC, and HEVC. The sub-bitstream extraction process relates to converting a bitstream, typically by removing NAL units, to a sub-bitstream, which may also be referred to as a bitstream subset. The sub-bitstream still remains conforming to the standard. For example, in HEVC, the bitstream created by excluding all VCL NAL units having a TemporalId value greater than a selected value and including all  
10 other VCL NAL units remains conforming. In HEVC, the sub-bitstream extraction process takes a TemporalId and/or a list of nuh\_layer\_id values as input and derives a sub-bitstream (also known as a bitstream subset) by removing from the bitstream all NAL units with TemporalId greater than the input TemporalId value or nuh\_layer\_id value not among the values in the input list of nuh\_layer\_id values.

[0271] A coding standard or system may refer to a term operation point or alike, which may indicate  
15 the scalable layers and/or sub-layers under which the decoding operates and/or may be associated with a sub-bitstream that includes the scalable layers and/or sub-layers being decoded. Some non-limiting definitions of an operation point are provided in the following.

[0272] In HEVC, an operation point is defined as bitstream created from another bitstream by operation  
20 of the sub-bitstream extraction process with the another bitstream, a target highest TemporalId, and a target layer identifier list as inputs.

[0273] In SHVC and MV-HEVC, an operation point definition may include a consideration a target  
output layer set. In SHVC and MV-HEVC, an operation point may be defined as a bitstream that is created from another bitstream by operation of the sub-bitstream extraction process with the another bitstream, a target highest TemporalId, and a target layer identifier list as inputs, and that is associated with a set of  
25 target output layers.

[0274] In MVC, an operation point may be defined as follows: An operation point is identified by a  
temporal\_id value representing the target temporal level and a set of view\_id values representing the target  
output views. One operation point is associated with a bitstream subset, which consists of the target output  
views and all other views the target output views depend on, that is derived using the sub-bitstream  
30 extraction process with tIdTarget equal to the temporal\_id value and viewIdTargetList consisting of the set  
of view\_id values as inputs. More than one operation point may be associated with the same bitstream  
subset. When "an operation point is decoded", a bitstream subset corresponding to the operation point may  
be decoded and subsequently the target output views may be output.

[0275] When a bitstream, such as an HEVC bitstream, starts at a CRA or BLA picture, it is not possible to decode the RASL pictures associated with the CRA or BLA picture correctly, because some reference pictures of these RASL pictures might not have been decoded. These RASL pictures are therefore not output by the decoding process and/or HRD. It may also be possible to provide external means to impact the decoding process, such as an interface or an API to the decoder, through which the decoder can be controlled to treat a CRA picture similarly to a BLA picture or a CRA picture initiating a bitstream and hence omit the output of the associated RASL pictures. The decoding process may for example associate a variable NoRaslOutputFlag with each IRAP picture and derive a value for the variable for example as follows:

- If the current picture is an IDR picture, a BLA picture, the first picture in the bitstream in decoding order, or the first picture that follows an end of sequence NAL unit in decoding order, the variable NoRaslOutputFlag is set equal to 1.
- Otherwise, if some external means are available to set the variable HandleCraAsBlaFlag to a value for the current picture, the variable HandleCraAsBlaFlag is set equal to the value provided by the external means and the variable NoRaslOutputFlag is set equal to HandleCraAsBlaFlag.
- Otherwise, the variable HandleCraAsBlaFlag is set equal to 0 and the variable NoRaslOutputFlag is set equal to 0.

[0276] Available media file format standards include ISO base media file format (ISO/IEC 14496-12, which may be abbreviated ISOBMFF), MPEG-4 file format (ISO/IEC 14496-14, also known as the MP4 format), file format for NAL unit structured video (ISO/IEC 14496-15) and 3GPP file format (3GPP TS 26.244, also known as the 3GP format). The SVC and MVC file formats are specified as amendments to the AVC file format. The ISO file format is the base for derivation of all the above mentioned file formats (excluding the ISO file format itself). These file formats (including the ISO file format itself) are generally called the ISO family of file formats.

[0277] The basic building block in the ISO base media file format is called a box. Each box has a header and a payload. The box header indicates the type of the box and the size of the box in terms of bytes. A box may enclose other boxes, and the ISO file format specifies which box types are allowed within a box of a certain type. Furthermore, the presence of some boxes may be mandatory in each file, while the presence of other boxes may be optional. Additionally, for some box types, it may be allowable to have more than one box present in a file. Thus, the ISO base media file format may be considered to specify a hierarchical structure of boxes.

[0278] According to the ISO family of file formats, a file includes media data and metadata that are enclosed in separate boxes. In an example embodiment, the media data may be provided in a media data (mdat) box and the movie (moov) box may be used to enclose the metadata. In some cases, for a file to be operable, both of the mdat and moov boxes must be present. The movie (moov) box may include one or

more tracks, and each track may reside in one corresponding track box. A track may be one of the following types: media, hint, timed metadata. A media track refers to samples formatted according to a media compression format (and its encapsulation to the ISO base media file format). A hint track refers to hint samples, containing cookbook instructions for constructing packets for transmission over an indicated communication protocol. The cookbook instructions may include guidance for packet header construction and include packet payload construction. In the packet payload construction, data residing in other tracks or items may be referenced. As such, for example, data residing in other tracks or items may be indicated by a reference as to which piece of data in a particular track or item is instructed to be copied into a packet during the packet construction process. A timed metadata track may refer to samples describing referred media and/or hint samples. For the presentation of one media type, typically one media track is selected. Samples of a track may be implicitly associated with sample numbers that are incremented by 1 in the indicated decoding order of samples. The first sample in a track may be associated with sample number 1.

[0279] An example of a simplified file structure according to the ISO base media file format may be described as follows. The file may include the moov box and the mdat box and the moov box may include one or more tracks that correspond to video and audio, respectively.

[0280] The ISO base media file format does not limit a presentation to be contained in one file. As such, a presentation may be comprised within several files. As an example, one file may include the metadata for the whole presentation and may thereby include all the media data to make the presentation self-contained. Other files, if used, may not be required to be formatted to ISO base media file format, and may be used to include media data, and may also include unused media data, or other information. The ISO base media file format concerns the structure of the presentation file only. The format of the media-data files may be constrained by the ISO base media file format or its derivative formats only in that the media-data in the media files is formatted as specified in the ISO base media file format or its derivative formats.

[0281] The ability to refer to external files may be realized through data references. In some examples, a sample description box included in each track may provide a list of sample entries, each providing detailed information about the coding type used, and any initialization information needed for that coding. All samples of a chunk and all samples of a track fragment may use the same sample entry. A chunk may be defined as a contiguous set of samples for one track. The Data Reference (dref) box, also included in each track, may define an indexed list of uniform resource locators (URLs), uniform resource names (URNs), and/or self-references to the file containing the metadata. A sample entry may point to one index of the Data Reference box, thereby indicating the file containing the samples of the respective chunk or track fragment.

[0282] Movie fragments may be used when recording content to ISO files in order to avoid losing data if a recording application crashes, runs out of memory space, or some other incident occurs. Without movie fragments, data loss may occur because the file format may typically require that all metadata, e.g., the movie box, be written in one contiguous area of the file. Furthermore, when recording a file, there may not be sufficient amount of memory space (e.g., RAM) to buffer a movie box for the size of the storage available, and re-computing the contents of a movie box when the movie is closed may be too slow. Moreover, movie fragments may enable simultaneous recording and playback of a file using a regular ISO file parser. Finally, a smaller duration of initial buffering may be required for progressive downloading, e.g., simultaneous reception and playback of a file, when movie fragments are used and the initial movie box is smaller compared to a file with the same media content but structured without movie fragments.

[0283] The movie fragment feature may enable splitting the metadata that conventionally would reside in the movie box into multiple pieces. Each piece may correspond to a certain period of time for a track. In other words, the movie fragment feature may enable interleaving file metadata and media data. Consequently, the size of the movie box may be limited and the use cases mentioned above be realized.

[0284] In some examples, the media samples for the movie fragments may reside in an mdat box, as usual, if they are in the same file as the moov box. For the metadata of the movie fragments, however, a moof box may be provided. The moof box may include the information for a certain duration of playback time that would previously have been in the moov box. The moov box may still represent a valid movie on its own, but in addition, it may include an mvex box indicating that movie fragments will follow in the same file. The movie fragments may extend the presentation that is associated to the moov box in time.

[0285] Within the movie fragment there may be a set of track fragments, including anywhere from zero to a plurality per track. The track fragments may in turn include anywhere from zero to a plurality of track runs, each of which document is a contiguous run of samples for that track. Within these structures, many fields are optional and can be defaulted. The metadata that may be included in the moof box may be limited to a subset of the metadata that may be included in a moov box and may be coded differently in some cases. Details regarding the boxes that can be included in a moof box may be found from the ISO base media file format specification.

[0286] ISO/IEC 14496-15 specifies Aggregators and Extractors are file format internal structures enabling efficient grouping of NAL units or extraction of NAL units from other tracks. While presently ISO/IEC 14496-15 specifies the use of Aggregators and Extractors for AVC extensions, such as SVC and MVC, similar NAL-unit-like structures could be applied for HEVC extensions too. Aggregators and Extractors use the NAL unit syntax, but their payload has not necessarily been subject to start code emulation prevention. These structures are seen as NAL units in the context of the sample structure of the file format. While accessing a sample, Aggregators must be removed (leaving their contained or referenced

NAL Units) and Extractors must be replaced by the data they reference. Aggregators and Extractors must not be present in a stream outside the file format.

[0287] ISO/IEC 14496-15 supports aggregation of multiple NAL units into one aggregator NAL unit (the NAL unit type value of which was taken from the value range that is unspecified in the H.264/AVC standard). Aggregator NAL units can both aggregate by inclusion NAL units within them (within the size indicated by their length) and also aggregate by reference NAL units that follow them (within the area indicated by the additional\_bytes field within them). When the stream is scanned by an AVC file reader, only the included NAL units are seen as “within” the aggregator. This permits, for example, an AVC file reader to skip a whole set of unneeded SVC or MVC NAL units. SVC NAL units refer to the SVC specific NAL units for which the NAL unit type values are reserved by the AVC specification. MVC NAL units refer to the MVC specific NAL units for which the NAL unit type values are reserved by the AVC specification. Similarly, if AVC NAL units are aggregated by reference, the AVC reader will not skip them and they remain in-stream for that reader.

[0288] Another benefit achieved through using aggregators is presented in this paragraph. H.264/AVC, HEVC and their extensions allow an access unit to be coded in multiple NAL units. The number of NAL units can vary. In order to address a dependency representation (of SVC), a layer representation (of SVC) or a coded view (of MVC) as one logical unit by an ISOMBFF parser, one or more aggregator NAL units can be used. Aggregators help in organizing file format samples (e.g. access units) in constant patterns of number of NAL units per a logical scalability unit. For example, if all base layer NAL units of an access unit are aggregated into one aggregator, it can be considered that the base layer consists of one NAL unit. Some of the scalability and/or multiview properties are indicated in SVC and MVC file formats through a concept called tiers, which are specified through the sample grouping mechanism of ISOBMFF. Each Scalable Group Entry or Multiview Group Entry, included as a sample group description entry, documents a subset of the SVC stream or the MVC stream, respectively. Each of the subsets is associated with a tier and may contain one or more operating points. The number of Scalable Group Entry or Multiview Group Entry is equal to the number of different NAL unit sequence pattern to tier assignments. For example, one sample group description may indicate that a file format sample consists of one NAL unit belonging to a first tier and another NAL unit belonging to a second tier. Another sample group description entry would be needed for a file format sample consisting of two NAL units belonging to the first tier followed by a NAL unit belonging to the second tier. However, if an aggregator is used to aggregate the two NAL units belonging to the first tier, they are considered as one NAL unit in the sample group description entries. Consequently, aggregators help in reducing the amount of signaling needed to specify tiers.

[0289] Two forms of aggregation are supported currently by the aggregators. In the first, all samples that are aggregated are included within the aggregator NAL unit themselves. This form of aggregation is

called aggregation by inclusion. In the second form, all aggregated samples which are always stored as continued bytes in the ISOMBBFF storage format are signaled by reference. The aggregated NAL units are not contained within the Aggregator NAL unit but are referred to from within the Aggregator NAL unit. The syntax of the aggregator NAL unit includes an `additional_bytes` field, which indicates the number of bytes aggregated by reference. It is possible to have a single aggregator that aggregates both by inclusion and by reference by defining the length or size of the aggregator to include NAL units and by defining a non-zero `additional_bytes` value to aggregate NAL units by reference.

[0290] HEVC enable encoders and/or other entities, such as a splicers, to control picture output from the decoding process and/or HRD with the values of the `pic_output_flag` and/or `no_output_of_prior_pics_flag` syntax elements as described in the following.

[0291] The `pic_output_flag` syntax element, which may be present in the slice segment header, affects the derivation of the variable `PicOutputFlag`. Each decoded picture is associated with a value of `PicOutputFlag` (which may differ from the value of `PicOutputFlag` of other pictures). `PicOutputFlag` is set to 1 for pictures that are not RASL pictures associated with an IRAP picture with `NoRaslOutputFlag` equal to 1. Otherwise, `PicOutputFlag` is set to be equal to `pic_output_flag`. Pictures with `PicOutputFlag` equal to 0 are not output. Picture with `PicOutputFlag` equal to 1 are output, unless they become affected by `NoOutputOfPriorPicsFlag`, as explained in the following.

[0292] The `no_output_of_prior_pics_flag` syntax element, which is present for IRAP pictures in the slice segment header, affects the output of previously-decoded pictures in the decoded picture buffer after the decoding of an IDR or a BLA picture that is not the first picture in the bitstream as follows: The removal of pictures from the DPB before decoding of the current picture (but after parsing the slice header of the first slice of the current picture) happens instantaneously at the CPB removal time of the first decoding unit of access unit `n` (containing the current picture) and proceeds as follows:

- The reference picture set of the current picture is decoded.
- When the current picture is an IRAP picture with `NoRaslOutputFlag` equal to 1 that is not picture 0, the following ordered steps are applied:
  - o The variable `NoOutputOfPriorPicsFlag` is derived as follows:
    - If the current picture is a CRA picture, `NoOutputOfPriorPicsFlag` is set equal to 1 (regardless of the value of `no_output_of_prior_pics_flag`).
    - Otherwise, if the width, height, or the DPB capacity of the current picture differ from those of the previous picture, `NoOutputOfPriorPicsFlag` may but need not be set equal to 1 or may be set equal to `no_output_of_prior_pics_flag`.
    - Otherwise, `NoOutputOfPriorPicsFlag` is set equal to `no_output_of_prior_pics_flag`.

- The value of NoOutputOfPriorPicsFlag derived for the decoder under test is applied for the HRD, such that when the value of NoOutputOfPriorPicsFlag is equal to 1, all picture storage buffers in the DPB are emptied without output of the pictures they contain, and the DPB fullness is set equal to 0.
- 5     - When both of the following conditions are true for any pictures k in the DPB, all such pictures k in the DPB are removed from the DPB:
  - Picture k is marked as "unused for reference".
  - Picture k has PicOutputFlag equal to 0 or its DPB output time is less than or equal to the CPB removal time of the first decoding unit (denoted as decoding unit m) of the current picture n.
- 10     - For each picture that is removed from the DPB, the DPB fullness is decremented by one.

[0293] Splicers or other entities concatenating coded video sequences or bitstreams may set no\_output\_of\_prior\_pics\_flag equal to 1 for example when the output of a spliced coded video sequence might overlap with the output of some picture in the previous coded video sequence.

- 15 [0294] A coded video sequence (CVS) in scalable extensions of HEVC may be specified as follows: A coded video sequence is a sequence of access units that consists, in decoding order, of an IRAP access unit containing an IRAP picture having nuh\_layer\_id equal to 0 and NoRaslOutputFlag equal to 1, followed by zero or more access units that are not IRAP access units containing an IRAP picture having nuh\_layer\_id equal to 0 and NoRaslOutputFlag equal to 1, including all subsequent access units up to but not including
- 20 any subsequent access unit that is an IRAP access unit containing an IRAP picture having nuh\_layer\_id equal to 0 and NoRaslOutputFlag equal to 1.

- [0295] Adaptive Resolution Change (ARC) refers to dynamically changing the resolution within the video sequence, for example in video-conferencing use-cases. Adaptive Resolution Change may be used e.g. for better network adaptation and error resilience. For better adaptation to changing network requirements for different content, it may be desired to be able to change both the temporal/spatial resolution in addition to quality. The Adaptive Resolution Change may also enable a fast start, wherein the start-up time of a session may be able to be increased by first sending a low resolution frame and then increasing the resolution. The Adaptive Resolution Change may further be used in composing a conference. For example, when a person starts speaking, his/her corresponding resolution may be increased. Doing this
- 25 with an IDR frame may cause a "blip" in the quality as IDR frames need to be coded at a relatively low quality so that the delay is not significantly increased.

- 30 [0296] In the following some details of an adaptive resolution change use-cases are described in more detail using the scalable video coding framework. As scalable video coding inherently includes mechanisms

for resolution change, the adaptive resolution change could efficiently be supported as indicated in Figure 4.

[0297] In the example of Figure 4, switching happens at picture 3 (BL3) and the decoder receives the bitstream with following pictures: BL0-BL1-BL2-BL3-EL3-EL4-EL6-EL6... Pictures BL0—BL3 are pictures of a base layer and pictures EL3—EL6... are pictures of an enhancement layer. In this example pictures BL0—BL3 and EL4—EL6 use only intra-layer prediction and the picture EL3 uses either or both intra prediction and inter-layer prediction.

[0298] A sequence level signaling may be used to indicate the decoder if there is a resolution change in the bitstream. This may be performed e.g. by using a flag `single_layer_for_non_irap_flag`. The `single_layer_for_non_irap_flag` may be used to signal that a coded video sequence is constrained to perform the adaptive resolution change operation. The `single_layer_for_non_irap_flag` specifically indicates that, except for switching pictures, each AU in the sequence contains a single picture from a single layer (which may or may not be BL picture); and that access units where switching happens include pictures from two layers, the EL picture in such an access unit is an IRAP picture, and inter-layer prediction may be used for the EL IRAP picture.

[0299] Indicating `single_layer_for_non_irap_flag` in VPS allows the decoder to know that scalability is not used except for achieving resolution change, so that the decoder resources can be allocated accordingly upon a session start.

[0300] Scalable bitstreams with IRAP pictures or similar that are not aligned across layers may be used for example more frequent IRAP pictures can be used in the base layer, where they may have a smaller coded size due to e.g. a smaller spatial resolution. A process or mechanism for layer-wise start-up of the decoding may be included in a video decoding scheme. Decoders may hence start decoding of a bitstream when a base layer contains an IRAP picture and step-wise start decoding other layers when they contain IRAP pictures. In other words, in a layer-wise start-up of the decoding process, decoders progressively increase the number of decoded layers (where layers may represent an enhancement in spatial resolution, quality level, views, additional components such as depth, or a combination) as subsequent pictures from additional enhancement layers are decoded in the decoding process. The progressive increase of the number of decoded layers may be perceived for example as a progressive improvement of picture quality (in case of quality and spatial scalability).

[0301] A layer-wise start-up mechanism may generate unavailable pictures for the reference pictures of the first picture in decoding order in a particular enhancement layer. Alternatively, a decoder may omit the decoding of pictures preceding the IRAP picture from which the decoding of a layer can be started. These pictures that may be omitted may be specifically labeled by the encoder or another entity within the

bitstream. For example, one or more specific NAL unit types may be used for them. These pictures may be referred to as cross-layer random access skip (CL-RAS) pictures.

[0302] A layer-wise start-up mechanism may start the output of enhancement layer pictures from an IRAP picture in that enhancement layer, when all reference layers of that enhancement layer have been initialized similarly with an IRAP picture in the reference layers. In other words, any pictures (within the same layer) preceding such an IRAP picture in output order might not be output from the decoder and/or might not be displayed. In some cases, decodable leading pictures associated with such an IRAP picture may be output while other pictures preceding such an IRAP picture might not be output.

[0303] A layer-wise start-up mechanism may be initiated in one or more of the following cases:

- At the beginning of a bitstream.
- For base-layer IRAP pictures, when specifically controlled by means external to the decoding process through a notification to the decoder e.g. referred to as NoClrasOutputFlag (set equal to 1). This notification may be for example set by a player e.g. when a decoding process is started or re-started e.g. as response to tuning into a broadcast or seeking to a position in a file or stream.
- At a BLA picture in a base layer.
- At an IDR picture in a base layer, where the IDR picture may be specifically marked to initiate a layer-wise start-up mechanism for example using a flag in the slice header, which may for example be referred to as cross\_layer\_bla\_flag.
- At an IRAP picture in a base layer, where an associated indication indicates the invocation of the layer-wise start-up mechanism. The associated indication may for example reside in one or more of the following syntax structures: NAL unit header, slice header, slice segment header, picture parameter set, group of slices parameter set or similar, picture header, access unit delimiter, prefix NAL unit, suffix NAL unit. The indication in the NAL unit header may for example be a nal\_unit\_type value that indicates an IRAP picture causing the invocation of a layer-wise start-up mechanism.

[0304] In some syntax alternatives, the cross\_layer\_bla\_flag may be present only in slice segment headers of certain NAL unit types only. In some syntax alternatives, the cross\_layer\_bla\_flag may be present only in slice segment headers of certain layer(s) only, such as the base layer only or each layer not depending on any other layer.

[0305] It may be specified that when cross\_layer\_bla\_flag is equal to 1, a layer-wise start-up process is used in decoding. Likewise, it may be specified that when cross\_layer\_bla\_flag is equal to 0, a layer-wise start-up process is not used in decoding.

[0306] A decoding process may be specified in a manner that a certain variable controls whether or not a layer-wise start-up process is used. For example, a variable NoClrasOutputFlag may be used, which, when

equal to 0, indicates a normal decoding operation, and when equal to 1, indicates a layer-wise start-up operation. NoCrasOutputFlag may be set for example using one or more of the following steps:

- 1) If the current picture is an IRAP picture that is the first picture in the bitstream, NoCrasOutputFlag is set equal to 1.
- 5 2) Otherwise, if some external means are available to set the variable NoCrasOutputFlag equal to a value for a base-layer IRAP picture, the variable NoCrasOutputFlag is set equal to the value provided by the external means.
- 3) Otherwise, if the current picture is a BLA picture that is the first picture in a coded video sequence (CVS), NoCrasOutputFlag is set equal to 1.
- 10 4) Otherwise, if the current picture is an IDR picture that is the first picture in a coded video sequence (CVS) and cross\_layer\_bla\_flag is equal to 1, NoCrasOutputFlag is set equal to 1. 5) Otherwise, NoCrasOutputFlag is set equal to 0.

[0307] Step 4 above may alternatively may be phrased more generally for example as follows:

- [0308] Otherwise, if the current picture is an IRAP picture that is the first picture in a CVS and an indication of layer-wise start-up process is associated with the IRAP picture, NoCrasOutputFlag is set equal to 1. It should be understood that other ways to phrase the condition are possible and equally applicable.

- [0309] A decoding process for layer-wise start-up may be for example controlled by two array variables LayerInitialisedFlag[ i ] and FirstPicInLayerDecodedFlag[ i ] which may have entries for each layer (possibly excluding the base layer and possibly other independent layers too). When the layer-wise start-up process is invoked, for example as response to NoCrasOutputFlag being equal to 1, these array variables may be reset to their default values. For example, when there 64 layers are enabled (e.g. with a 6-bit nuh\_layer\_id), the variables may be reset as follows: the variable LayerInitialisedFlag[ i ] is set equal to 0 for all values of i from 0 to 63, inclusive, and the variable FirstPicInLayerDecodedFlag[ i ] is set equal to 0 for all values of i from 1 to 63, inclusive.

[0310] The decoding process may include the following or similar to control the output of RASL pictures. When the current picture is an IRAP picture, the following applies:

- If LayerInitialisedFlag[ nuh\_layer\_id ] is equal to 0, the variable NoRaslOutputFlag is set equal to 1.
- 30 – Otherwise, if some external means is available to set the variable HandleCraAsBlaFlag to a value for the current picture, the variable HandleCraAsBlaFlag is set equal to the value provided by the external means and the variable NoRaslOutputFlag is set equal to HandleCraAsBlaFlag.
- Otherwise, the variable HandleCraAsBlaFlag is set equal to 0 and the variable NoRaslOutputFlag is set equal to 0.

[0311] The decoding process may include the following to update the LayerInitialisedFlag for a layer. When the current picture is an IRAP picture and either one of the following is true, LayerInitialisedFlag[ nuh\_layer\_id ] is set equal to 1.

- nuh\_layer\_id is equal to 0.
- 5 – LayerInitialisedFlag[ nuh\_layer\_id ] is equal to 0 and LayerInitialisedFlag[refLayerId] is equal to 1 for all values of refLayerId equal to RefLayerId[ nuh\_layer\_id ][ j ], where j is in the range of 0 to NumDirectRefLayers[ nuh\_layer\_id ] – 1, inclusive.

[0312] When FirstPicInLayerDecodedFlag[ nuh\_layer\_id ] is equal to 0, the decoding process for generating unavailable reference pictures may be invoked prior to decoding the current picture. The decoding process for generating unavailable reference pictures may generate pictures for each picture in a reference picture set with default values. The process of generating unavailable reference pictures may be primarily specified only for the specification of syntax constraints for CL-RAS pictures, where a CL-RAS picture may be defined as a picture with nuh\_layer\_id equal to layerId and LayerInitialisedFlag[ layerId ] is equal to 0. In HRD operations, CL-RAS pictures may need to be taken into consideration in derivation of CPB arrival and removal times. In some embodiments, decoders may ignore any CL-RAS pictures, as these pictures are not specified for output and have no effect on the decoding process of any other pictures that are specified for output.

[0313] In some syntax alternatives, CL-RAS pictures may be specifically identified in the bitstream, e.g. by an encoder or by a splicer or another entity. CL-RAS pictures may have one or more certain nal\_unit\_type values, such as CL\_RAS\_N for CL-RAS pictures that are also sub-layer non-reference pictures and CL\_RAS\_R for CL-RAS pictures that are no sub-layer non-reference pictures. In another option, the same syntax element to indicate cross\_layer\_bla\_flag or similar may be used to indicate CL-RAS pictures or potential CL-RAS pictures. The cross\_layer\_bla\_flag or similar is applicable only to base-layer IRAP pictures, hence the same syntax element may be reused for indicating CL-RAS pictures, which appear only on enhancement layers.

[0314] CL-RAS pictures may have the property that when a layer-wise start-up mechanism is invoked (e.g. when NoCllasOutputFlag is equal to 1), the CL-RAS pictures are not output and may not be correctly decodable, as the CL-RAS picture may contain references to pictures that are not present in the bitstream. It may be specified that CL-RAS pictures are not used as reference pictures for the decoding process of non-CL-RAS pictures.

[0315] In contrast to the MVC/SVC design, the MV-HEVC/SHVC design allows access units in which the pictures may represent a subset of the layers present in the bitstream. This characteristic can be used for example for the following purposes:

- Different picture rate in different layers, e.g. the base layer at a lower spatial resolution with 60Hz and the enhancement layer at a higher spatial resolution with 30Hz;
- Use of SHVC for enabling adaptive resolution change of otherwise single-layer bitstream;
- Layer-wise start-up, where the number of layers is step-wise increased after the start of the

5

[0316] The MV-HEVC/SHVC design includes output layer sets specified in the VPS extension, each of which specifies the layers that are output by the decoder when that output layer set is in use. The selection of the output layer set in use is intended to be performed by external means (e.g. through an interface to the decoding), similarly to selecting the target temporal level for example. Each specified output layer set may be tested in conformance checking with HRD. In order to enable conformance checking in the HRD, the selected output layer set must remain unchanged for the entire coded video sequence (CVS). Term target output layer set may be used for the output layer set that has been selected into use by external means or by the HRD. Term target output layer may be used for each output layer included in a target output layer set.

10

[0317] The MV-HEVC/SHVC draft specification (in documents JCT3V-E1004 and JCTVC-N1008, respectively) derives a default output layer set for each layer set specified in the VPS. The VPS extension flag `default_one_target_output_layer_flag`, when equal to 0, specifies that each layer is an output layer, and when equal to 1, specifies that only the highest layer is an output layer. In addition, to the default output layer sets, the VPS extension enables to specify additional output layer sets with selected layers indicated to be output layers.

15

[0318] The target output layer set affects the picture output from the decoding process and/or or HRD. Only pictures that are on target output layers may be output (depending on their value of `PicOutputFlag` and depending on impacts of pictures having `NoOutputOfPriorPicsFlag` equal to 1, if any). Pictures that are not on target output layers are not output. This may be enforced for example by setting `PicOutputFlag` equal to 0 for pictures that are not included in any target output layer.

20

[0319] While output layer sets suit well use cases and bitstreams where the highest layer stays unchanged in each access unit, they do not support use cases where the highest layer changes from one access unit to another. Referring to the first example above, the output layer set could be specified in either of the following ways: The output layer set containing the enhancement layer may only cause the decoder to output at 30Hz picture range. The output layer set containing both the base layer and the enhancement may cause the decoder to output to picture per every other access unit, which might be difficult to handle in some standard APIs and by some renderers.

25

30

[0320] An approach to overcome the static nature of output layer sets may be to output a picture of the next lower present spatial or SNR scalable layer in the absence of a picture the specified output layer within an access unit. This approach is suboptimal due to the following reasons:

– It makes the MV-HEVC/SHVC design dependent on the current scalability types. If possible, the design should be independent of scalability types, which would enable the use of the approach for temporally varying output layers also for scalability types that are potentially specified in extensions beyond MV-HEVC/SHVC.

5 – It can be content-dependent whether the output of lower spatial and/or SNR layers is desired in the absence of the specified output layer. For example (referring to example 1 above), in some cases the content provider can intend the enhancement layer to be consumed at 30Hz (not interleaving the base layer pictures into the output e.g. due to too heavy quality fluctuation), while in other cases the content provider can intend the enhancement layer  
10 decoding to produce a 60Hz output.

[0321] In SVC, access units were required to have the same dependency\_id values present consistently. In other words, SVC disallowed layers having different picture rate, for example.

[0322] For so-called medium grain quality scalability (MGS), SVC included so-called layer representations that were identified by the quality\_id syntax element in the NAL unit header. SVC decoders  
15 decode the target dependency representation and up to the highest quality representation present in each access unit. The SVC syntax allowed removing any layer representations with quality\_id greater than 0 from the bitstream, while the resulting bitstream remained conforming. However, decoded pictures of such resulting bitstream suffer from a mismatch compared to the reconstructed pictures of the encoder. The encoder can stop the propagation and accumulation of the mismatch by indicating the use of base quality  
20 layer representations (with quality\_id equal to 0) in certain access units.

[0323] The SVC approach of using a scalability dimension (quality\_id in SVC) for indicating layers that may or may not be present in access units couples the layer output behavior with the scalability dimension. Consequently, the scalability dimension value itself would not be sufficient e.g. to differentiate between spatial and SNR scalability, as it merely indicates that the layers are alternatives when it comes to  
25 output and the preference order of outputting the layers (e.g. the layer representation with the highest present quality\_id is decoded and output). Moreover, the SVC approach required the base layer representation (quality\_id equal to 0) to be present in each access unit, which is undesirable e.g. in temporal scalability use cases.

[0324] According to a first aspect, each output layer in an output layer set is associated with a minimum  
30 alternative output layer. The minimum alternative output layer for each output layer of each output layer set is indicated in a sequence-level syntax structure (e.g. VPS in H.265/HEVC and its extensions). It is used to specify the minimum layer identifier value of a picture that can be output instead of the specified output layer, if a picture at the output layer is not present in an access unit. The first direct or indirect inter-layer

reference picture present in the access unit in descending layer identifier order down to the indicated minimum alternative output layer is output.

[0325] In a method according to an embodiment and shown in Figure 5 pictures are encoded into a bitstream 1040. The bitstream comprises at least two scalability layers and pictures being associated with access units. A first indication and a second indication are encoded into the bitstream 1042, 1044. The first indication is configured to indicate an output layer, and the second indication is configured to indicate at least one alternative output layer. The first indication and the second indication jointly have an impact that a first picture of said at least one alternative output layer is output by a decoding process of the bitstream when no picture of the output layer is in an access unit containing said first picture of said at least one alternative output layer. In a further embodiment, the layers are associated with one or more scalability dimensions and with different value sets of the one or more scalability dimensions. The second indication is coded independently of the one or more scalability dimensions.

[0326] In a method according to an embodiment and shown in Figure 6, pictures are decoded from a bitstream 1050. The bitstream comprises at least two scalability layers and at least two pictures being associated with access units. A first indication and a second indication are decoded from the bitstream 1052, 1054. The first indication indicates an output layer, and the second indication indicates at least one alternative output layer. A first picture of the at least one alternative layer is output 1056 when no picture of the output layer is in a first access unit containing said first picture. In a further embodiment, a second picture of the output layer is output, the second picture being associated with a second access unit. In a further embodiment, the layers are associated with one or more scalability dimensions and with different value sets of the one or more scalability dimensions. The second indication is coded independently of the one or more scalability dimensions. Hence, the assignment of one or more scalability dimensions and their values to layers is decoded independently of decoding the second indication to conclude at least one alternative output layer.

[0327] Yet in other words, each output layer is associated with a minimum alternative output layer. For the default output layer sets, the minimum output layer is the same as the output layer; i.e. when a default output layer set is in use, the decoder outputs only the specified output layers. For the additional output layer sets, the syntax element `min_alt_output_layer_idx[i][j]` is added to the VPS extension syntax for each output layer of the non-default output layer sets.

[0327] Yet in other words, each output layer is associated with a minimum alternative output layer. For the default output layer sets, the minimum output layer is the same as the output layer; i.e. when a default output layer set is in use, the decoder outputs only the specified output layers. For the additional output layer sets, the syntax element `min_alt_output_layer_idx[i][j]` is added to the VPS extension syntax for each output layer of the non-default output layer sets.

<code>if( output_layer_flag[ i ][ j ] )</code>	
<code>min_alt_output_layer_idx[ i ][ j ]</code>	<code>ue(v)</code>
<code>}</code>	

[0328] It is used to specify the minimum `nuh_layer_id` value of a picture that can be output instead of the specified output layer, if a picture at the output layer is not present in an access unit. The first direct or

indirect inter-layer reference picture present in the access unit in descending nuh\_layer\_id order down to the minimum value derived from min\_alt\_output\_layer\_idx[i][j] is output.

[0329] The variables NumRecursiveRefLayers[i] and RecursiveRefLayerFlag[i][j] are derived and indicate the number of direct and indirect reference layers and the direct and indicated reference layer mask, respectively, for layer i. The variables NumRecursiveRefLayers[i] and RecursiveRefLayerFlag[i][j] are derived as follows:

- NumRecursiveRefLayers[ i ] and RecursiveRefLayerFlag[ i ][ j ] are first initialized to 0 for all values of i and j in the range of 0 to 63, inclusive.

- The function setRefLayerFlags( currLayerId ) is specified as follows. The function is used to modify RecursiveRefLayerFlag[ i ][ j ].

```
for( j = 0; j < NumDirectRefLayers[ currLayerId ]; j++ )
```

```
    refLayerId = RefLayerId[ currLayerId ][ j ]
```

```
    setRefLayerFlags( refLayerId )
```

```
    for( k = 0; k <= 63; k++ )
```

```
        RecursiveRefLayerFlag[ currLayerId ][ k ] =
```

```
            RecursiveRefLayerFlag[ currLayerId ][ k ]|RecursiveRefLayerFlag[ refLayerId ][ k ]
```

- The function setRefLayerFlags( currLayerId ) is called as follows:

```
for( i = 0; i <= vps_max_layers_minus1; i++ ) {
```

```
    iNuhLId = layer_id_in_nuh[ i ]
```

```
    setRefLayerFlags( iNuhLId )
```

```
    for( j = 0; j < i; j++ )
```

```
        NumRecursiveRefLayers[ iNuhLId ] += RecursiveRefLayerFlag[ iNuhLId ][ j ]
```

[0330] The variable OutputLayerFlag[i][j] for output layer set with index i and for layer j is derived and indicates whether the layer is an output layer.

[0331] The variable MinAltOutputLayerId[i][j] for output layer set with index i and for layer j is derived and indicates the minimum nuh\_layer\_id value of a picture that can be output instead of the specified output layer, if a picture at the output layer is not present in an access unit. min\_alt\_output\_layer\_idx[i][j] is used to derive variable MinAltOutputLayerId[i][j]. The value of min\_alt\_output\_layer\_idx[i][j] shall be in the range of 0 to NumRecursiveRefLayers[ LayerSetLayerIdList[IsIdx][j] – 1, inclusive.

[0332] Variables OutputLayerFlag[i][j] and MinAltOutputLayerId[i][j] where i is the index of the output layer set and j is a nuh\_layer\_id value are derived as follows:

```
for( i = 0; i < numOutputLayerSets; i++ ) {
```

```
    for( j = 0; j <= 63; j++ ) {
```

```

        OutputLayerFlag[ i ][ j ] = 0
        MinAltOutputLayerId[ i ][ j ] = 0
    }
    if( i <= vps_number_layer_sets_minus1 )
5       if( default_one_target_output_layer_flag ) {
            layerId = LayerSetLayerIdList[ i ][ NumLayersInIdList[ i ] - 1 ]
            OutputLayerFlag[ i ][ layerId ] = 1
            MinAltOutputLayerId[ i ][ layerId ] = layerId
        } else
10        for( j = 0; j < NumLayersInIdList[ i ]; j++ ) {
            layerId = LayerSetLayerIdList[ i ][ j ]
            OutputLayerFlag[ i ][ layerId ] = 1
            MinAltOutputLayerId[ i ][ layerId ] = layerId
        }
15        else {
            lsIdx = output_layer_set_idx_minus1[ i ] + 1
            for( j = 0 ; j < NumLayersInIdList[ lsIdx ] - 1; j++ ) {
                layerId = LayerSetLayerIdList[ i ][ j ]
                OutputLayerFlag[ i ][ layerId ] = output_layer_flag[ i ][ j ]
20                MinAltOutputLayerId[ i ][ layerId ] =
                    LayerSetLayerIdList[ i ][ min_alt_output_layer_idx[ i ][ j ] ]
            }
        }
    }
25

```

[0333] The variable TargetOutputLayerSetIdx is set by external means, or (if not set by external means) by the HRD, or (if neither of the previous) set to 0.

- If some external means is available to set TargetOutputLayerSetIdx, TargetOutputLayerSetIdx is set by the external means.
- 30 – Otherwise, if the decoding process is invoked in a bitstream conformance test .
- Otherwise, TargetOutputLayerSetIdx is set equal to 0.

[0334] In an embodiment the derivation of PicOutputFlag stays unchanged compared to the current MV-HEVC and SHVC drafts (documents JCT3V-E1004 and JCTVC-N1008, respectively) when each picture is decoded. This process basically sets the value of PicOutputFlag as if the picture belonged to a

layer that is output. When all pictures of an access unit have been decoded the values of PicOutputFlag are updated depending on the output layers in the output layer set in use and if output layers are present in the access unit.

[0335] When the decoding of the last picture of an access unit has been completed PicOutputFlag for each picture in the access unit is updated as follows:

```

5   for( i = 0; i <= 63; i++ )
      outputLayerIdNeeded[ i ] = 0
   for( i = 0; i <= 63; i++ )
      if( OutputLayerFlag[ TargetOutputLayerSetIdx ][ i ] )
10          if( picture with nuh_layer_id equal to i is present in the access unit &&
              PicOutputFlag of the picture with with nuh_layer_id equal to i in the access unit
              is equal to 1 )
              outputLayerIdNeeded[ i ] = 1
          else {
15              foundFlag = 0
              for( j = i - 1; j >= MinAltOutputLayerId[ TargetOutputLayerSetIdx ][ i ] &&
                  !foundFlag; j-- )
                  if( RecursiveRefLayerFlag[ i ][ j ] &&
                      picture with nuh_layer_id equal to j is present in the access unit
20          &&
                      PicOutputFlag of the picture with with nuh_layer_id equal to i in
                      access unit is equal to 1 ) {
                          outputLayerIdNeeded[ j ] = 1
                          foundFlag = 1
25          }
          }
   for( i = 0; i <= 63; i++ )
      if( picture with nuh_layer_id equal to i is present in the access unit &&
          PicOutputFlag of the picture with nuh_layer_id equal to i in the access unit is equal to 1
30  &&
          outputLayerIdNeeded[ i ] is equal to 0 )
          PicOutputFlag of of the picture with nuh_layer_id equal to i in the access unit is set
          equal to 0

```

[0336] An embodiment which may applied independently of or together with other embodiments when a scalable bitstream is stored in a container file e.g. according to ISO/BMFF is described in this paragraph. Each track (or alike) of the file may be associated with an output layer set. The association may be explicit, i.e. a file creator may include for example an index of the output layer set or list the output layers in a file format structure such as an optional box within a sample entry. Alternatively or additionally, the association may be implicit, i.e. the track may contain certain layers, which may be indicated e.g. by tier definitions of ISO/IEC 14496-15, and the layers may be mapped to output layer set(s) for example on the basis of VPS. Similarly, each track (or alike) and/or each output layer set represented by a track may be associated with a set of alternative output layers as described above. For example, a sample entry may include, for example in an optional box, an identifier of a minimum alternative output layer. A file parser or player may interpret the above-mentioned information on output layers and alternative output layers. A file parser or player may use external means to control the decoder on the selection of an output layer set and/or output layers and/or alternative output layers. For example, a file parser or player may determine a track to be played, obtain the output layer set associated with the track, and input the above-mentioned TargetOutputLayerSetIdx according to the output layer set represented by the selected track to the decoder. Alternatively or additionally, when a file parser or player receives decoded pictures, it may determine on the basis of output layers and alternative output layers, which decoded pictures are passed to a displaying process.

[0337] According to a second aspect, instead of (or in addition to) specifying the output layers of an output layer set in a sequence-level structure, an output layer set is characterized in a sequence-level structure for example in terms of the number of output layers (i.e. output pictures) per access unit its use results into. An index or an identifier is also associated with the output layer sets specified in the sequence-level structure. A slice- or picture-level structure includes indications, e.g. a list of output layer set indices, indicating which output layer sets the slice or picture belongs to. For example, a slice segment header extension may include a minimum output layer set index and a maximum output layer set index, and the corresponding picture is included in all output layers sets in range from the indicated minimum to the indicated maximum, inclusive.

[0338] According to an embodiment, shown in Figure 7, pictures are encoded into bitstream 1070. The bitstream comprises at least two scalability layers. At least one indication is encoded into the bitstream 1072. The first indication provides one or more properties for an output layer set and associates an index for the output layer set. A first picture of a first layer is encoded into the bitstream 1074. At least one second indication is encoded into a picture-level structure within the bitstream 1076. The second indication associates the first picture to the output layer set. A second picture of a second layer is encoded into the bitstream 1078. At least one third indication is encoded into the picture-level structure within the bitstream 1079. The third indication associates the second picture to the output layer set. The first, second and third

indications jointly have an impact that the pictures being associated with the output layer set are output by decoding process of the bitstream wherein the output layer set has been selected.

[0339] According to an embodiment and shown in Figure 8, pictures are decoded from a bitstream 1080. The bitstream comprises at least two scalability layers. At least one first indication is decoded from the bitstream 1082. The first indication provides one or more properties for an output layer set and associates an index for the output layer set. It is determined that the output layer set is used in outputting from the decoder. A first picture of a first layer is decoded from the bitstream 1084. At least one second indication is decoded from a picture-level structure within the bitstream 1086. The second indication associates the first picture to the output layer set. A second picture of a second layer is decoded from the bitstream 1087. At least one third indication is decoded from the picture-level structure within the bitstream 1088. The third indication associates the second picture to the output layer set. The first picture and the second picture are output as they are pictures that are associated with the output layer set 1089.

[0340] An embodiment which may applied independently of or together with other embodiments when a scalable bitstream is stored in a container file e.g. according to ISO BMFF is described in this paragraph. An Aggregator NAL unit and/or an Extractor NAL unit according to the embodiment includes indications, e.g. a list of output layer set indices, indicating which output layer sets the contained or referred NAL units belong to. A file creator may aggregate NAL units using an Aggregator NAL unit and/or refer to NAL units from another track using an Extractor NAL unit on the basis that they belong to the same output layer set(s). A file parser or player may conclude the output layer set(s) for NAL units based on the information provided in the Aggregator and/or Extractor NAL units.

[0341] In both above aspects, the indication of output layers is independent of the scalability dimensions.

[0342] In an embodiment, which may be applied independently of or together with other embodiments, a decoder or an HRD or another entity may use information on alternative output layers to provide tolerance to picture losses and/or careless bitstream manipulation e.g. in middle-boxes as follows. The reference picture set design enables decoders to detect unintentional losses of reference pictures and the layer-wise start-up process enables decoders to increment the number of decoded layers in a step-wise manner. In the presently described embodiment, if a missing reference picture for an enhancement-layer picture is detected (e.g. by a reference to missing reference picture in an applied reference picture set) and an alternative output layer exists for that enhancement layer, a layer-wise start-up process is initiated for the enhancement layer and pictures of one or more alternative output layers are output until the enhancement layer contains an IRAP picture. The mentioned impacts may be realized in MV-HEVC and/or SHVC and/or other HEVC extensions for example as follows: LayerInitialisedFlag is set equal to 0 for the enhancement layer. This causes PicOutputFlag to be set to 0 for the enhancement layer until the next IRAP picture in the

enhancement layer and the use of an alternative output layer for picture output. FirstPicInLayerDecodedFlag is set equal to 0 for the picture referring to a missing reference picture. This causes the generation of unavailable reference pictures to be invoked. If a missing reference picture in a reference picture set for an enhancement-layer picture is detected and no alternative output layer exists for that enhancement layer, a bitstream conformance violation may be concluded and a decoder may for example invoke error concealment and/or interactive error control or correction, such as send feedback messages to the far-end encoder. The embodiment may be realized for example as follows: if there is an entry equal to "no reference picture" in any such subsets of the reference picture set that contains pictures that may be used as reference for the current picture (e.g. RefPicSetInterLayer0, RefPicSetInterLayer1, RefPicSetStCurrBefore, RefPicSetStCurrAfter, or RefPicSetLtCurr), the following applies with currLayerId being equal to nuh\_layer\_id of the current picture: If any such picture with nuh\_layer\_id layerIdA in the current access unit has been decoded for which nuhLayerIdA is greater than or equal to MinAltOutputLayerId[ TargetOutputLayerSetIdx ][ currLayerId ] and RecursiveRefLayerFlag[ currLayerId ][ layerIdA ] is equal to 1, LayerInitialisedFlag[ currLayerId ] is set equal to 0, FirstPicInLayerDecodedFlag[ currLayerId ] is set equal to 0, and noRefPicViolationFlag is set equal to 0.

[0343] In an embodiment, which may be applied independently of or together with other embodiments, an encoder may indicate in the bitstream prediction constraints, which may be taken advantage of when a layer is not among output layers. Examples of such prediction constraints and their signaling in the bitstream are provided in the following.

[0344] A picture in a reference layer that need to be decoded even if the reference layer is not among output layers may be referred to as a key picture. A picture in an output layer may for example use a key picture as a reference for inter-layer prediction, or another picture in the reference layer is used as a reference for inter-layer prediction of a picture in an output layer and the key picture may be used as a reference for that another picture. Pictures in a reference layer that are not key pictures may be referred to as non-key pictures.

[0345] Non-key pictures may be classified or characterized for example as follows:

- A non-key picture that is not used versus a non-key picture that may be used for sample prediction of other base-layer (BL) non-key pictures. (If a non-key picture is not used for sample prediction of other BL pictures, it need not be maintained in the DPB if the base layer is not among output layers. If a non-key picture is used for sample prediction of other BL pictures, in a single-loop decoding operation it needs to be replaced by the respective decoded EL picture, which introduces drift.)

- A non-key picture that is not used versus a non-key picture that may be used for motion prediction of BL key pictures and/or other BL non-key pictures.
- A non-key picture that is not used versus a non-key picture that may be used for inter-layer sample prediction.
- 5     - A non-key picture that is not used versus a non-key picture that may be used for inter-layer motion prediction.

[0346] In the embodiment presently described, an encoder may indicate in the bitstream whether a picture is a key picture or a non-key picture and/or it may indicate for a non-key picture one or more of the above-mentioned characteristics. Some detailed examples of indicating prediction characteristics follow.

10 [0347] In a first example, 2-bit fixed-length coded syntax element, i.e.  $u(2)$ , here referred to `pic_ref_idc`, is included in the syntax for a picture, for example in the slice segment header syntax structure, for example by taking certain two bit positions of the `slice_reserved[ i ]` syntax element of HEVC slice segment header into use. The semantics of `pic_ref_idc` may be specified for example as follows:

- 15     - `pic_ref_idc` equal to 3 indicates a discardable picture and specifies that the coded picture is not used as a reference picture for inter prediction and is not used as an inter-layer reference picture in the decoding process of subsequent pictures in decoding order
- `pic_ref_idc` equal to 2 indicates an intra-layer non-reference picture and specifies that the picture is not used as a reference picture for inter prediction of subsequent pictures in decoding order within the same layer and that the picture may be used as an inter-layer (sample or motion) prediction reference.
- 20     - `pic_ref_idc` equal to 1 indicates a non-output-layer skip (NOLS) picture and specifies that the picture may be used as a reference for inter prediction of subsequent pictures in decoding order that have `pic_ref_idc` equal to 1 or 3 until the next picture in decoding order with `pic_ref_idc` equal to 0 and the same or lower TemporalId value than that of the current NOLS picture. `pic_ref_idc` equal to 1 also specifies that the picture is not used as a reference for inter-layer prediction and not used as a reference for inter prediction of any picture with `pic_ref_idc` equal to 0.
- 25     - `pic_ref_idc` equal to 0 indicates a key picture and specifies that the picture may be used as a reference for inter prediction of subsequent pictures in decoding order within the same layer and may be used as a reference for inter-layer prediction.

30 [0348] A picture marked as intra-layer non-reference picture and either not used or no longer used as a reference picture for inter-layer prediction may be marked as “unused for reference”.

[0349] The signaling of the first example or any similar signaling enabling the identification of NOLS pictures enables removal of NOLS pictures from layers that are not among output layers. It may be required that either all discardable pictures and NOLS pictures (of a certain layer) are removed between two key

pictures (of the same certain layer) or that no NOLS pictures between two key pictures (of the same layer) are removed. The removal of NOLS pictures may be done by different entities including but not limited to the following:

- The removal may be done by a decoder or a bitstream pre-processor connected to a decoder, which have knowledge on the output layer set under which the decoding process operates. Consequently, the output layers can be concluded from the output layer set in use and NOLS pictures from layers that are not output but which are among the decoded layers can be removed. Such removal of NOLS pictures reduces decoding complexity and the memory required for decoded picture buffering.
- The removal may be done by an entity modifying the bitstream, such as a media-aware network element (MANE). The NOLS pictures may be removed from such layers that are not output layers among any specified output layer sets. The entity may modify the indicated output layer sets, for example based on receivers' properties and/or mode requests, to exclude certain output layer sets originally present in the indications, e.g. in VPS. For example, those output layer sets that are concluded to be unused among receivers or are concluded to be suboptimal for receivers' properties and/or mode requests may be removed. Consequently, the number of layers from which NOLS pictures can be removed may be increased.

[0350] A second example is otherwise the same as the first example, but it is specified that NOLS pictures may be used for inter-layer motion prediction but are not used for inter-layer sample prediction. Consequently, only the picture motion field of the NOLS pictures need to be decoded when the layer containing the NOLS pictures is not a target output layer, while the sample arrays of these NOLS pictures need not be decoded or need not be maintained in the DPB.

[0351] In a third example, NOLS pictures of both the first and second example are separately indicated. For example, in the 2-bit `pic_ref_idc` syntax element, a value indicating an intra-layer non-reference picture may be replaced by one of the NOLS picture types (of either example 1 or example 2).

[0352] In a fourth example, the NOLS picture is otherwise specified as in the first, second or third example, but the semantics of NOLS pictures are independent of the `TemporalId` value.

[0353] In an embodiment, which may be applied independently of or together with other embodiments, an encoder or another entity may include commands or alike into the bitstream that cause reference picture marking as "unused for reference" of a picture on a certain layer sooner than when the decoding of the next picture of that layer is started. Examples of such commands include but are not limited to the following:

- Include the reference picture set (RPS) to be applied after the decoding of the picture within the layer into the bitstream. Such an RPS may be referred to as a post-decoding RPS. A post-decoding RPS may be applied for example after the decoding of the access unit has been finished (which guarantees that no picture that is still used as a reference for inter-layer prediction becomes marked

as “unused for reference”). A post-decoding RPS may be included for example in a specific NAL unit, within a suffix NAL unit or a prefix NAL unit, and/or within slice header extension. It may be required that the post-decoding RPS is identical to or causes the same pictures to be maintained in the DPB as the RPS of the next picture in the same layer.

- 5       - Include a flag, e.g. in the slice segment header e.g. using a bit position of the slice\_reserved[ i ] syntax element of HEVC slice segment header, that causes marking of all pictures within the layer (including the current picture for which the flag is set to 1) as “unused for reference” after the decoding of the current picture for example when the access unit containing the current picture has been entirely decoded.
- 10       - Include a flag, e.g. in the slice segment header e.g. using a bit position of the slice\_reserved[ i ] syntax element of HEVC slice segment header, that causes marking of all non-key pictures (e.g. with key\_pic\_idc unequal to 0 in some above-mentioned examples) within the current layer “as unused for reference” after the decoding of the current picture for example when the access unit containing the current picture has been entirely decoded.
- 15       - The above-mentioned flags may be specific to TemporalId, i.e. cause pictures of the same and higher TemporalId value as that of the current picture to be marked as “unused for reference” (while the semantics of the flags are otherwise the same as above).
- An MMCO command or alike causing decoded reference picture marking.

[0354]     A decoder and/or HRD and/or another entity, such as a media-aware network element, may  
20     decode one or more of above-mentioned commands or alike from the bitstream and consequently mark reference pictures as “unused for reference”. The marking of a picture as “unused for reference” may affect the emptying or deallocation of picture storage buffers in the DPB as described earlier.

[0355]     In an embodiment, which may be applied independently of or together with other embodiments,  
25     the control of whether or not to output pictures preceding, in decoding order, an IRAP picture is applied over multiple layers. A base-layer picture that initiates the layer-wise start-up process causes the control whether or not to output pictures preceding, in decoding order, the base-layer picture to be applied across multiple layers, such as all layers present in the bitstream or all layers depending on the base layer. Similarly to applying the embodiment to a base-layer picture initiating a layer-wise start-up process, the embodiment may be applied to a picture on an independent layer (i.e. a layer having no reference layers) that initiates a  
30     layer-wise start-up process. When an IRAP picture initiating a layer-wise start-up process (e.g. an IRAP picture with NoClrasOutputFlag equal to 1) is decoded, all pictures in the DPB (in all layers) may be marked as “unused for reference” prior to including the IRAP picture in the DPB. As described earlier, the control on whether or not to output pictures preceding, in decoding order, an IRAP picture may be controlled through the NoOutputOfPriorPicsFlag variable on the derivation of which the

no\_output\_of\_prior\_pics\_flag syntax element has an impact. Moreover, as described earlier, the variable NoClrasOutputFlag can be used to indicate the start of a layer-wise start-up process. In a decoding process and/or HRD that is controlled by NoOutputOfPriorPicsFlag and NoClrasOutputFlag, this embodiment may be realized by setting NoOutputOfPriorPicsFlags of all pictures of an access unit to no\_output\_of\_prior\_pics\_flag of the base-layer picture for which NoClrasOutputFlag is set equal to 1 and which is contained in the same access unit. This embodiment enables a splicer or another entity to splice a coded video sequence starting with an access unit containing an IRAP picture in the base layer but not in all layers and to control the output of the pictures in the previous coded video sequence (preceding the spliced coded video sequence). In other words, thanks to this embodiment the splicer or another entity is able to control if pictures at any layer of the previous coded video sequence are output after the start of the decoding of the IRAP picture initiating a layer-wise start-up process.

[0356] In an embodiment, which may be applied independently of or together with other embodiments, the encoder or another entity, such as an HRD verifier, may indicate buffering parameters for one or both of the following types of bitstreams: bitstreams where CL-RAS pictures of IRAP pictures for which NoClrasOutputFlag is equal to 1 are present and bitstreams where CL-RAS picture of IRAP pictures for which NoClrasOutputFlag is equal to 1 are not present. For example, CPB buffer size(s) and bitrate(s) may be indicated separately e.g. in VUI for either or both mentioned types of bitstreams. Additionally or alternatively, the encoder or another entity may indicate initial CPB and/or DPB buffering delay and/or other buffering and/or timing parameters for either or both mentioned types of bitstreams. The encoder or another entity may, for example, include a buffering period SEI message into a scalable nesting SEI message, which may indicate the sub-bitstream, the layer set or the output layer set to which the contained buffering period SEI message applies. The buffering period SEI message of HEVC supports indicating two sets of parameters, one for the case where the leading pictures associated with the IRAP picture (for which the buffering period SEI message is also associated with) are present and another for the case where the leading pictures are not present. In the case when a buffering period SEI message is contained within a scalable nesting SEI message, the latter (alternative) set of parameters may be considered to concern a bitstream where CL-RAS pictures associated with the IRAP picture (for which the buffering period SEI message is also associated with) are not present. Generally, the latter set of buffering parameters may concern a bitstream where CL-RAS pictures associated with an IRAP picture for which NoClrasOutputFlag is equal to 1 are not present. It is to be understood that while specific terms and variable names are used in the description of this embodiment, it can be similarly realized with other terminology and need not use the same or similar variables as long as the decoder operation is similar.

[0357] In an embodiment, which may be applied independently of or together with other embodiments, the encoder or another entity, such as an HRD verifier, may indicate buffering parameters for output layer

sets within which NOLS pictures are removed on layers that are not output. In other words, buffering parameters may be indicated for a sub-bitstream containing the layers included in an output layer set where NOLS pictures of the layers not among the target output layers of the output layer set are removed. For example, CPB buffer size(s) and bitrate(s) may be indicated e.g. in VUI for such sub-bitstreams.

5 Additionally or alternatively, the encoder or another entity may indicate initial CPB and/or DPB buffering delay and/or other buffering and/or timing parameters for such sub-bitstreams. The encoder or another entity may, for example, include a buffering period SEI message into a nesting SEI message, which may indicate the output layer set to which the contained buffering period SEI message applies. The buffering period SEI message of HEVC supports indicating two sets of parameters, one for the case where the leading

10 pictures associated with the IRAP picture (for which the buffering period SEI message is also associated with) are present and another for the case where the leading pictures are not present. In the case when a buffering period SEI message is contained within a nesting SEI message applying to an output layer set, the latter (alternative) set of parameters may be considered to concern a bitstream where NOLS pictures of non-output layers and/or CL-RAS pictures associated with the IRAP picture (for which the buffering period

15 SEI message is also associated with) are not present. Generally, the latter set of buffering parameters may concern a bitstream where CL-RAS pictures associated with an IRAP picture for which NoClrasOutputFlag is equal to 1 are not present. It is to be understood that while specific terms and variable names are used in the description of this embodiment, it can be similarly realized with other terminology and need not use the same or similar variables as long as the decoder operation is similar.

20 [0358] In an embodiment, which may be applied independently of or together with other embodiments, a sub-bitstream extraction process is specified where inputs are a bitstream, an output layer set (or a list of output layers), and possibly a maximum TemporalId value (highestTid). An output layer set may be used to infer a target layer set that is required to be decoded, or the target layer set may be provided as input. An output of the process is a sub-bitstream. The sub-bitstream extraction process includes in the sub-bitstream

25 all VCL NAL units within the output layer set having TemporalId smaller than or equal to highestTid and the non-VCL NAL units associated with the included VCL NAL units. Furthermore, the sub-bitstream extraction process included those VCL NAL units from the non-output layers (but included in the target layer set) that are not NOLS pictures and not discardable pictures. Furthermore, the sub-bitstream extraction process may modify the syntax structure containing information on output layer sets, such as VPS, to

30 exclude those output layer sets where the layers from which NOLS pictures and discardable pictures are on output layers. Furthermore, the sub-bitstream extraction process may modify the signaling of HRD parameters such that the HRD parameters that apply for output layer sets without the presence of NOLS pictures (and discardable pictures) of the non-output layers are in force. Furthermore, the sub-bitstream

extraction process may exclude those HRD parameters from the bitstream that assume the presence of NOLS pictures (and discardable pictures) of the non-output layers.

[0359] In an embodiment, which may be applied independently of or together with other embodiments, a sub-bitstream extraction process is appended or specified to exclude CL-RAS pictures associated with IRAP pictures for which NoClrasOutputFlag is equal to 1. Furthermore, the sub-bitstream extraction process may exclude those HRD parameters from the bitstream that assume the presence of CL-RAS pictures that have been removed from the output bitstream.

[0360] In an embodiment, which may be applied independently of or together with other embodiments when a scalable bitstream is stored in a container file e.g. according to ISOBMFF is described in this paragraph. As described above, a track may be associated with one or more output layer sets explicitly or implicitly. If the track contains layers, which are not output layers, the NOLS pictures from these output layers need not be present. A track representing more than one layers may for example include base-layer data by reference from another track using Extractor NAL units. However, a file creator may omit the creation of Extractor NAL units for NOLS pictures of the base layer when the base layer is not among the output layers of a track containing the Extractor NAL units. A file creator may include an indication whether or not the track includes NOLS pictures of non-output layers for example within an optional box in a sample entry. Alternatively or additionally, a file creator may include (by inclusion or reference) NAL units of NOLS pictures into such Aggregator NAL units where the layer containing the included NOLS pictures is indicated not to be included among the output layers. A file parser or player may determine the output layer set and/or output layers to be played. It may subsequently select a track that represents the selected output layer set and/or output layers. If there is a track where NOLS pictures of non-output layers are excluded, the file parser or player may select that track to reduce decoding complexity and/or memory use. Additionally or alternatively, the file parser or player may skip (omit processing) such Aggregator NAL units that are indicated not to be mapped to the selected output layer set and/or selected output layers.

[0361] In an embodiment, which may be applied independently of or together with other embodiments, an encoder and/or a decoder and/or an HRD may derive a layer-wise picture order count value for pictures in a scalable bitstream. The layer-wise picture order count (POC) variable may be denoted as LayerPicOrderCntVal. The use of the layer-wise POC value may be used to provide continuity and/or integrity in POC difference values of enhancement layers when the POC in the base layer is reset, for example by a base-layer IDR picture that sets the POC value to 0, while the access unit containing the base-layer IDR picture does not contain IDR pictures in one or more other layers. The encoder may encode into the bitstream and the decoder/HRD may decode from the bitstream information on layer-wise POC, such as a layer-wise differential POC value. For example, the following syntax or alike may be used within the slice segment header syntax structure:

slice_segment_header( ) {
...
if( nuh_layer_id > 0 ) {
<b>layer_delta_poc_present_flag</b>
if( layer_delta_poc_present_flag )
<b>layer_delta_pic_order_cnt_minus1</b>
}
...

[0362] In the presently described embodiment, the variable LayerPicOrderCntVal may be derived independently of the picture order count value (PicOrderCntVal), which may be the same across all pictures of an access unit. Hence, if a base-layer picture, such as a base-layer IDR picture, resets the picture order count value, the PicOrderCntVal for the enhancement-layer pictures of the same access unit may also be reset. The variable LayerPicOrderCntVal for an enhancement-layer picture is derived independently of PicOrderCntVal values of other layers if there is a preceding picture, in decoding order, in the same enhancement layer and no layer-wise start-up process has been initiated in between. The variable LayerPicOrderCntVal may be derived for example as follows:

- If the current picture is an IRAP picture with NoRaslOutputFlag equal to 1 or FirstPicInLayerDecodedFlag[ nuh\_layer\_id ] is equal to 0, LayerPicOrderCntVal is set equal to PicOrderCntVal, where PicOrderCntVal is the picture order count value of the current picture.
- Otherwise, the following applies:
  - o Let prevTid0Pic be the previous picture in decoding order that has TemporalId equal to 0 and nuh\_layer\_id equal to nuh\_layer\_id of the current picture and that is not a RASL picture, a RADL picture, or a sub-layer non-reference picture.
  - o Let prevLayerPicOrderCnt be equal to LayerPicOrderCntVal of prevTid0Pic.
  - o Let prevPicOrderCntVal be equal to PicOrderCntVal of prevTid0Pic.
  - o If layer\_delta\_poc\_present\_flag is equal to 1, LayerPicOrderCntVal is set equal to prevLayerPicOrderCnt + layer\_delta\_pic\_order\_cnt\_minus1 + 1.
  - o Otherwise, LayerPicOrderCntVal is set equal to prevLayerPicOrderCnt + PicOrderCntVal - prevPicOrderCntVal.

[0363] In the presently described embodiment, the difference of POCs between two pictures of the same enhancement layer may be derived as the difference of their layer-wise POC values. The difference of POCs may be used for example in the inter prediction process, such as for deriving motion parameters in the merge mode and motion vector prediction. The difference of POCs of two pictures of the same enhancement layer may be specified as a function DiffPicOrderCnt( picA, picB ) as follows:

$$\text{DiffPicOrderCnt}( \text{picA}, \text{picB} ) = \text{layerPicOrderCnt}( \text{picA} ) - \text{layerPicOrderCnt}( \text{picB} ),$$

where the function layerPicOrderCnt( picX ) is specified as follows:

layerPicOrderCnt( picX ) = LayerPicOrderCntVal of the picture picX

[0364] In the presently described embodiment, the all pictures (in any layers) of a first coded video sequence may be required to precede in output order all pictures (in any layers) of a second coded video sequence, when the second coded video sequence follows, in decoding order, the first coded video sequence. The decoder and/or the HRD may therefore determine the output order per each coded video sequence using the picture order count values of the pictures and may determine the relative output order of two coded video sequences to be the same as their relative decoding order.

[0365] The draft SHVC design enables to have layers that cover a different area in the field of view. For example, an enhancement layer picture may enhance a region of the reference layer picture rather the reference base layer picture. An enhancement layer need not apply inter-layer prediction. An region of interest (ROI) could be enhanced not by utilizing inter-layer prediction but instead simulcasting the base layer picture and the enhancement layer picture, where the EL picture corresponds to a region of the base layer picture. This scenario would allow ROI enhancement but with low complexity decoding as the entire BL picture need not be decoded. Mixed chroma format coding would be a useful feature for many applications, such as wireless displays. In mixed chroma format coding, a certain region of video is coded in 4:2:0 format and other regions in 4:4:4 format. This use-case could be supported by coding the base layer in 4:2:0 format and the 4:4:4 region could be coded with the auxiliary chroma layers.

[0366] In an embodiment, which may be applied independently of or together with other embodiments, an output picture overlaying process is included or connected with the output picture cropping and/or picture output process of the decoding process and/or HRD. The output picture overlaying process may take as input two or more pictures from different layers that represent the same output time or the same output order. Said two or more pictures may for example reside on different scalability layers in the same access unit and both scalability layers may be output layers in an output layer set. In the embodiment, the encoder or another entity may indicate in the bitstream, for example in a sequence-level syntax structure such as VPS, that said scalability layers are subject to the output picture overlaying process (rather than outputting decoded pictures of these layers separately). In the output picture overlaying process, scaled reference offsets and/or sampling grid offset and/or alike are applied to locate said two or more pictures of said scalability layers on the same coordinates. A z-order or an overlaying order may be indicated by the encoder or another entity in the bitstream or may be inferred for example to be an ascending layer identifier order. Scaled reference offsets, which may be indicated e.g. by the encoder e.g. in SPS and/or VPS, may indicate the top, left, bottom and right offset values between a current picture of a current layer and a resampled reference layer picture, where the reference layer can be for example a layer preceding the current layer in the a z-order. The output picture of the output picture overlaying process may be formed by locating the sample arrays of the two or more pictures in the z-order on top of each in such a manner

that the sample array later in the z-order covers or replaces the sample values in collocated positions of the sample arrays earlier in the z-order. For example, an enhancement-layer picture representing a region of interest (ROI) of the base layer would replace that ROI area of the base layer picture in the output overlaying process. In addition to overlaying said two or more pictures, the process may include aligning the decoded  
5 representations of said two or more pictures. For example, if one picture is represented by the YUV 4:2:0 chroma format and the other one, later in the z-order, is represented by the YUV 4:4:4 chroma format, the first one may be upsampled to YUV 4:4:4 as part of the process. Likewise, if one picture is represented by a first color gamut or format, such as ITU-R BT.709, and another one, later in the z-order, is represented  
10 by a second color gamut or format, such as ITU-R BT.2020, the first one may be converted to the second color gamut or format as part of the process. Eventually, when all of the said two or more pictures are overlaid and processed as described above, the resulting picture forms the output picture and is output instead of said two or more pictures. A conformance check point may be specified additionally or alternatively to be the output of the output picture overlaying process (in addition to or rather than the  
cropped decoded two or more decoded pictures).

15 [0367] In the above, some embodiments have been described in relation to encoding indications, syntax elements, and/or syntax structures into a bitstream or into a coded video sequence and/or decoding indications, syntax elements, and/or syntax structures from a bitstream or from a coded video sequence. It needs to be understood, however, that embodiments could be realized when encoding indications, syntax elements, and/or syntax structures into a syntax structure or a data unit that is external from a bitstream or  
20 a coded video sequence comprising video coding layer data, such as coded slices, and/or decoding indications, syntax elements, and/or syntax structures from a syntax structure or a data unit that is external from a bitstream or a coded video sequence comprising video coding layer data, such as coded slices.

[0368] In the above, the example embodiments have been described with the help of syntax of the bitstream. It needs to be understood, however, that the corresponding structure and/or computer program  
25 may reside at the encoder for generating the bitstream and/or at the decoder for decoding the bitstream. Likewise, where the example embodiments have been described with reference to an encoder, it needs to be understood that the resulting bitstream and the decoder have corresponding elements in them. Likewise, where the example embodiments have been described with reference to a decoder, it needs to be understood that the encoder has structure and/or computer program for generating the bitstream to be decoded by the  
30 decoder.

[0369] Where example embodiments have been described with reference to a splicer, it needs to be understood that a splicer could likewise be an encoder, a middle-box, or any other entity that creates or modifies a coded video bitstream.

[0370] In the above, some example embodiments have been described with reference to a container file. It needs to be understood that embodiments could be similarly realized with reference to payloads (of a transmission protocol or format), such as segments in dynamic adaptive streaming over HTTP or RTP payloads. In the above, some example embodiments have been described with reference to a file creator. It  
5 needs to be understood that embodiments could be similarly realized with reference to packetizers or transmitters. In the above, some example embodiments have been described with reference to a file parser or player. It needs to be understood that embodiments could be similarly realized with reference to a depacketizer or receiver.

[0371] In the above, some embodiments have been described with reference to an enhancement layer and a reference layer, where the reference layer may be for example a base layer. In the above, some  
10 enhancements have been described with reference to an enhancement layer and a base layer, where the base layer may be considered to be any reference layer of the enhancement layer.

[0372] It needs to be understood that embodiments may be applicable to any types of layered coding, for example for multiview coding, quality scalability, spatial scalability, and for multiview video plus depth  
15 coding.

[0373] Embodiments of the present invention may be implemented in software, hardware, application logic or a combination of software, hardware and application logic. In an example embodiment, the application logic, software or an instruction set is maintained on any one of various conventional computer-readable media. In the context of this document, a “computer-readable medium” may be any media or  
20 means that can contain, store, communicate, propagate or transport the instructions for use by or in connection with an instruction execution system, apparatus, or device, such as a computer, with one example of a computer described and depicted in Figures 1 and 2. A computer-readable medium may comprise a computer-readable storage medium that may be any media or means that can contain or store the instructions for use by or in connection with an instruction execution system, apparatus, or device, such  
25 as a computer.

[0374] If desired, the different functions discussed herein may be performed in a different order and/or concurrently with each other. Furthermore, if desired, one or more of the above-described functions may be optional or may be combined.

[0375] Although the above examples describe embodiments of the invention operating within a codec  
30 within an electronic device, it would be appreciated that the invention as described below may be implemented as part of any video codec. Thus, for example, embodiments of the invention may be implemented in a video codec which may implement video coding over fixed or wired communication paths.

[0376] Thus, user equipment may comprise a video codec such as those described in embodiments of the invention above. It shall be appreciated that the term user equipment is intended to cover any suitable type of wireless user equipment, such as mobile telephones, portable data processing devices or portable web browsers.

5 [0377] Furthermore elements of a public land mobile network (PLMN) may also comprise video codecs as described above.

[0378] In general, the various embodiments of the invention may be implemented in hardware or special purpose circuits, software, logic or any combination thereof. For example, some aspects may be implemented in hardware, while other aspects may be implemented in firmware or software which may be  
10 executed by a controller, microprocessor or other computing device, although the invention is not limited thereto. While various aspects of the invention may be illustrated and described as block diagrams, flow charts, or using some other pictorial representation, it is well understood that these blocks, apparatuses, systems, techniques or methods described herein may be implemented in, as non-limiting examples, hardware, software, firmware, special purpose circuits or logic, general purpose hardware or controller or  
15 other computing devices, or some combination thereof.

[0379] The embodiments of this invention may be implemented by computer software executable by a data processor of the mobile device, such as in the processor entity, or by hardware, or by a combination of software and hardware. Further in this regard it should be noted that any blocks of the logic flow as in the  
20 Figures may represent program steps, or interconnected logic circuits, blocks and functions, or a combination of program steps and logic circuits, blocks and functions. The software may be stored on such physical media as memory chips, or memory blocks implemented within the processor, magnetic media such as hard disk or floppy disks, and optical media such as for example DVD and the data variants thereof, CD.

[0380] The various embodiments of the invention can be implemented with the help of computer  
25 program code that resides in a memory and causes the relevant apparatuses to carry out the invention. For example, a terminal device may comprise circuitry and electronics for handling, receiving and transmitting data, computer program code in a memory, and a processor that, when running the computer program code, causes the terminal device to carry out the features of an embodiment. Yet further, a network device may comprise circuitry and electronics for handling, receiving and transmitting data, computer program code in  
30 a memory, and a processor that, when running the computer program code, causes the network device to carry out the features of an embodiment.

[0381] The memory may be of any type suitable to the local technical environment and may be implemented using any suitable data storage technology, such as semiconductor based memory devices, magnetic memory devices and systems, optical memory devices and systems, fixed memory and removable

memory. The data processors may be of any type suitable to the local technical environment, and may include one or more of general purpose computers, special purpose computers, microprocessors, digital signal processors (DSPs) and processors based on multi core processor architecture, as non-limiting examples.

5 [0382] Embodiments of the inventions may be practiced in various components such as integrated circuit modules. The design of integrated circuits is by and large a highly automated process. Complex and powerful software tools are available for converting a logic level design into a semiconductor circuit design ready to be etched and formed on a semiconductor substrate.

10 [0383] Programs, such as those provided by Synopsys Inc., of Mountain View, California and Cadence Design, of San Jose, California automatically route conductors and locate components on a semiconductor chip using well established rules of design as well as libraries of pre stored design modules. Once the design for a semiconductor circuit has been completed, the resultant design, in a standardized electronic format (e.g., Opus, GDSII, or the like) may be transmitted to a semiconductor fabrication facility or "fab" for fabrication.

15 [0384] The foregoing description has provided by way of exemplary and non-limiting examples a full and informative description of the exemplary embodiment of this invention. However, various modifications and adaptations may become apparent to those skilled in the relevant arts in view of the foregoing description, when read in conjunction with the accompanying drawings and the appended claims. However, all such and similar modifications of the teachings of this invention will still fall within the scope  
20 of this invention.

[0385] According to a first example there is provided a method comprising:

- encoding pictures into a bitstream, the bitstream comprising at least two scalability layers and pictures being associated with access units;
- encoding, into the bitstream, a first indication, which indicates an output layer; and
- 25 – encoding, into the bitstream, a second indication, which indicates at least one alternative output layer
- wherein the first indication and the second indication jointly have an impact that a first picture of the at least one alternative output layer is output by a decoding process of the bitstream when no picture of the output layer is in an access unit containing the first picture.

30 [0386] In an embodiment, the method further comprises:

- associating layers with one or more scalability dimensions; and
- associating layers with different value sets of the one or more scalability dimensions;
- wherein the second indication is independent of the one or more scalability dimensions.

[0387] According to a second example, there is provided an apparatus comprising at least one processor and at least one memory, said at least one memory stored with code thereon, which when executed by said at least one processor, causes an apparatus to perform at least the following:

- encode pictures into a bitstream, the bitstream comprising at least two scalability layers and pictures being associated with access units;
- encode, into the bitstream, a first indication, which indicates an output layer; and
- encode, into the bitstream, a second indication, which indicates at least one alternative output layer
- wherein the first indication and the second indication jointly have an impact that a first picture of the at least one alternative output layer is output by a decoding process of the bitstream when no picture of the output layer is in an access unit containing the first picture.

[0388] In an embodiment of the apparatus, said at least one memory stored with code thereon, which when executed by said at least one processor, causes the apparatus to

- associate layers with one or more scalability dimensions; and
- associate layers with different value sets of the one or more scalability dimensions;
- wherein the second indication is independent of the one or more scalability dimensions.

[0389] According to a third example, there is provided a computer program product embodied on a non-transitory computer readable medium, comprising computer program code configured to, when executed on at least one processor, cause an apparatus or a system to:

- encode pictures into a bitstream, the bitstream comprising at least two scalability layers and pictures being associated with access units;
- encode, into the bitstream, a first indication, which indicates an output layer; and
- encode, into the bitstream, a second indication, which indicates at least one alternative output layer
- wherein the first indication and the second indication jointly have an impact that a first picture of the at least one alternative output layer is output by a decoding process of the bitstream when no picture of the output layer is in an access unit containing the first picture.

[0390] In an embodiment of the computer program product, said at least one memory stored with code thereon, which when executed by said at least one processor, causes the apparatus to

- associate layers with one or more scalability dimensions; and
- associate layers with different value sets of the one or more scalability dimensions;
- wherein the second indication is independent of the one or more scalability dimensions.

[0391] According to a fourth example, there is provided a method comprising

- decoding pictures from a bitstream, the bitstream comprising at least two scalability layers and pictures being associated with access units;
- decoding, from the bitstream, a first indication, which indicates an output layer;
- decoding, from the bitstream, a second indication, which indicates at least one alternative output layer; and
- outputting a first picture of the at least one alternative output layer when no picture of the output layer is in a first access unit containing the first picture.

5 [0392] In an embodiment, the method further comprises:

- outputting a second picture of the output layer, the second picture being associated with a second access unit.

10 [0393] In an embodiment, the method further comprises:

- associating layers with one or more scalability dimensions; and
- associating layers with different value sets of the one or more scalability dimensions;
- wherein the second indication is independent of the one or more scalability dimensions.

15 [0394] According to a fifth example, there is provided an apparatus comprising at least one processor and at least one memory, said at least one memory stored with code thereon, which when executed by said at least one processor, causes an apparatus to perform at least the following:

- decode pictures from a bitstream, the bitstream comprising at least two scalability layers and pictures being associated with access units;
- decode, from the bitstream, a first indication, which indicates an output layer;
- decode, from the bitstream, a second indication, which indicates at least one alternative output layer; and
- output a first picture of the at least one alternative output layer when no picture of the output layer is in a first access unit containing the first picture.

20 [0395] In an embodiment of the apparatus, said at least one memory stored with code thereon, which when executed by said at least one processor, causes the apparatus to

- output a second picture of the output layer, the second picture being associated with a second access unit.

25 [0396] In an embodiment of the apparatus, said at least one memory stored with code thereon, which when executed by said at least one processor, causes the apparatus to

- associate layers with one or more scalability dimensions; and
- associate layers with different value sets of the one or more scalability dimensions;
- wherein second indication is independent of the one or more scalability dimensions.

[0397] According to a sixth example, there is provided a computer program product embodied on a non-transitory computer readable medium, comprising computer program code configured to, when executed on at least one processor, cause an apparatus or a system to:

- decode pictures from a bitstream, the bitstream comprising at least two scalability layers and pictures being associated with access units;
- decode, from the bitstream, a first indication, which indicates an output layer;
- decode, from the bitstream, a second indication, which indicates at least one alternative output layer; and
- output a first picture of the at least one alternative output layer when no picture of the output layer is in a first access unit containing the first picture.

[0398] In an embodiment of the computer program product, said at least one memory stored with code thereon, which when executed by said at least one processor, causes the apparatus to

- output a second picture of the output layer, the second picture being associated with a second access unit.

[0399] In an embodiment of the computer program product, said at least one memory stored with code thereon, which when executed by said at least one processor, causes the apparatus to

- associate layers with one or more scalability dimensions; and
- associate layers with different value sets of the one or more scalability dimensions;
- wherein second indication is independent of the one or more scalability dimensions.

[0400] According to a seventh example there is provided a method comprising:

- encoding pictures into a bitstream, the bitstream comprising at least two scalability layers;
- encoding, into the bitstream, at least one first indication that provides one or more properties for an output layer set and associates an index for the output layer set;
- encoding, into the bitstream, a first picture of a first layer;
- encoding, into a picture-level structure within the bitstream, at least one second indication that associates the first picture to the output layer set;
- encoding, into the bitstream, a second picture of a second layer;
- encoding, into a picture-level structure within the bitstream, at least one third indication that associates the second picture to the output layer set;
- wherein the first, the second and the third indications jointly have an impact that the pictures being associated with the output layer set are output by a decoding process of the bitstream wherein the output layer set has been selected.

[0401] According to an eighth example, there is provided an apparatus comprising at least one processor and at least one memory, said at least one memory stored with code thereon, which when executed by said at least one processor, causes an apparatus to perform at least the following:

- encode pictures into a bitstream, the bitstream comprising at least two scalability layers;
- 5 – encode, into the bitstream, at least one first indication that provides one or more properties for an output layer set and associates an index for the output layer set;
- encode, into the bitstream, a first picture of a first layer;
- encode, into a picture-level structure within the bitstream, at least one second indication that associates the first picture to the output layer set;
- 10 – encode, into the bitstream, a second picture of a second layer;
- encode, into a picture-level structure within the bitstream, at least one third indication that associates the second picture to the output layer set;
- wherein the first, the second and the third indications jointly have an impact that the pictures being associated with the output layer set are output by a decoding process of the bitstream wherein the output layer set has been selected.
- 15

[0402] According to a ninth example, there is provided a computer program product embodied on a non-transitory computer readable medium, comprising computer program code configured to, when executed on at least one processor, cause an apparatus or a system to:

- encode pictures into a bitstream, the bitstream comprising at least two scalability layers;
- 20 – encode, into the bitstream, at least one first indication that provides one or more properties for an output layer set and associates an index for the output layer set;
- encode, into the bitstream, a first picture of a first layer;
- encode, into a picture-level structure within the bitstream, at least one second indication that associates the first picture to the output layer set;
- 25 – encode, into the bitstream, a second picture of a second layer;
- encode, into a picture-level structure within the bitstream, at least one third indication that associates the second picture to the output layer set;
- wherein the first, the second and the third indications jointly have an impact that the pictures being associated with the output layer set are output by a decoding process of the bitstream wherein the output layer set has been selected.
- 30

[0403] According to a tenth example there is provided a method comprising:

- decoding pictures from a bitstream, the bitstream comprising at least two scalability layers;

- decoding, from the bitstream, at least one first indication that provides one or more properties for an output layer set and associates an index for the output layer set;
- determining that the output layer set is used in outputting from the decoder;
- decoding, from the bitstream, a first picture of a first layer;
- 5    – decoding, from a picture-level structure within the bitstream, at least one second indication that associates the first picture to the output layer set;
- decoding, from the bitstream, a second picture of a second layer;
- decoding, from a picture-level structure within the bitstream, at least one third indication that associates the second picture to the output layer set;
- 10   – outputting the first picture and the second picture, as they are pictures that are associated with the output layer set.

[0404] According to a eleventh example, there is provided an apparatus comprising at least one processor and at least one memory, said at least memory stored with code thereon, which when executed by said at least one processor, causes an apparatus to perform at least the following:

- 15    – decode pictures from a bitstream, the bitstream comprising at least two scalability layers;
- decode, from the bitstream, at least one first indication that provides one or more properties for an output layer set and associates an index for the output layer set;
- determine that the output layer set is used in outputting from the decoder;
- decode, from the bitstream, a first picture of a first layer;
- 20    – decode, from a picture-level structure within the bitstream, at least one second indication that associates the first picture to the output layer set;
- decode, from the bitstream, a second picture of a second layer;
- decode, from a picture-level structure within the bitstream, at least one third indication that associates the second picture to the output layer set;
- 25    – output the first picture and the second picture, as they are pictures that are associated with the output layer set.

[0405] According to a twelfth example, there is provided a computer program product embodied on a non-transitory computer readable medium, comprising computer program code configured to, when executed on at least one processor, cause an apparatus or a system to:

- 30    – decode pictures from a bitstream, the bitstream comprising at least two scalability layers;
- decode, from the bitstream, at least one first indication that provides one or more properties for an output layer set and associates an index for the output layer set;
- determine that the output layer set is used in outputting from the decoder;

- decode, from the bitstream, a first picture of a first layer;
  - decode, from a picture-level structure within the bitstream, at least one second indication that associates the first picture to the output layer set;
  - decode, from the bitstream, a second picture of a second layer;
- 5
- decode, from a picture-level structure within the bitstream, at least one third indication that associates the second picture to the output layer set;
  - output the first picture and the second picture, as they are pictures that are associated with the output layer set.

**CLAIMS:**

1. A method comprising:
  - encoding a scalable nesting supplemental enhancement information (SEI) message;
  - associating the scalable nesting SEI message with a bitstream comprising at least two layers and with a first intra random access point (IRAP) picture within the bitstream;
  - including a buffering period SEI message in the scalable nesting SEI message;
  - indicating, in the scalable nesting SEI message, that the buffering period SEI message applies to more than one layer of the associated bitstream;
  - indicating, in the buffering period SEI message, at least one buffering parameter applying when cross-layer random access skipped (CL-RAS) pictures of the first IRAP picture are not present in the bitstream, where CL-RAS picture in a particular layer is a picture preceding, in decoding order, an IRAP picture in that particular layer from which the decoding of that particular layer can be initialized.
2. A method according to claim 1, wherein pictures and SEI message are included in the bitstream, an access unit comprises consecutive data units in the bitstream, and said associating comprises including the scalable nesting SEI message in an access unit including the first IRAP picture.
3. A method according to claim 1, wherein the at least one buffering parameter comprises an initial coded picture buffering delay.
4. A method according to claim 1, wherein the at least one buffering parameter is in accordance with a hypothetical reference decoder model.
5. An apparatus comprising at least one processor and at least one memory, said at least one memory stored with code thereon, which when executed by said at least one processor, causes an apparatus to perform at least the following:
  - encoding a scalable nesting supplemental enhancement information (SEI) message;
  - associating the scalable nesting SEI message with a bitstream comprising at least two layers and with a first intra random access point (IRAP) picture within the bitstream;
  - including a buffering period SEI message in the scalable nesting SEI message;

- indicating, in the scalable nesting SEI message, that the buffering period SEI message applies to more than one layer of the associated bitstream;
  - indicating, in the buffering period SEI message, at least one buffering parameter applying when cross-layer random access skipped (CL-RAS) pictures of the first IRAP picture are not present in the bitstream, where CL\_RAS picture in a particular layer is a picture preceding, in decoding order, an IRAP picture in that particular layer from which the decoding of that particular layer can be initialized.
6. An apparatus according to claim 5, wherein pictures and SEI message are included in the bitstream, an access unit comprises consecutive data units in the bitstream, and said associating comprises including the scalable nesting SEI message in an access unit including the first IRAP picture.
7. An apparatus according to claim 5, wherein the at least one buffering parameter comprises an initial coded picture buffering delay.
8. An apparatus according to claim 5, wherein the at least one buffering parameter is in accordance with a hypothetical reference decoder model.
9. An apparatus comprising at least:
- processing means and memory means;
  - means for encoding a scalable nesting supplemental enhancement information (SEI) message;
  - means for associating the scalable nesting SEI message with a bitstream comprising at least two layers and with a first intra random access point (IRAP) picture within the bitstream;
  - means for including a buffering period SEI message in the scalable nesting SEI message;
  - means for indicating, in the scalable nesting SEI message, that the buffering period SEI message applies to more than one layer of the associated bitstream;
  - means for indicating, in the buffering period SEI message, at least one buffering parameter applying when cross-layer random access skipped (CL-RAS) pictures of the first IRAP picture are not present in the bitstream, where CL\_RAS picture in a particular layer is a picture preceding, in decoding order, an IRAP picture in that particular layer from which the decoding of that particular layer can be initialized.

10. A computer program product embodied on a non-transitory computer readable medium, comprising computer program code configured to, when executed on at least one processor, cause an apparatus or a system to:
- encode a scalable nesting supplemental enhancement information (SEI) message;
  - associate the scalable nesting SEI message with a bitstream comprising at least two layers and with a first intra random access point (IRAP) picture within the bitstream;
  - include a buffering period SEI message in the scalable nesting SEI message;
  - indicate, in the scalable nesting SEI message, that the buffering period SEI message applies to more than one layer of the associated bitstream;
  - indicate, in the buffering period SEI message, at least one buffering parameter applying when cross-layer random access skipped (CL-RAS) pictures of the first IRAP picture are not present in the bitstream, where CL\_RAS picture in a particular layer is a picture preceding, in decoding order, an IRAP picture in that particular layer from which the decoding of that particular layer can be initialized.




Fig. 2




Fig. 1




FIG. 3




Fig.4

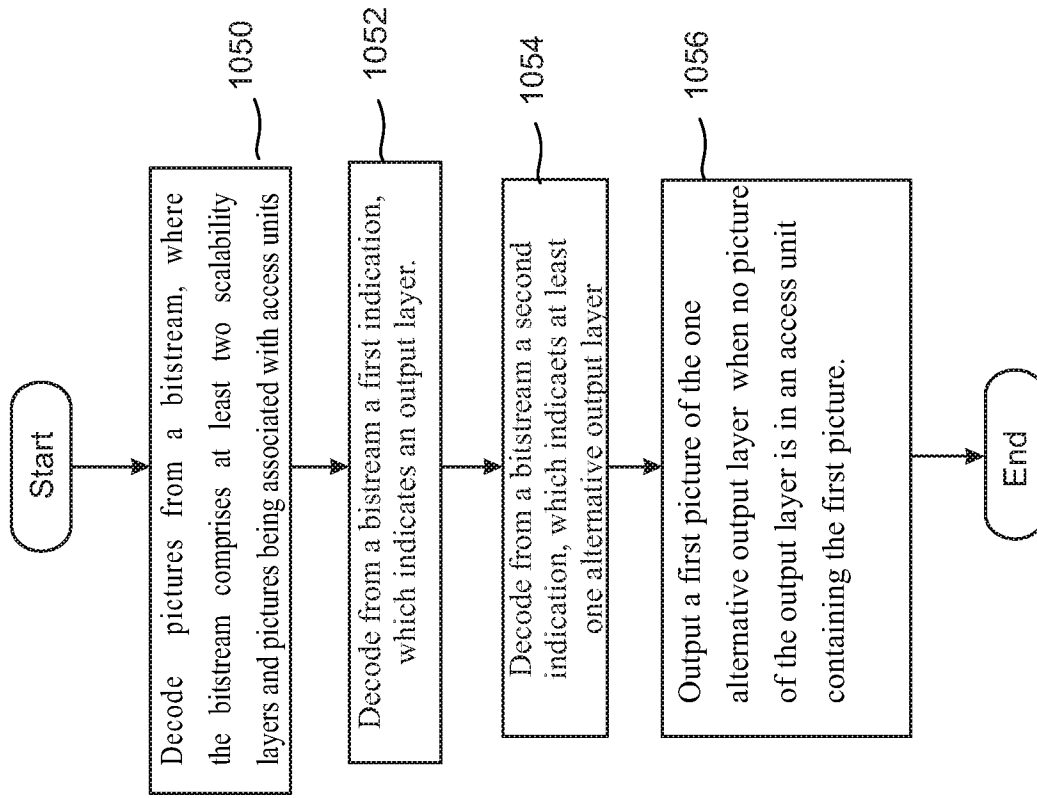


Fig.6

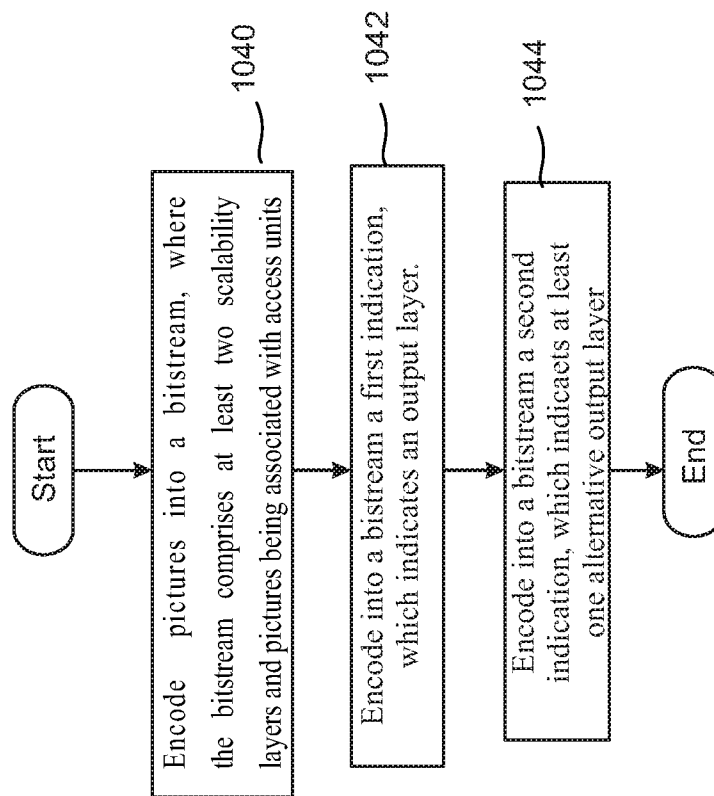


Fig.5

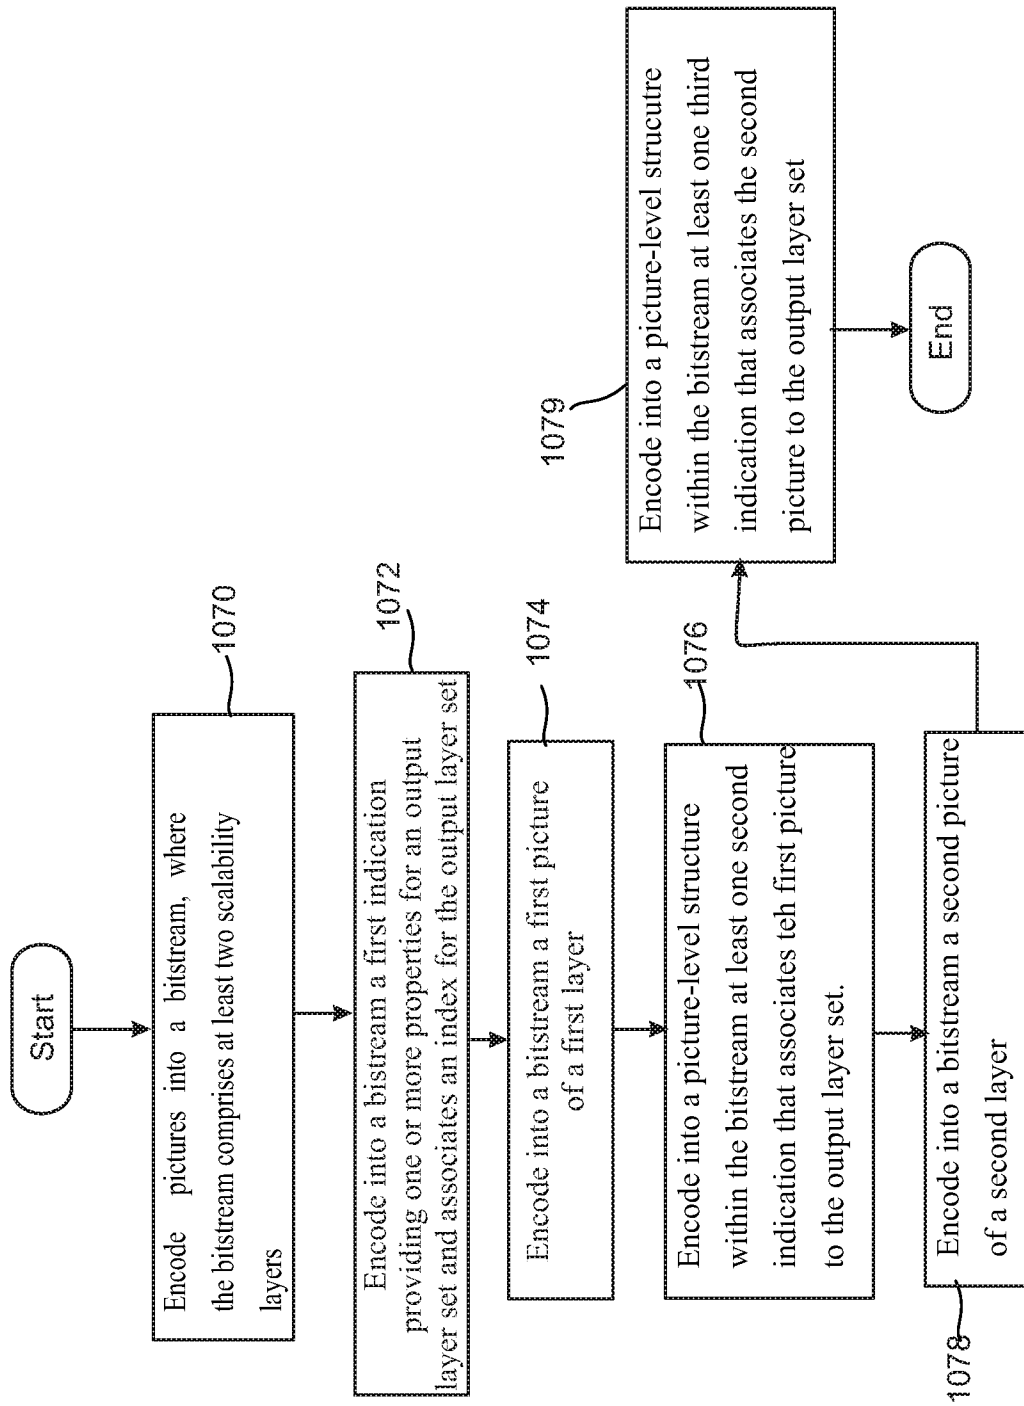


Fig. 7

Fig. 8

