



(51) International Patent Classification:

H04N 7/26 (2006.01)

(21) International Application Number:

PCT/GB2013/051074

(22) International Filing Date:

26 April 2013 (26.04.2013)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

1207459.7 26 April 2012 (26.04.2012) GB

1211074.8 22 June 2012 (22.06.2012) GB

(71) Applicant: SONY CORPORATION [JP/JP]; 1-7-1 Kon-

an, Minato-ku, Tokyo 108-0075 (JP).

(71) Applicant (for JP only): SONY EUROPE LIMITED

[GB/GB]; The Heights, Brooklands, Weybridge, Surrey KT13 0XW (GB).

(72) Inventors: GAMEI, James Alexander; 34 Norbiton Hall,

London Road, Kingston-upon-Thames, Surrey KT2 6RB (GB). SAUNDERS, Nicholas Ian; 21 Mongers Piece,

Chineham, Basingstoke, Hampshire RG24 8RL (GB).

SHARMAN, Karl James; 6 Tadburn Close, Chandler's

Ford, Eastleigh, Hampshire SO53 2NF (GB). SILCOCK,

Paul James; 50 Sevenfields, Highworth, Swindon SN6

7NF (GB).

(74) Agent: TURNER, James; 120 Holborn, London EC1N

2DY (GB).

(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, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, 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, 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, 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, ML, MR, NE, SN, TD, TG).

Published:

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

(54) Title: QUANTISATION PARAMETER SELECTION FOR DIFFERENT COLOUR SAMPLING FORMATS

Index	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
Value	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25

Index	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51
Value	26	27	28	29	29	30	31	32	32	33	34	34	35	35	36	36	37	37	38	38	38	39	39	39	39	39

(4:2:0)

FIG. 29A

Index	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
Value	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25

Index	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51
Value	26	27	28	29	30	31	32	33	33	34	35	36	37	38	39	40	40	41	42	42	43	44	44	45	45	45

(4:2:2)

FIG. 29B

(57) Abstract: A video decoding apparatus operable to dequantise blocks of frequency-transformed luminance and chrominance component video data in a 4:4:4 or a 4:2:2 format according to a selected quantisation parameter which defines a quantisation step size comprises a controller configured to define a quantisation parameter association between luminance and chrominance quantisation parameters, where the association is such that a chrominance quantisation step size is less than a corresponding luminance quantisation step size for the 4:2:2 format but equal to the corresponding luminance quantisation step size for the 4:4:4 format, over at least a predetermined range of quantization parameters; and a multiplier configured to multiply each component of the frequency-transformed data by a respective value derived from the respective quantisation step size.



QUANTISATION PARAMETER SELECTION FOR DIFFERENT COLOUR SAMPLING FORMATS

CROSS REFERENCE TO RELATED APPLICATIONS

The present application claims the benefit of the earlier filing date of GB1211074.8 and GB 1207459.7 filed in the United Kingdom Intellectual Property Office on 22 June 2012 and 26 April 2012 respectively, the entire contents of which applications are incorporated herein by reference.

BACKGROUND

Field

This disclosure relates to data encoding and decoding.

Description of Related Art

The "background" description provided herein is for the purpose of generally presenting the context of the disclosure. Work of the presently named inventors, to the extent it is described in this background section, as well as aspects of the description which may not otherwise qualify as prior art at the time of filing, is neither expressly or impliedly admitted as prior art against the present disclosure.

There are several video data encoding and decoding systems which involve transforming video data into a frequency domain representation, quantising the frequency domain coefficients and then applying some form of entropy encoding to the quantised coefficients. This can achieve compression of the video data. A corresponding decoding or decompression technique is applied to recover a reconstructed version of the original video data.

Current video codecs (coder-decoders) such as those used in H.264/MPEG-4 Advanced Video Coding (AVC) achieve data compression primarily by only encoding the differences between successive video frames. These codecs use a regular array of so-called macroblocks, each of which is used as a region of comparison with a corresponding macroblock in a previous video frame, and the image region within the macroblock is then encoded according to the degree of motion found between the corresponding current and previous macroblocks in the video sequence, or between neighbouring macroblocks within a single frame of the video sequence.

High Efficiency Video Coding (HEVC), also known as H.265 or MPEG-H Part 2, is a proposed successor to H.264/MPEG-4 AVC. It is intended for HEVC to improve video quality and double the data compression ratio compared to H.264, and for it to be scalable from 128 × 96 to 7680 × 4320 pixels resolution, roughly equivalent to bit rates ranging from 128kbit/s to 800Mbit/s.

In HEVC a so-called 4:2:0 block structure is proposed for consumer equipment, in which the amount of data used in each chroma channel is one quarter that in the luma channel. This is

because subjectively people are more sensitive to brightness variations than to colour variations, and so it is possible to use greater compression and/or less information in the colour channels without a subjective loss of quality.

5 HEVC replaces the macroblocks found in existing H.264 and MPEG standards with a more flexible scheme based upon coding units (CUs), which are variable size structures.

Consequently, when encoding the image data in video frames, the CU sizes can be selected responsive to the apparent image complexity or detected motion levels, instead of using uniformly distributed macroblocks. Consequently far greater compression can be achieved in regions with little motion between frames and with little variation within a frame, 10 whilst better image quality can be preserved in areas of high inter-frame motion or image complexity.

Each CU contains one or more variable-block-sized prediction units (PUs) of either intra-picture or inter-picture prediction type, and one or more transform units (TUs) which contain coefficients for spatial block transform and quantisation.

15 Moreover, PU and TU blocks are provided for each of three channels; luma (Y), being a luminance or brightness channel, and which may be thought of as a greyscale channel, and two colour difference or chrominance (chroma) channels; Cb and Cr. These channels provide the colour for the greyscale image of the luma channel. The terms Y, luminance and luma are used interchangeably in this description, and similarly the terms Cb and Cr, chrominance and 20 chroma, are used interchangeably as appropriate, noting that chrominance or chroma can be used generically for "one or both of Cr and Cb", whereas when a specific chrominance channel is being discussed it will be identified by the term Cb or Cr.

Generally PUs are considered to be channel independent, except that a PU has a luma part and a chroma part. Generally, this means that the samples forming part of the PU for each 25 channel represent the same region of the image, so that there is a fixed relationship between the PUs between the three channels. For example, for 4:2:0 video, an 8x8 PU for Luma always has a corresponding 4x4 PU for chroma, with the chroma parts of the PU representing the same area as the luma part, but containing a smaller number of pixels because of the subsampled nature of the 4:2:0 chroma data compared to the luma data in 4:2:0 video. (In 4:4:4 the 30 sampling rates are the same and so the same PU structures can be used for luma and chroma samples). The two chroma channels share intra-prediction information; and the three channels share inter-prediction information. Similarly, the TU structure also has a fixed relationship between the three channels.

However, for professional broadcast and digital cinema equipment, it is desirable to have 35 less compression (or more information) in the chroma channels, and this may affect how current and proposed HEVC processing operates.

SUMMARY

The present disclosure addresses or mitigates problems arising from this processing.

Respective aspects and features of the present disclosure are defined in the appended claims.

5 It is to be understood that both the foregoing general description and the following detailed description are exemplary, but are not restrictive, of the present technology.

BRIEF DESCRIPTION OF THE DRAWINGS

A more complete appreciation of the disclosure and many of the attendant advantages thereof will be readily obtained as the same becomes better understood by reference to the
10 following detailed description when considered in connection with the accompanying drawings, wherein:

Figure 1 schematically illustrates an audio/video (A/V) data transmission and reception system using video data compression and decompression;

Figure 2 schematically illustrates a video display system using video data
15 decompression;

Figure 3 schematically illustrates an audio/video storage system using video data compression and decompression;

Figure 4 schematically illustrates a video camera using video data compression;

Figure 5 provides a schematic overview of a video data compression and
20 decompression apparatus;

Figure 6 schematically illustrates the generation of predicted images;

Figure 7 schematically illustrates a largest coding unit (LCU);

Figure 8 schematically illustrates a set of four coding units (CU);

Figures 9 and 10 schematically illustrate the coding units of Figure 8 sub-divided into
25 smaller coding units;

Figure 11 schematically illustrates an array of prediction units (PU);

Figure 12 schematically illustrates an array of transform units (TU);

Figure 13 schematically illustrates a partially-encoded image;

Figure 14 schematically illustrates a set of possible intra-prediction directions;

30 Figure 15 schematically illustrates a set of prediction modes;

Figure 16 schematically illustrates an up-right diagonal scan;

Figure 17 schematically illustrates a video compression apparatus;

Figure 18a and 18b schematically illustrate possible block sizes;

Figure 19 schematically illustrates the use of co-located information from chroma and
35 luma blocks;

Figure 20 schematically illustrates a situation in which co-located information from one chroma channel is used in respect of another chroma channel;

Figure 21 schematically illustrates pixels used for an LM-CHROMA mode;

Figure 22 schematically illustrates a set of luma prediction directions;

Figure 23 schematically illustrates the directions of Figure 22, as applied to a horizontally sparse chroma channel;

5 Figure 24 schematically illustrates the directions of Figure 22 mapped to a rectangular chroma pixel array;

Figures 25-28 schematically illustrate luma and chroma pixel interpolation;

Figures 29a and 29b schematically illustrates quantisation parameter tables for 4:2:0 and 4:2:2 respectively;

10 Figure 29c schematically illustrates further quantisation parameter tables; and

Figures 30 and 31 schematically illustrate quantisation variation tables.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

Referring now to the drawings, Figures 1-4 are provided to give schematic illustrations of apparatus or systems making use of the compression and/or decompression apparatus to be
15 described below in connection with embodiments of the present technology.

All of the data compression and/or decompression apparatus to be described below may be implemented in hardware, in software running on a general-purpose data processing apparatus such as a general-purpose computer, as programmable hardware such as an application specific integrated circuit (ASIC) or field programmable gate array (FPGA) or as
20 combinations of these. In cases where the embodiments are implemented by software and/or firmware, it will be appreciated that such software and/or firmware, and non-transitory data storage media by which such software and/or firmware are stored or otherwise provided, are considered as embodiments of the present technology.

Figure 1 schematically illustrates an audio/video data transmission and reception system
25 using video data compression and decompression.

An input audio/video signal 10 is supplied to a video data compression apparatus 20 which compresses at least the video component of the audio/video signal 10 for transmission along a transmission route 30 such as a cable, an optical fibre, a wireless link or the like. The compressed signal is processed by a decompression apparatus 40 to provide an output
30 audio/video signal 50. For the return path, a compression apparatus 60 compresses an audio/video signal for transmission along the transmission route 30 to a decompression apparatus 70.

The compression apparatus 20 and decompression apparatus 70 can therefore form one node of a transmission link. The decompression apparatus 40 and decompression
35 apparatus 60 can form another node of the transmission link. Of course, in instances where the transmission link is uni-directional, only one of the nodes would require a compression apparatus and the other node would only require a decompression apparatus.

Figure 2 schematically illustrates a video display system using video data decompression. In particular, a compressed audio/video signal 100 is processed by a decompression apparatus 110 to provide a decompressed signal which can be displayed on a display 120. The decompression apparatus 110 could be implemented as an integral part of the display 120, for example being provided within the same casing as the display device. Alternatively, the decompression apparatus 110 maybe provided as (for example) a so-called set top box (STB), noting that the expression "set-top" does not imply a requirement for the box to be sited in any particular orientation or position with respect to the display 120; it is simply a term used in the art to indicate a device which is connectable to a display as a peripheral device.

Figure 3 schematically illustrates an audio/video storage system using video data compression and decompression. An input audio/video signal 130 is supplied to a compression apparatus 140 which generates a compressed signal for storing by a store device 150 such as a magnetic disk device, an optical disk device, a magnetic tape device, a solid state storage device such as a semiconductor memory or other storage device. For replay, compressed data is read from the store device 150 and passed to a decompression apparatus 160 for decompression to provide an output audio/video signal 170.

It will be appreciated that the compressed or encoded signal, and a storage medium storing that signal, are considered as embodiments of the present technology.

Figure 4 schematically illustrates a video camera using video data compression. In Figure 4, an image capture device 180, such as a charge coupled device (CCD) image sensor and associated control and read-out electronics, generates a video signal which is passed to a compression apparatus 190. A microphone (or plural microphones) 200 generates an audio signal to be passed to the compression apparatus 190. The compression apparatus 190 generates a compressed audio/video signal 210 to be stored and/or transmitted (shown generically as a schematic stage 220).

The techniques to be described below relate primarily to video data compression and decompression. It will be appreciated that many existing techniques may be used for audio data compression in conjunction with the video data compression techniques which will be described, to generate a compressed audio/video signal. Accordingly, a separate discussion of audio data compression will not be provided. It will also be appreciated that the data rate associated with video data, in particular broadcast quality video data, is generally very much higher than the data rate associated with audio data (whether compressed or uncompressed). It will therefore be appreciated that uncompressed audio data could accompany compressed video data to form a compressed audio/video signal. It will further be appreciated that although the present examples (shown in Figures 1-4) relate to audio/video data, the techniques to be described below can find use in a system which simply deals with (that is to say, compresses,

decompresses, stores, displays and/or transmits) video data. That is to say, the embodiments can apply to video data compression without necessarily having any associated audio data handling at all.

Figure 5 provides a schematic overview of a video data compression and decompression apparatus.

A controller 343 controls the overall operation of the apparatus and, in particular when referring to a compression mode, controls the trial encoding processes (to be described below) by acting as a selector to select various modes of operation such as CU, PU and TU block sizes and whether the video data is to be encoded losslessly or otherwise.

Successive images of an input video signal 300 are supplied to an adder 310 and to an image predictor 320. The image predictor 320 will be described below in more detail with reference to Figure 6. The adder 310 in fact performs a subtraction (negative addition) operation, in that it receives the input video signal 300 on a "+" input and the output of the image predictor 320 on a "-" input, so that the predicted image is subtracted from the input image. The result is to generate a so-called residual image signal 330 representing the difference between the actual and projected images.

One reason why a residual image signal is generated is as follows. The data coding techniques to be described, that is to say the techniques which will be applied to the residual image signal, tend to work more efficiently when there is less "energy" in the image to be encoded. Here, the term "efficiently" refers to the generation of a small amount of encoded data; for a particular image quality level, it is desirable (and considered "efficient") to generate as little data as is practicably possible. The reference to "energy" in the residual image relates to the amount of information contained in the residual image. If the predicted image were to be identical to the real image, the difference between the two (that is to say, the residual image) would contain zero information (zero energy) and would be very easy to encode into a small amount of encoded data. In general, if the prediction process can be made to work reasonably well, the expectation is that the residual image data will contain less information (less energy) than the input image and so will be easier to encode into a small amount of encoded data.

The remainder of the apparatus acting as an encoder (to encode the residual or difference image) will now be described. The residual image data 330 is supplied to a transform unit 340 which generates a discrete cosine transform (DCT) representation of the residual image data. The DCT technique itself is well known and will not be described in detail here. There are however aspects of the techniques used in the present apparatus which will be described in more detail below, in particular relating to the selection of different blocks of data to which the DCT operation is applied. These will be discussed with reference to Figures 7-12 below.

The output of the transform unit 340, which is to say, a set of DCT coefficients for each transformed block of image data, is supplied to a quantiser 350. Various quantisation techniques are known in the field of video data compression, ranging from a simple multiplication by a quantisation scaling factor through to the application of complicated lookup tables under the control of a quantisation parameter. The general aim is twofold. Firstly, the quantisation process reduces the number of possible values of the transformed data. Secondly, the quantisation process can increase the likelihood that values of the transformed data are zero. Both of these can make the entropy encoding process, to be described below, work more efficiently in generating small amounts of compressed video data.

A data scanning process is applied by a scan unit 360. The purpose of the scanning process is to reorder the quantised transformed data so as to gather as many as possible of the non-zero quantised transformed coefficients together, and of course therefore to gather as many as possible of the zero-valued coefficients together. These features can allow so-called run-length coding or similar techniques to be applied efficiently. So, the scanning process involves selecting coefficients from the quantised transformed data, and in particular from a block of coefficients corresponding to a block of image data which has been transformed and quantised, according to a "scanning order" so that (a) all of the coefficients are selected once as part of the scan, and (b) the scan tends to provide the desired reordering. One example scanning order which can tend to give useful results is a so-called up-right diagonal scanning order.

The scanned coefficients are then passed to an entropy encoder (EE) 370. Again, various types of entropy encoding may be used. Two examples are variants of the so-called CABAC (Context Adaptive Binary Arithmetic Coding) system and variants of the so-called CAVLC (Context Adaptive Variable-Length Coding) system. In general terms, CABAC is considered to provide a better efficiency, and in some studies has been shown to provide a 10-20% reduction in the quantity of encoded output data for a comparable image quality compared to CAVLC. However, CAVLC is considered to represent a much lower level of complexity (in terms of its implementation) than CABAC. Note that the scanning process and the entropy encoding process are shown as separate processes, but in fact can be combined or treated together. That is to say, the reading of data into the entropy encoder can take place in the scan order. Corresponding considerations apply to the respective inverse processes to be described below. Note that the current HEVC documents under consideration at the time of filing no longer include the possibility of a CAVLC coefficient encoder.

The output of the entropy encoder 370, along with additional data (mentioned above and/or discussed below), for example defining the manner in which the predictor 320 generated the predicted image, provides a compressed output video signal 380.

However, a return path is also provided because the operation of the predictor 320 itself depends upon a decompressed version of the compressed output data.

The reason for this feature is as follows. At the appropriate stage in the decompression process (to be described below) a decompressed version of the residual data is generated. This decompressed residual data has to be added to a predicted image to generate an output image (because the original residual data was the difference between the input image and a predicted image). In order that this process is comparable, as between the compression side and the decompression side, the predicted images generated by the predictor 320 should be the same during the compression process and during the decompression process. Of course, at decompression, the apparatus does not have access to the original input images, but only to the decompressed images. Therefore, at compression, the predictor 320 bases its prediction (at least, for inter-image encoding) on decompressed versions of the compressed images.

The entropy encoding process carried out by the entropy encoder 370 is considered to be "lossless", which is to say that it can be reversed to arrive at exactly the same data which was first supplied to the entropy encoder 370. So, the return path can be implemented before the entropy encoding stage. Indeed, the scanning process carried out by the scan unit 360 is also considered lossless, but in the present embodiment the return path 390 is from the output of the quantiser 350 to the input of a complimentary inverse quantiser 420.

In general terms, an entropy decoder 410, the reverse scan unit 400, an inverse quantiser 420 and an inverse transform unit 430 provide the respective inverse functions of the entropy encoder 370, the scan unit 360, the quantiser 350 and the transform unit 340. For now, the discussion will continue through the compression process; the process to decompress an input compressed video signal will be discussed separately below.

In the compression process, the scanned coefficients are passed by the return path 390 from the quantiser 350 to the inverse quantiser 420 which carries out the inverse operation of the scan unit 360. An inverse quantisation and inverse transformation process are carried out by the units 420, 430 to generate a compressed-decompressed residual image signal 440.

The image signal 440 is added, at an adder 450, to the output of the predictor 320 to generate a reconstructed output image 460. This forms one input to the image predictor 320, as will be described below.

Turning now to the process applied to decompress a received compressed video signal 470, the signal is supplied to the entropy decoder 410 and from there to the chain of the reverse scan unit 400, the inverse quantiser 420 and the inverse transform unit 430 before being added to the output of the image predictor 320 by the adder 450. So, at the decoder side, the decoder reconstructs a version of the residual image and then applies this (by the adder 450) to the predicted version of the image (on a block by block basis) so as to decode each block. In

straightforward terms, the output 460 of the adder 450 forms the output decompressed video signal 480. In practice, further filtering may be applied before the signal is output.

So, the apparatus of Figures 5 and 6 can act as a compression apparatus or a decompression apparatus. The functions of the two types of apparatus overlap very heavily.

5 The scan unit 360 and entropy encoder 370 are not used in a decompression mode, and the operation of the predictor 320 (which will be described in detail below) and other units follow mode and parameter information contained in the received compressed bitstream rather than generating such information themselves.

10 Figure 6 schematically illustrates the generation of predicted images, and in particular the operation of the image predictor 320.

There are two basic modes of prediction carried out by the image predictor 320: so-called intra-image prediction and so-called inter-image, or motion-compensated (MC), prediction. At the encoder side, each involves detecting a prediction direction in respect of a current block to be predicted, and generating a predicted block of samples according to other
15 samples (in the same (intra) or another (inter) image). By virtue of the units 310 or 450, the difference between the predicted block and the actual block is encoded or applied so as to encode or decode the block respectively.

Intra-image prediction bases a prediction of the content of a block of the image on data from within the same image. This corresponds to so-called I-frame encoding in other video
20 compression techniques. In contrast to I-frame encoding, where the whole image is intra-encoded, in the present embodiments the choice between intra- and inter- encoding can be made on a block-by-block basis, though in other embodiments the choice is still made on an image-by-image basis.

Motion-compensated prediction is an example of inter-image prediction and makes use
25 of motion information which attempts to define the source, in another adjacent or nearby image, of image detail to be encoded in the current image. Accordingly, in an ideal example, the contents of a block of image data in the predicted image can be encoded very simply as a reference (a motion vector) pointing to a corresponding block at the same or a slightly different position in an adjacent image.

30 Returning to Figure 6, two image prediction arrangements (corresponding to intra- and inter-image prediction) are shown, the results of which are selected by a multiplexer 500 under the control of a mode signal 510 so as to provide blocks of the predicted image for supply to the adders 310 and 450. The choice is made in dependence upon which selection gives the lowest "energy" (which, as discussed above, may be considered as information content requiring
35 encoding), and the choice is signalled to the encoder within the encoded output datastream. Image energy, in this context, can be detected, for example, by carrying out a trial subtraction of an area of the two versions of the predicted image from the input image, squaring each pixel

value of the difference image, summing the squared values, and identifying which of the two versions gives rise to the lower mean squared value of the difference image relating to that image area.

5 The actual prediction, in the intra-encoding system, is made on the basis of image blocks received as part of the signal 460, which is to say, the prediction is based upon encoded-decoded image blocks in order that exactly the same prediction can be made at a decompression apparatus. However, data can be derived from the input video signal 300 by an intra-mode selector 520 to control the operation of the intra-image predictor 530.

10 For inter-image prediction, a motion compensated (MC) predictor 540 uses motion information such as motion vectors derived by a motion estimator 550 from the input video signal 300. Those motion vectors are applied to a processed version of the reconstructed image 460 by the motion compensated predictor 540 to generate blocks of the inter-image prediction.

15 Accordingly, the units 530 and 540 (operating with the estimator 550) each act as detectors to detect a prediction direction in respect of a current block to be predicted, and as a generator to generate a predicted block of samples (forming part of the prediction passed to the units 310 and 450) according to other samples defined by the prediction direction.

20 The processing applied to the signal 460 will now be described. Firstly, the signal is filtered by a filter unit 560, which will be describe in greater detail below. This involves applying a "deblocking" filter to remove or at least tend to reduce the effects of the block-based processing carried out by the transform unit 340 and subsequent operations. A sample adaptive offsetting (SAO) filter (described further below) may also be used. Also, an adaptive loop filter is applied using coefficients derived by processing the reconstructed signal 460 and the input video signal 300. The adaptive loop filter is a type of filter which, using known techniques, applies adaptive filter coefficients to the data to be filtered. That is to say, the filter coefficients can vary in dependence upon various factors. Data defining which filter coefficients to use is included as part of the encoded output datastream.

30 Adaptive filtering represents in-loop filtering for image restoration. An LCU can be filtered by up to 16 filters, with a choice of filter and an ALF on/off status (adaptive loop filter – see below)being derived in respect of each CU within the LCU. Currently the control is at the LCU level, not the CU level.

35 The filtered output from the filter unit 560 in fact forms the output video signal 480 when the apparatus is operating as a compression apparatus. It is also buffered in one or more image or frame stores 570; the storage of successive images is a requirement of motion compensated prediction processing, and in particular the generation of motion vectors. To save on storage requirements, the stored images in the image stores 570 may be held in a compressed form and then decompressed for use in generating motion vectors. For this particular purpose, any known compression / decompression system may be used. The stored images are passed to

an interpolation filter 580 which generates a higher resolution version of the stored images; in this example, intermediate samples (sub-samples) are generated such that the resolution of the interpolated image is output by the interpolation filter 580 is 4 times (in each dimension) that of the images stored in the image stores 570 for the luminance channel of 4:2:0 and 8 times (in each dimension) that of the images stored in the image stores 570 for the chrominance channels of 4:2:0. The interpolated images are passed as an input to the motion estimator 550 and also to the motion compensated predictor 540.

In embodiments, a further optional stage is provided, which is to multiply the data values of the input video signal by a factor of four using a multiplier 600 (effectively just shifting the data values left by two bits), and to apply a corresponding divide operation (shift right by two bits) at the output of the apparatus using a divider or right-shifter 610. So, the shifting left and shifting right changes the data purely for the internal operation of the apparatus. This measure can provide for higher calculation accuracy within the apparatus, as the effect of any data rounding errors is reduced.

The way in which an image is partitioned for compression processing will now be described. At a basic level, an image to be compressed is considered as an array of blocks of samples. For the purposes of the present discussion, the largest such block under consideration is a so-called largest coding unit (LCU) 700 (Figure 7), which represents a square array of typically 64 x 64 samples (the LCU size is configurable by the encoder, up to a maximum size such as defined by the HEVC documents). Here, the discussion relates to luminance samples. Depending on the chrominance mode, such as 4:4:4, 4:2:2, 4:2:0 or 4:4:4:4 (GBR plus key data), there will be differing numbers of corresponding chrominance samples corresponding to the luminance block.

Three basic types of blocks will be described: coding units, prediction units and transform units. In general terms, the recursive subdividing of the LCUs allows an input picture to be partitioned in such a way that both the block sizes and the block coding parameters (such as prediction or residual coding modes) can be set according to the specific characteristics of the image to be encoded.

The LCU may be subdivided into so-called coding units (CU). Coding units are always square and have a size between 8x8 samples and the full size of the LCU 700. The coding units can be arranged as a kind of tree structure, so that a first subdivision may take place as shown in Figure 8, giving coding units 710 of 32x32 samples; subsequent subdivisions may then take place on a selective basis so as to give some coding units 720 of 16x16 samples (Figure 9) and potentially some coding units 730 of 8x8 samples (Figure 10). Overall, this process can provide a content-adapting coding tree structure of CU blocks, each of which may be as large as the LCU or as small as 8x8 samples. Encoding of the output video data takes place on the basis of

the coding unit structure, which is to say that one LCU is encoded, and then the process moves to the next LCU, and so on.

Figure 11 schematically illustrates an array of prediction units (PU). A prediction unit is a basic unit for carrying information relating to the image prediction processes, or in other words the additional data added to the entropy encoded residual image data to form the output video signal from the apparatus of Figure 5. In general, prediction units are not restricted to being square in shape. They can take other shapes, in particular rectangular shapes forming half of one of the square coding units (for example, 8x8 CUs can have 8x4 or 4x8 PUs). Employing PUs which align to image features is not a compulsory part of the HEVC system, but the general aim would be to allow a good encoder to align the boundary of adjacent prediction units to match (as closely as possible) the boundary of real objects in the picture, so that different prediction parameters can be applied to different real objects. Each coding unit may contain one or more prediction units.

Figure 12 schematically illustrates an array of transform units (TU). A transform unit is a basic unit of the transform and quantisation process. Transform units may or may not be square and can take a size from 4x4 up to 32x32 samples. Each coding unit can contain one or more transform units. The acronym SDIP-P in Figure 12 signifies a so-called short distance intra-prediction partition. In this arrangement only one dimensional transforms are used, so a 4xN block is passed through N transforms with input data to the transforms being based upon the previously decoded neighbouring blocks and the previously decoded neighbouring lines within the current SDIP-P. SDIP-P is currently not included in HEVC at the time of filing the present application.

As mentioned above, coding takes place as one LCU, then a next LCU, and so on. Within an LCU, coding is carried out CU by CU. Within a CU, coding is carried out for one TU, then a next TU and so on.

The intra-prediction process will now be discussed. In general terms, intra-prediction involves generating a prediction of a current block (a prediction unit) of samples from previously-encoded and decoded samples in the same image. Figure 13 schematically illustrates a partially encoded image 800. Here, the image is being encoded from top-left to bottom-right on an LCU basis. An example LCU encoded partway through the handling of the whole image is shown as a block 810. A shaded region 820 above and to the left of the block 810 has already been encoded. The intra-image prediction of the contents of the block 810 can make use of any of the shaded area 820 but cannot make use of the unshaded area below that. Note however that for an individual TU within the current LCU, the hierarchical order of encoding (CU by CU then TU by TU) discussed above means that there may be previously encoded samples in the current LCU and available to the coding of that TU which are, for example, above-right or below-left of that TU.

The block 810 represents an LCU; as discussed above, for the purposes of intra-image prediction processing, this may be subdivided into a set of smaller prediction units and transform units. An example of a current TU 830 is shown within the LCU 810.

5 The intra-image prediction takes into account samples coded prior to the current TU being considered, such as those above and/or to the left of the current TU. Source samples, from which the required samples are predicted, may be located at different positions or directions relative to the current TU. To decide which direction is appropriate for a current prediction unit, the mode selector 520 of an example encoder may test all combinations of available TU structures for each candidate direction and select the PU direction and TU
10 structure with the best compression-efficiency.

The picture may also be encoded on a "slice" basis. In one example, a slice is a horizontally adjacent group of LCUs. But in more general terms, the entire residual image could form a slice, or a slice could be a single LCU, or a slice could be a row of LCUs, and so on. Slices can give some resilience to errors as they are encoded as independent units. The
15 encoder and decoder states are completely reset at a slice boundary. For example, intra-prediction is not carried out across slice boundaries; slice boundaries are treated as image boundaries for this purpose.

Figure 14 schematically illustrates a set of possible (candidate) prediction directions. The full set of 34 candidate directions is available to a prediction unit of 8x8, 16x16 or 32x32
20 samples. The special cases of prediction unit sizes of 4x4 and 64x64 samples have a reduced set of candidate directions available to them (17 candidate directions and 5 candidate directions respectively). The directions are determined by horizontal and vertical displacement relative to a current block position, but are encoded as prediction "modes", a set of which is shown in Figure 15. Note that the so-called DC mode represents a simple arithmetic mean of the surrounding
25 upper and left-hand samples.

In general terms, after detecting a prediction direction in respect of each prediction unit, the systems are operable to generate a predicted block of samples according to other samples defined by the prediction direction.

Figure 16 schematically illustrates a so-called up-right diagonal scan, being an example
30 scan pattern which may be applied by the scan unit 360. In Figure 16, the pattern is shown for an example block of 8x8 DCT coefficients, with the DC coefficient being positioned at the top left position 840 of the block, and increasing horizontal and vertical spatial frequencies being represented by coefficients at increasing distances downwards and to the right of the top-left position 840. Other alternative scan orders may be used instead.

35 Variations of the block arrangements and of the CU, PU and TU structures will be discussed below. These will be discussed in the context of the apparatus of Figure 17, which is

similar in many respects to that illustrated in Figures 5 and 6 discussed above. Indeed, many of the same reference numerals have been used, and these parts will not be discussed further.

The main substantive differences with respect to Figures 5 and 6 relate to the filter 560 (Figure 6), which in Figure 17 is shown in more detail as comprising a deblocking filter 1000 and associated encoding decision block 1030, a sample adaptive offsetting (SAO) filter 1010 and associated coefficient generator 1040, and an adaptive loop filter (ALF) 1020 and associated coefficient generator 1050.

The deblocking filter 1000 attempts to reduce distortion and to improve visual quality and prediction performance by smoothing the sharp edges which can form between CU, PU and TU boundaries when block coding techniques are used.

The SAO filter 1010 classifies reconstructed pixels into different categories and then attempts to reduce distortion by simply adding an offset for each category of pixels. The pixel intensity and edge properties are used for pixel classification. To further improve the coding efficiency, a picture can be divided into regions for localization of offset parameters.

The ALF 1020 attempts to restore the compressed picture such that the difference between the reconstructed and source frames is minimized. The coefficients of ALF are calculated and transmitted on a frame basis. The ALF can be applied to the entire frame or to local areas.

As noted above, the proposed HEVC documents use a particular chroma sampling scheme known as the 4:2:0 scheme. The 4:2:0 scheme can be used for domestic/consumer equipment. However, several other schemes are possible.

In particular, a so-called 4:4:4 scheme would be suitable for professional broadcasting, mastering and digital cinema, and in principle would have the highest quality and data rate.

Similarly, a so-called 4:2:2 scheme could be used in professional broadcasting, mastering and digital cinema with some loss of fidelity.

These schemes and their corresponding possible PU and TU block structures are described below.

In addition, other schemes include the 4:0:0 monochrome scheme.

In the 4:4:4 scheme, each of the three Y, Cb and Cr channels have the same sample rate. In principle therefore, in this scheme there would be twice as much chroma data as luma data.

Hence in HEVC, in this scheme each of the three Y, Cb and Cr channels would have corresponding PU and TU blocks that are the same size; for example an 8x8 luma block would have corresponding 8x8 chroma blocks for each of the two chroma channels.

Consequently in this scheme there would generally be a direct 1:1 relationship between block sizes in each channel.

In the 4:2:2 scheme, the two chroma components are sampled at half the sample rate of luma (for example using vertical or horizontal subsampling, but for the purposes of the present description, horizontal subsampling is assumed). In principle therefore, in this scheme there would be as much chroma data as luma data, though the chroma data would be split between the two chroma channels.

Hence in HEVC, in this scheme the Cb and Cr channels would have different size PU and TU blocks to the luma channel; for example an 8x8 luma block could have corresponding 4 wide x 8 high chroma blocks for each chroma channel.

Notably therefore in this scheme the chroma blocks could be non-square, even though they correspond to square luma blocks.

In the currently proposed HEVC 4:2:0 scheme, the two chroma components are sampled at a quarter of the sample rate of luma (for example using vertical and horizontal subsampling). In principle therefore, in this scheme there is half as much chroma data as luma data, the chroma data being split between the two chroma channels.

Hence in HEVC, in this scheme again the Cb and Cr channels have different size PU and TU blocks to the luma channel. For example an 8x8 luma block would have corresponding 4x4 chroma blocks for each chroma channel.

The above schemes are colloquially known in the art as 'channel ratios', as in 'a 4:2:0 channel ratio'; however it will be appreciated from the above description that in fact this does not always mean that the Y, Cb and Cr channels are compressed or otherwise provided in that ratio. Hence whilst referred to as a channel ratio, this should not be assumed to be literal. In fact, the correct ratios for the 4:2:0 scheme are 4:1:1 (the ratios for the 4:2:2 scheme and 4:4:4 scheme are in fact correct).

Before discussing particular arrangements with reference to Figures 18a and 18b, some general terminology will be summarised or revisited.

A Largest Coding Unit (LCU) is a root picture object. Typically, it covers the area equivalent to 64 x 64 luma pixels. It is recursively split to form a tree-hierarchy of Coding Units (CUs). In general terms, the three channels (one luma channel and two chroma channels) have the same CU tree-hierarchy. Having said this, however, depending upon the channel ratio, a particular luma CU may comprise a different number of pixels to the corresponding chroma CUs.

The CUs at the end of the tree-hierarchy, which is to say, the smallest CUs resulting from the recursive splitting process (which may be referred to as leaf CUs) are then split into Prediction Units (PUs). The three channels (luma and two chroma channels) have the same PU structure, except when the corresponding PU for a chroma channel would have too few samples, in which case just one PU for that channel is available. This is configurable, but commonly the minimum dimension of an intra PU is 4 samples; the minimum dimension of an

inter PU is 4 luma samples (or 2 chroma samples for 4:2:0). The restriction on the minimum CU size always is large enough for at least one PU for any channel.

The leaf CUs are also split into Transform Units (TUs). The TUs can - and, when they are too big (for example, over 32x32 samples), must - be split into further TUs. A limit is applied so that TUs can be split down to a maximum tree depth, currently configured as 2 levels. i.e. there can be no more than 16 TUs for each CU. An illustrative smallest allowable TU size is 4x4 samples and the largest allowable TU size is 32x32 samples. Again, the three channels have the same TU structure wherever possible, but if a TU cannot be split to a particular depth for a given channel due to the size restriction, it remains at the larger size. The so-called non-square quad-tree transform arrangement (NSQT) is similar, but the method of splitting into four TUs need not be 2x2, but can be 4x1 or 1x4.

Referring to Figures 18a and 18b, the different block sizes possible are summarised for CU, PU and TU blocks, with 'Y' referring to luma blocks and 'C' referring in a generic sense to a representative one of the chroma blocks, and the numbers referring to pixels. 'Inter' refers to inter-frame prediction PUs (as opposed to intra-frame prediction PUs). In many cases, only the block sizes for the luma blocks are shown. The corresponding sizes of the associated chroma blocks are related to the luma block sizes according to the channel ratios.

So, for 4:4:4, the chroma channels have the same block sizes as the luma blocks shown in Figures 18a and 18b. Accordingly, the chroma PU structure mirrors (or is respectively identical to) the luma PU structure at all split levels (where a "split level" refers to the choice of CU size options and, within a CU size option, the choice of PU size and shape options). While it would be possible to employ the same set of possible block sizes and shapes but to allow different ones (of that set) to be selected in respect of luma and chroma components of a CU, in some embodiments of the disclosure the same PU block size and shape is selected for the chroma and the luma components corresponding to a CU. Therefore, whichever PU size(s) and shape(s) is (are) selected for a luma CU, the identical size and shape PU(s) is (are) selected for the chroma components of that CU. Note that as discussed below, the selection of block sizes and shapes and PUs is an encoder based decision, under the control of the controller 343. In such arrangements, the luminance and chrominance samples are in a 4:4:4 format and are arranged as plural coding units each comprising a block of luminance and chrominance samples; and a prediction direction is detected in respect of a current prediction unit, being a block of luminance or chrominance samples representing at least a subset of a respective coding unit. For a particular coding unit, a size and shape of one or more prediction units each comprising luminance and chrominance samples from that coding unit is made (for example by the controller 343) such that the selection of the prediction unit size and shape are the same for luminance samples and for chrominance samples.

For 4:2:2 and 4:2:0, the chroma blocks will each have fewer pixels than the corresponding luma block, according to the channel ratio.

The arrangements shown in Figures 18a and 18b concern four possible CU sizes: 64x64, 32x32, 16x16 and 8x8 luma pixels respectively. Each of these CUs has a corresponding row of PU options (shown in a column 1140) and TU options (shown in a column 1150). For the possible CU sizes defined above, the rows of options are referenced as 1100, 1110, 1120 and 1130 respectively.

Note that 64x64 is currently a maximum CU size but this restriction could change.

Within each row 1100... 1130, different PU options are shown applicable to that CU size.

The TU options applicable to those PU configurations are shown horizontally aligned with the respective PU option(s).

Note that in several cases, multiple PU options are provided. As discussed above, the aim of the apparatus in selecting a PU configuration is to match (as closely as possible) the boundary of real objects in the picture, so that different prediction parameters can be applied to different real objects.

The block sizes and shapes and PUs are an encoder based decision, under the control of the controller 343. The current method involves conducting trials of many TU tree structures for many directions, getting the best "cost" at each level. Here, the cost may be expressed as a measure of the distortion, or noise, or errors, or bit rate resulting from each block structure. So, the encoder may try two or more (or even all available) permutations of block sizes and shapes within those allowed under the tree structures and hierarchies discussed above, before selecting the one of the trials which gives the lowest bit rate for a certain required quality measure, or the lowest distortion (or errors, or noise, or combinations of these measures) for a required bit rate, or a combination of these measures.

Given the selection of a particular PU configuration, various levels of splitting may be applied to generate the corresponding TUs. Referring to the row 1100, in the case of a 64x64 PU, this block size is too large for use as a TU and so a first level of splitting (from "level 0" (not split) to "level 1") is compulsory, resulting in an array of four 32x32 luma TUs. Each of these may be subjected to further splitting in a tree hierarchy (from "level 1" to "level 2") as required, with the splitting being carried out before transforming or quantising that TU is performed. The maximum number of levels in the TU tree is limited by (for example) the HEVC documents.

Other options are provided for PU sizes and shapes in the case of a 64x64 luma pixel CU. These are restricted to use only with inter-coded pictures and, in some cases, with the so-called AMP option enabled. AMP refers to Asymmetric Motion Partitioning and allows for PUs to be partitioned asymmetrically.

Similarly, in some cases options are provided for TU sizes and shapes. If NQST (non-square quad-tree transform, basically allowing a non-square TU) is enabled, then splitting to

level 1 and/or level 2 can be carried out as shown, whereas if NQST is not enabled, the TU sizes follow the splitting pattern of the respective largest TU for that CU size.

Similar options are provided for other CU sizes.

- 5 In addition to the graphical representation shown in Figures 18a and 18b, the numerical part of the same information is provided in the following table, though the presentation in Figures 18a and 18b is considered definitive. "n/a" indicates a mode which is not allowed. The horizontal pixel size is recited first. If a third figure is given, it relates to the number of instances of that block size, as in (horizontal) x (vertical) x (number of instances) blocks. N is an integer.

CU Size	PU Options	TU Options		
		Level 0	Level 1	Level 2
64x64	64x64	n/a	32x32x4	16x16x4
	64x32x2 (horizontal configuration) 64x16 + 64x48 (2 horizontal configurations)	n/a	32x32x4	32x8x4
	32x64x2 (vertical configuration) 16x64 + 48x64 (2 vertical configurations)	n/a	32x32x4	8x32x4
32x32	32 x 32	32x32	16x16x4	8x8x4
	32x16x2 (horizontal configuration) 32x8 + 32x24 (2 horizontal configurations)	n/a	32x8x4	16x4x4 (luma) + 4x4x4 (chroma, 4:2:0 or 4:2:2) or 8x4x4 (chroma, 4:2:2)
	16x32x2 (vertical configuration) 8x32 + 24x32 (2 vertical configurations)	n/a	8x32x4	4x16x4 (luma) + 4x4x4 (chroma)
16x16	16x16	16x16	8x8x4	4x4x4 (luma) + 4x8x4 (chroma)
	16x8x2 (horizontal configuration) 16x4 + 16x12 (2 horizontal configurations)	n/a	16x4x4 (luma) + 4x8x4 (chroma) (4:2:0 or 4:2:2) 16x4x4 (luma) + 8x4x4 (chroma) (4:2:2)	4x4x4 (luma) + 4x8x1 (chroma) (4:2:0 or 4:2:2) 4x4x4 (luma) + 8x4x1 (chroma) (4:2:2)
	8x16x2 (vertical configuration) 4x16 + 12x16 (2 vertical configurations)	n/a		
8x8	8x8 4x4x4 8x4x2 (horizontal configuration) 4x8x2 (vertical configuration)	8x8	4x4x4 (luma) + 4x8x1 (chroma)	n/a
	4x4x4 (luma) + 4 x N (chroma)	n/a	4x4x4 (luma) + 4x8x1 (chroma)	n/a

10

4:2:0, 4:2:2 and 4:4:4 block structure variants

It has been appreciated that both 4:2:0 and 4:4:4 schemes have square PU blocks for intra-prediction coding. Moreover, currently the 4:2:0 scheme permits 4x4 pixel PU & TU blocks.

In embodiments, it is consequently proposed that for the 4:4:4 scheme the recursion for CU blocks is permitted down to 4x4 pixels rather than 8x8 pixels, since as noted above in the 4:4:4 mode the luma and chroma blocks will be the same size (i.e. the chroma data is not subsampled) and so for a 4x4 CU no PU or TU will need to be less than the already allowed minimum of 4x4 pixels. This is therefore an example of selecting, for a particular coding unit, a size and shape of one or more prediction units each comprising luminance or chrominance samples from at least a subset of that coding unit, the selection of the prediction unit size and shape being the same for luminance samples and for chrominance samples.

Similarly, in the 4:4:4 scheme, in an embodiment each of the Y, Cr, Cb channels, or the Y and the two Cr, Cb channels together, could have respective CU tree-hierarchies. A flag may then be used to signal which hierarchy or arrangement of hierarchies is to be used. This approach could also be used for a 4:4:4 RGB colour space scheme. However, in an alternative, the tree hierarchies for chroma and luma may instead be independent.

In the example of an 8x8 CU in the 4:2:0 scheme, this results in four 4x4 luma PUs and one 4x4 chroma PU. Hence in the 4:2:2 scheme, having twice as much chroma data, one option is in this case is to have two 4x4 chroma PUs, where (for example) the bottom chroma block would correspond in position to the bottom left luma block. However, it is has been appreciated that using one non-square 4x8 chroma PU in this case would be more consistent with arrangements for the 4:2:0 chroma format.

In the 4:2:0 scheme there are in principle some non-square TU blocks permitted for certain classes of inter-prediction coding, but not for intra-prediction coding. However in inter-prediction coding, when non-square quad-tree transforms (NSQT) are disabled (which is the current default for the 4:2:0 scheme), all TUs are square. Hence in effect the 4:2:0 scheme currently enforces square TUs. For example, a 16x16 4:2:0 luma TU would correspond with respective Cb & Cr 8x8 4:2:0 Chroma TUs.

However, as noted previously, the 4:2:2 scheme can have non-square PUs. Consequently in an embodiment it is proposed to allow non-square TUs for the 4:2:2 scheme.

For example, whilst a 16x16 4:2:2 luma TU could correspond with two respective Cb & Cr 8x8 4:2:2 Chroma TUs, in this embodiment it could instead correspond with respective Cb & Cr 8x16 4:2:2 Chroma TUs.

Similarly, four 4x4 4:2:2 luma TUs could correspond with two respective 4x4 Cb+Cr 4:2:2 TUs, or in this embodiment could instead correspond with respective 4x8 Cb & Cr 4:2:2 TUs.

Having non-square chroma TUs, and hence fewer TUs, may be more efficient as they are likely to contain less information. However this may affect the transformation and scanning processes of such TUs, as will be described later.

Finally, for the 4:4:4 scheme it may be preferable to have the TU structure channel-independent, and selectable at the sequence, picture, slice or finer level.

As noted above, NSQT is currently disabled in the 4:2:0 scheme of HEVC. However, if for inter-picture prediction, NSQT is enabled and asymmetric motion partitioning (AMP) is permitted, this allows for PUs to be partitioned asymmetrically; thus for example a 16x16 CU may have a 4x16 PU and a 12x16 PU. In these circumstances, further considerations of block structure are important for each of the 4:2:0 and 4:2:2 schemes.

For the 4:2:0 scheme, in NSQT the minimum width/height of a TU may be restricted to 4 luma/chroma samples:

Hence in a non-limiting example a 16x4/16x12 luma PU structure has four 16x4 luma TUs and four 4x4 chroma TUs, where the luma TUs are in a 1x4 vertical block arrangement and the chroma TUs are in a 2x2 block arrangement.

In a similar arrangement where the partitioning was vertical rather than horizontal, a 4x16/12x16 luma PU structure has four 4x16 luma TUs and four 4x4 chroma TUs, where the luma TUs are in a 4x1 horizontal block arrangement and the chroma TUs are in a 2x2 block arrangement.

For the 4:2:2 scheme, in NSQT as a non-limiting example a 4x16/12x16 luma PU structure has four 4x16 luma TUs and four 4x8 chroma TUs, where the luma TUs are in a 4x1 horizontal block arrangement; the chroma TUs are in a 2x2 block arrangement.

However, it has been appreciated that a different structure can be considered for some cases. Hence in an embodiment, in NSQT as a non-limiting example 16x4/16x12 luma PU structure has four 16x4 luma TUs and four 8x4 chroma TUs, but now the luma and chroma TUs are in a 1x4 vertical block arrangement, aligned with the PU layout (as opposed to the 4:2:0 style arrangement of four 4x8 chroma TUs in a 2x2 block arrangement).

Similarly 32x8 PU can have four 16x4 luma TUs and four 8x4 chroma TUs, but now the luma and chroma TUs are in a 2x2 block arrangement.

Hence more generally, for the 4:2:2 scheme, in NSQT the TU block sizes are selected to align with the asymmetric PU block layout. Consequently the NSQT usefully allows TU boundaries to align with PU boundaries, which reduces high frequency artefacts that may otherwise occur.

In general terms, embodiments of the disclosure can relate to a video coding method, apparatus or program operable in respect of images of a 4:2:2 format video signal. An image to be encoded is divided into coding units, prediction units and transform units for encoding, a coding unit being a square array of luminance samples and the corresponding chrominance samples, there being one or more prediction units in a coding unit, and there being one or more transform units in a coding unit; in which a prediction unit is an elementary unit of prediction so

that all samples within a single prediction unit are predicted using a common prediction technique, and a transform unit is a basic unit of transformation and quantisation.

A Non-square transform mode (such as an NSQT mode) is enabled so as to allow non-square prediction units. Optionally, asymmetric motion partitioning is enabled so as to allow asymmetry between two or more prediction units corresponding to a single coding unit.

The controller 343 controls the selection of transform unit block sizes to align with the prediction unit block layout, for example by detecting image features in the portion of the image corresponding to a PU and selecting TU block sizes in respect of that PU so as to align TU boundaries with edges of image features in the portion of the image.

The rules discussed above dictate which combinations of block sizes are available. The encoder may just try different combinations. As discussed above, a trial may include two or more, through to all available options. The trial encode processes can be carried out according to a cost function metric and a result selected according to an assessment of the cost function.

Given that there are three levels of variation, according to the CU size and shape, the PU size and shape and the TU size and shape, this could lead to a large number of permutations to be trial-encoded. To reduce this variation, the system could trial encode for a CU size by using an arbitrarily selected one of the PU/TU configurations allowable for each CU size; then, having selected a CU size, a PU size and shape could be selected by trial encoding the different PU options each with a single arbitrarily chosen TU configuration. Then, having selected a CU and PU, the system could try all applicable TU configurations to select a final TU configuration.

Another possibility is that some encoders may use a fixed choice of block configuration, or may allow a limited subset of the combinations set out in the discussions above.

Intra-prediction

4:2:0 intra-prediction

Turning now to Figure 22, for intra-prediction, HEVC allows for angular chroma prediction.

By way of introduction, Figure 22 illustrates 35 prediction modes applicable to luma blocks, 33 of which specify directions to reference samples for a current predicted sample position 110. The remaining two modes are mode 0 (planar) and mode 1 (dc).

HEVC allows chroma to have DC, Vertical, Horizontal, Planar, DM_CHROMA and LM_CHROMA modes.

DM_CHROMA indicates that the prediction mode to be used is the same as that of the co-located luma PU (i.e. one of the 35 shown in Figure 22).

LM_CHROMA (linear mode chroma) indicates that co-located luma samples (downsampled as appropriate to the channel ratios) are used to derive the predicted chroma

samples. In this case, if the luma PU from which the DM_CHROMA prediction mode would be taken selected DC, Vertical, Horizontal or Planar, that entry in the chroma prediction list is replaced using mode 34. In the LM_CHROMA mode, the luma pixels from which the chroma pixels are predicted are scaled (and have an offset applied if appropriate) according to a linear relationship between luma and chroma. This linear relationship is derived from surrounding pixels, and the derivation can be carried out on a block by block basis, with the decoder finishing decoding one block before moving on to the next.

It is notable that the prediction modes 2-34 sample an angular range from 45 degrees to 225 degrees; that is to say, one diagonal half of a square. This is useful in the case of the 4:2:0 scheme, which as noted above only uses square chroma PUs for intra-picture prediction.

4:2:2 intra-prediction variants

However, also as noted above the 4:2:2 scheme could have rectangular (non-square) chroma PUs even when the luma PUs are square. Or indeed, the opposite could be true: a rectangular luma PU could correspond to a square chroma PU. The reason for the discrepancy is that in 4:2:2, the chroma is subsampled horizontally (relative to the luma) but not vertically. So the aspect ratio of a luma block and a corresponding chroma block would be expected to be different.

Consequently, in an embodiment, for chroma PUs having a different aspect ratio to the corresponding luma block, a mapping table may be required for the direction. Assuming (for example) a 1-to-2 aspect ratio for rectangular chroma PUs, then for example mode 18 (currently at an angle of 135 degrees) may be re-mapped to 123 degrees. Alternatively selection of current mode 18 may be remapped to a selection of current mode 22, to much the same effect.

Hence more generally, for non-square PUs, a different mapping between the direction of the reference sample and the selected intra prediction mode may be provided compared with that for square PUs.

More generally still, any of the modes, including the non-directional modes, may also be re-mapped based upon empirical evidence.

It is possible that such mapping will result in a many-to-one relationship, making the specification of the full set of modes redundant for 4:2:2 chroma PUs. In this case, for example it may be that only 17 modes (corresponding to half the angular resolution) are necessary. Alternatively or in addition, these modes may be angularly distributed in a non-uniform manner.

Similarly, the smoothing filter used on the reference sample when predicting the pixel at the sample position may be used differently; in the 4:2:0 scheme it is only used to smooth luma pixels, but not chroma ones. However, in the 4:2:2 and 4:4:4 schemes this filter may also be used for the chroma PUs. In the 4:2:2 scheme, again the filter may be modified in response to the different aspect ratio of the PU, for example only being used for a subset of near horizontal modes. An example subset of modes is preferably 2-18 and 34, or more preferably 7-14. In

4:2:2, smoothing of only the left column of reference samples may be carried out in embodiments.

These arrangements are discussed later in more detail.

4:4:4 intra-prediction variants

5 In the 4:4:4 scheme, the chroma and luma PUs are the same size, and so the intra-prediction mode for a chroma PU can be either the same as the co-located luma PU (so saving some overhead in the bit stream by not having to encode a separate mode), or alternatively, it can be independently selected.

10 In this latter case therefore, in an embodiment the system may have 1, 2 or 3 different prediction modes for each of the PUs in a CU;

In a first example, the Y, Cb and Cr PUs may all use the same intra-prediction mode.

In a second example, the Y PU may use one intra-prediction mode, and the Cb and Cr PUs both use another independently selected intra-prediction mode.

15 In a third example, the Y, Cb and Cr PUs each use a respective independently selected intra-prediction mode.

It will be appreciated that having independent prediction modes for the chroma channels (or each chroma channel) will improve the colour prediction accuracy. But this is at the expense of an additional data overhead to communicate the independent prediction modes as part of the encoded data.

20 To alleviate this, the selection of the number of modes could be indicated in the high-level syntax (e.g. at sequence, picture, or slice level). Alternatively, the number of independent modes could be derived from the video format; for example, GBR could have up to 3, whilst YCbCr could be restricted to up to 2.

25 In addition to independently selecting the modes, the available modes may be allowed to differ from the 4:2:0 scheme in the 4:4:4 scheme.

For example as the luma and chroma PUs are the same size in 4:4:4, the chroma PU may benefit from access to all of the 35 + LM_CHROMA + DM_CHROMA directions available. Hence for the case of Y, Cb and Cr each having independent prediction modes, then the Cb channel could have access to DM_CHROMA & LM_CHROMA, whilst the Cr channel could have access to DM_CHROMA_Y, DM_CHROMA_Cb, LM_CHROMA_Y and LM_CHROMA_Cb, where these replace references to the Luma channel with references to the Y or Cb chroma channels.

35 Where the luma prediction modes are signalled by deriving a list of most probable modes and sending an index for that list, then if the chroma prediction mode(s) are independent, it may be necessary to derive independent lists of most probable modes for each channel.

Finally, in a similar manner to that noted for the 4:2:2 case above, in the 4:4:4 scheme the smoothing filter used on the reference sample when predicting the pixel at the sample position may be used for chroma PUs in a similar manner to luma PUs. Currently, a [1,2,1] low-pass filter can be applied to the reference samples prior to intra-prediction. This is only used for luma TUs when using certain prediction modes.

One of the intra-prediction modes available to chroma TUs is to base the predicted samples on co-located luma samples. Such an arrangement is illustrated schematically in Figure 19, which shows an array of TUs 1200 (from a region of a source image) represented by small squares in the Cb, Cr and Y channels, showing the special alignment between image features (schematically indicated by dark and light shaded boxes 1200) in the Cb and Y channels and in the Cr and Y channels. In this example, it is of benefit to force the chroma TUs to base their predicted samples on co-located luma samples. However, it is not always the case that image features correspond between the three channels. In fact, certain features may appear only in one or two of the channels, and in general the image content of the three channels may differ.

In embodiments, for Cr TUs, LM_Chroma could optionally be based on co-located samples from the Cb channel (or, in other embodiments, the dependence could be the other way around). Such an arrangement is shown in schematic form in Figure 20. Here, spatially aligned TUs are illustrated between the Cr, Cb and Y channels. A further set of TUs labelled "source" is a schematic representation of the colour picture as seen as a whole. The image features (a top left triangle and a lower right triangle) seen in the source image do not in fact represent changes in the luminance, but only changes in chrominance between the two triangular regions. In this case, basing LM_Chroma for Cr on the luminance samples would produce a poor prediction, but basing it on the Cb samples could give a better prediction.

The decision as to which LM_Chroma mode to be used can be made by the controller 343 and/or the mode controller 520, based on trial encoding of different options (including the option of basing LM_Chroma on the co-located luma or co-located chroma samples), with the decision as to which mode to select being made by assessing a cost function, similar to that described above, with respect to the different trial encodings. Examples of the cost function are noise, distortion, error rate or bit rate. A mode from amongst those subjected to trial encoding which gives the lowest of any one or more of these cost functions is selected.

Figure 21 schematically illustrates a method used to obtain reference samples for intra-prediction in embodiments. In viewing Figure 21, it should be borne in mind that encoding is carried out according to a scanning pattern, so that in general terms encoded versions of the blocks above and to the left of a current block to be encoded are available to the encoding process. Sometimes samples below-left or to the above-right are used, if they have been

previously coded as part of other already-encoded TUs within the current LCU. Reference is made to Figure 13 as described above, for example.

A shaded area 1210 represents a current TU, that is to say, a TU which is currently being encoded.

5 In 4:2:0 and 4:2:2, the column of pixels immediately to the left of the current TU does not contain co-located luminance and chrominance samples because of horizontal subsampling. In other words, this is because the 4:2:0 and 4:2:2 formats have half as many chrominance pixels as luminance pixels (in a horizontal direction), so not every luminance sample position has a co-located chrominance sample. Therefore, although luminance samples may be present in the
10 column of pixels immediately to the left of the TU, chrominance samples are not present. Therefore, in some embodiments, the column located two samples to the left of the current TU is used to provide reference samples for LM_Chroma. Note that the situation is different in 4:4:4, in that the column immediately to the left of the current TU does indeed contain co-located luma and chroma samples. This column could therefore be used to provide reference
15 samples.

The reference samples are used as follows.

In the LM_Chroma mode, predicted chroma samples are derived from reconstructed luma samples according to a linear relationship. So, in general terms, it can be said that the predicted chrominance values within the TU are given by a relationship:

$$20 \quad P_C = a + bP_L$$

where P_C is a chrominance sample value, P_L is a reconstructed luminance sample value at that sample position, and a and b are constants. The constants are derived (for example, by the controller 343) and encoded (again, by the controller 343 for example, acting as an encoder) for a particular block by comparing, and detecting the relationship between, reconstructed luma
25 samples and chroma samples in the row just above that block and in the column just to the left of that block, these being sample positions which have already been encoded (see above).

In some embodiments, the constants a and b are derived as follows:

$$a = R(P_L', P_C') / R(P_L', P_L')$$

where R represents a linear (least squares) regression function, and P_L' and P_C' are luminance and chrominance samples respectively from the adjacent row and column as discussed above,
30 and:

$$b = \text{mean}(P_C') - a \cdot \text{mean}(P_L')$$

For 4:4:4, the P_L' and P_C' values are taken from the column immediately to the left of the current TU, and the row immediately above the current TU. For 4:2:2, the P_L' and P_C' values
35 are taken from the row immediately above the current TU and the column in the adjacent block which is two sample positions away from the left edge of the current TU. For 4:2:0 (which is subsampled vertically and horizontally) the P_L' and P_C' values would ideally be taken from a row

which is two rows above the current TU, but in fact are taken from a row in the adjacent block which is one sample positions above the current TU, and the column in the adjacent block which is two sample positions away from the left edge of the current TU. The reason is to avoid having to maintain an additional whole row of data in memory. So in this regard, 4:2:2 and 4:2:0 are treated in a similar way.

Accordingly, these techniques apply to video coding methods having a chrominance prediction mode in which a current block of chrominance samples representing a region of the image is encoded by deriving and encoding a relationship of the chrominance samples with respect to a co-sited block of luminance samples (such as reconstructed luminance samples) representing the same region of the image. The relationship (such as the linear relationship) is derived by comparing co-sited (otherwise expressed as correspondingly-sited) luminance and chrominance samples from adjacent already-encoded blocks. The chrominance samples are derived from luminance samples according to the relationship; and the difference between the predicted chrominance samples and the actual chrominance samples is encoded as residual data.

In respect of a first sampling resolution (such as 4:4:4) where the chrominance samples have the same sampling rate as the luminance samples, the co-sited samples are samples in sample positions adjacent to the current block.

In respect of a second sampling resolution (such as 4:2:2 or 4:2:0) where the chrominance samples have a lower sampling rate than that of the luminance samples, a nearest column or row of co-sited luminance and chrominance samples from the adjacent already-encoded block is used to provide the co-sited samples. Or, in the case of the second sampling resolution being a 4:2:0 sampling resolution, the correspondingly-sited samples are a row of samples adjacent to the current block and a nearest column or row of correspondingly-sited luminance and chrominance samples, from the adjacent already-encoded blocks.

The appropriate predictor as described above uses the derived relationship to predict the samples according to the LM_Chroma mode.

Figure 22 schematically illustrates the available prediction angles for luma samples. The current pixel being predicted as shown at the centre of the diagram as a pixel 1220. The smaller dots 1230 represent adjacent pixels. Those located on the top or left sides of the current pixel are available as reference samples to generate a prediction, because they have been previously encoded. Other pixels are currently unknown (at the time of predicting the pixel 1220) and will in due course be predicted themselves.

Each numbered prediction direction points to the reference samples 1230 on the top or left edges of the current block that are used to generate the current predicted pixel. In the case

of smaller blocks, where the prediction directions point to locations between reference samples, a linear interpolation between adjacent reference samples is used.

Turning now to intra-angular prediction for chroma samples, for 4:2:0, fewer prediction directions are available because of the relative scarcity of the chroma samples. However, if the DM_CHROMA mode is selected then the current chroma block will use the same prediction direction as the co-located luma block. In turn, this means that the luma directions for intra-prediction are also available to chroma.

However, for chroma samples in 4:2:2, it can be considered counter-intuitive to use the same prediction algorithm and direction as luma when DM_CHROMA is selected, given that chroma blocks now have a different aspect ratio to that of the luma blocks. For example, a 45° line for a square luma array of samples should still map to a 45° line for chroma samples, albeit with an array of rectangular sized samples. Overlaying the rectangular grid onto to a square grid indicates that the 45° line would then in fact map to a 26.6° line.

Figure 23 schematically illustrates luma intra-prediction directions as applied to chroma pixels in 4:2:2, in respect of a current pixel to be predicted 1220. Note that there are half as many pixels horizontally as there are vertically, because 4:2:2 has half the horizontal sample rate in the chroma channel as compared to the luma channel.

Figure 24 schematically illustrates the transformation or mapping of the 4:2:2 chroma pixels to a square grid, and subsequently how this transformation changes the prediction directions.

The luma prediction directions are shown as broken lines 1240. The chroma pixels are remapped to a square grid giving a rectangular array half the width of the corresponding luma array (such as that shown in Figure 22). The prediction directions shown in Figure 23 have been remapped to the rectangular array. It can be seen that for some pairs of directions (a pair being a luma direction and a chroma direction) there is either an overlap or a close relationship. For example, direction 2 in the luma array substantially overlies the direction 6 in the chroma array. However, it will also be noted that some luma directions, approximately half of them, have no corresponding chroma direction. An example is the luma direction numbered 3. Also, some chroma directions (2-5) have no equivalent in the luma array, and some luma directions (31-34) have no equivalent in the chroma array. But in general, the superposition as shown in Figure 24 demonstrates that it would be inappropriate to use the same angle for both the luma and chroma channels.

Accordingly, in order to derive the appropriate prediction angle for chroma when (a) DM_CHROMA is selected and (b) the DM_CHROMA mode currently in use indicates that the chroma prediction direction should be that of the co-located luma block, the following procedure is applied:

(i) derive the intra-prediction angle step and its inverse according to the luma direction according to usual HEVC rules

(ii) if the luma direction is predominantly vertical (that is, for example, a mode numbered from 18 to 34 inclusive) then the intra-prediction angle step is halved and its inverse is doubled.

5 (iii) otherwise, if the luma direction is predominantly horizontal (that is, for example, a mode numbered from 2 to 17 inclusive) then the intra-prediction angle step is doubled and its inverse halved.

Accordingly these embodiments relate to video coding or decoding methods, apparatus or programs in which luminance and chrominance samples are predicted from other respective reference samples according to a prediction direction associated with a sample to be predicted. In modes such as 4:2:2 the chrominance samples have a lower horizontal and/or vertical sampling rate than the luminance samples so that the ratio of luminance horizontal resolution to chrominance horizontal resolution is different to the ratio of luminance vertical resolution to chrominance vertical resolution. In short, this means that a block of luminance samples has a different aspect ratio to a corresponding block of chrominance samples.

10 The intra frame predictor 530, for example, is operable to detect a first prediction direction defined in relation to a grid of a first aspect ratio in respect of a set of current samples to be predicted; and to apply a direction mapping to the prediction direction so as to generate a second prediction direction defined in relation to a grid of samples of a different aspect ratio of the same set of current samples to be predicted.

In embodiments, the first prediction direction is defined in respect of one of luminance or chrominance samples, and the second prediction direction is defined in respect of the other of luminance or chrominance samples. In the particular examples discussed in the present description, the luminance prediction direction may be modified to provide the chrominance prediction direction. But the other way round could be used.

25 The technique is particularly applicable to intra-prediction, so that the reference samples are samples of the same respective image as the samples to be predicted.

In at least some arrangements the first prediction direction is defined with respect to a square block of luminance samples including the current luminance sample; and the second prediction direction is defined with respect to a rectangular block of chrominance samples including the current chrominance sample.

30 It is possible to provide independent prediction modes for the two chrominance components. In such an arrangement the chrominance samples comprise samples of first and second chrominance components, and the technique comprises applying the direction mapping discussed above step in respect of the first chrominance component (such as Cb); and providing a different prediction mode in respect of the second chrominance component (such as Cr).

The video data can be in a 4:2:2 format or a 4:4:4 format, for example.

In general terms, embodiments of the disclosure can provide for independent prediction modes for the chrominance components (for example, for each of the luminance and chrominance components separately). These embodiments relate to video coding methods in which luminance and chrominance samples of an image are predicted from other respective reference samples of the same image according to a prediction direction associated with a sample to be predicted, the chrominance samples having a lower horizontal and/or vertical sampling rate than the luminance samples so that the ratio of luminance horizontal resolution to chrominance horizontal resolution is different to the ratio of luminance vertical resolution to chrominance vertical resolution so that a block of luminance samples has a different aspect ratio to a corresponding block of chrominance samples, and the chrominance samples representing first and second chrominance components.

The intra frame mode selector 520 selects a prediction mode defining a selection of one or more reference samples for predicting a current chrominance sample of the first chrominance component (such as Cb). It also selects a different prediction mode defining a different selection of one or more reference samples for predicting a current chrominance sample of the second chrominance component (such as Cr), co-sited with the current chrominance sample of the first chrominance component.

A reference sample filter can optionally be applied, for example, as part of the operation of the frame stores 570 and/or the predictors 530, 540 acting as a filter arrangement, to horizontal samples or vertical samples (or both). The filter can be a 3-tap "1 2 1" filter, currently applied to all luma reference samples except the bottom left and top right (the samples of a NxN block are gathered together to form a single 1D array of size 2N+1, and then optionally filtered). In embodiments of the technology it is applied only the first (left hand edge) or last (top edge) N+1 chroma samples for 4:2:2, but noting that the bottom left, top right and top left would then not be adjusted; or all chroma samples (as for luma), for 4:2:2 and 4:4:4.

Embodiments can also provide video coding or decoding methods, apparatus or programs in which luminance and first and second chrominance component samples are predicted from other respective reference samples according to a prediction direction associated with a sample to be predicted, involving predicting samples of the second chrominance component from samples of the first chrominance component.

Embodiments can also provide video coding or decoding methods, apparatus or programs in which luminance and first and second chrominance component samples are predicted from other respective reference samples according to a prediction direction associated with a sample to be predicted, involving filtering the reference samples.

As discussed with reference to Figures 19 and 20, it is possible that the different prediction mode comprises a mode by which samples of the second chrominance component are predicted from samples of the first chrominance component.

Note that modes 0 and 1 are not angular prediction modes and so are not included in this procedure. The effect of the procedure shown above is to map the chroma prediction directions onto the luma prediction directions in Figure 24.

For 4:2:0, when either a purely horizontal prediction mode (luma mode 10) or a purely vertical prediction mode (luma mode 26) is selected, the top or left edges of the predicted TU are subject to filtering for the luma channel only. For the horizontal prediction mode, the top row is filtered in the vertical direction. For the vertical prediction mode, the left column is filtered in the horizontal direction.

Filtering a column of samples in the horizontal direction can be understood as applying a horizontally oriented filter to each sample in turn of the column of samples. So, for an individual sample, its value will be modified by the action of the filter, based on a filtered value generated from the current value of that sample and of one or more other samples at sample positions displaced from that sample in a horizontal direction (that is, one or more other samples to the left and/or right of the sample in question).

Filtering a row of samples in the vertical direction can be understood as applying a vertically oriented filter to each sample in turn of the row of samples. So, for an individual sample, its value will be modified by the action of the filter, based on a filtered value generated from the current value of that sample and of one or more other samples at sample positions displaced from that sample in a vertical direction (that is, one or more other samples above and/or below the sample in question).

One purpose of the edge pixel filtering process described above is to aim to reduce block based edge effects in the prediction thereby aiming to reduce energy in the residual image data.

In embodiments, a corresponding filtering process is also provided for chroma TUs in 4:4:4 and 4:2:2. Taking into account the horizontal subsampling, one proposal is only to filter the top row of the chroma TU in 4:2:2, but to filter both the top row and left column (as appropriate, according to the selected mode) in 4:4:4. It is considered appropriate to filter only in these regions so as to avoid filtering out too much useful detail, which (if filtered out) would lead to an increased energy of the residual data.

For 4:2:0, when DC mode is selected, one or both of the top and/or left edges of the predicted TU are subject to filtering for the luma channel only. Here, this is an example of a case where the luminance samples represent a luminance component and respective chrominance samples represent two chrominance components, the filtering step is applied to a subset of the three components, the subset being either one or two of the three components.

The subset may consist of the luminance component. The filtering may involve filtering one or both of the left column of samples in the predicted block of samples and the top row of samples in the predicted block of samples.

The filtering may be such that in DC Mode, the filter does a $(1 \times \text{neighbouring outside sample} + 3 \times \text{edge sample})/4$ averaging operation for all samples on both edges. However, for the top left the filter function is $(2 \times \text{current sample} + 1 \times \text{above sample} + 1 \times \text{left sample})/4$. This is an example of an operation in which in a DC mode in which a predicted sample is generated as a simple arithmetic mean of surrounding samples, the filtering step comprising filtering the left column of samples in the predicted block of samples and filtering the top row of samples in the predicted block of samples.

The H/V filter is an average between neighbouring outside sample and edge sample.

In some embodiments, this filtering process is also provided for chroma TUs in 4:4:4 and 4:2:2. Again, taking into account the horizontal subsampling, in some embodiments, only the top row of the chroma samples is filtered for 4:2:2, but the top row and left column of the chroma TU are filtered for 4:4:4.

Accordingly, this technique can apply in respect of a video coding or decoding method, apparatus or program in which luminance and chrominance samples in (for example) a 4:4:4 format or a 4:2:2 format are predicted from other respective samples according to a prediction direction associated with blocks of samples to be predicted.

In embodiments of the technique, a prediction direction is detected in respect of a current block to be predicted. A predicted block of chrominance samples is generated according to other chrominance samples defined by the prediction direction. If the detected prediction direction is substantially vertical (for example, being within $\pm n$ angle modes of the exactly vertical mode where n is (for example) 2), the left column of samples is filtered (for example, in a horizontal direction using a horizontally oriented filter) in the predicted block of chrominance samples. Or, if the detected prediction direction is substantially horizontal (for example, being within $\pm n$ angle modes of the exactly horizontal mode, where n is (for example) 2), the top row of samples is filtered (for example, in a vertical direction using a vertically oriented filter) in the predicted block of chrominance samples. In each case, the operation may apply only to the left column or the top row, respectively. Then the difference between the filtered predicted chrominance block and the actual chrominance block is encoded, for example as residual data. Alternatively, the test could be for an exactly vertical or horizontal mode rather than a substantially vertical or horizontal mode. The tolerance of $\pm n$ could be applied to one of the tests (vertical or horizontal) but not the other. In embodiments of the disclosure, only the left column or the top row of the predicted block may be filtered, and the filtering may be carried out by a horizontally oriented filter or a vertically oriented filter respectively.

The filtering may be carried out by the respective predictor 520, 530, acting as a filter in this regard.

After the filtering process, embodiments of the technique either encode a difference between the filtered predicted chrominance block and the actual chrominance block (at an encoder) or apply a decoded difference to the filtered predicted chrominance block so as to encode the block (at a decoder).

Inter-prediction

it is noted that inter prediction in HEVC already allows rectangular PUs, so 4:2:2 and 4:4:4 modes are already compatible with PU inter-prediction processing.

Each frame of a video image is a discrete sampling of a real scene, and as a result each pixel is a step-wise approximation of a real-world gradient in colour and brightness.

In recognition of this, when predicting the Y, Cb or Cr value of a pixel in a new video frame from a value in a previous video frame, the pixels in that previous video frame are interpolated to create a better estimate of the original real-world gradients, to allow a more accurate selection of brightness or colour for the new pixel. Consequently the motion vectors used to point between video frames are not limited to an integer pixel resolution. Rather, they can point to a sub-pixel position within the interpolated image.

4:2:0 inter-prediction

Referring now to Figures 25 and 26, in the 4:2:0 scheme as noted above typically an 8x8 luma PU 1300 will be associated with Cb and Cr 4x4 chroma PUs 1310. Consequently to interpolate the luma and chroma pixel data up to the same effective resolution, different interpolation filters are used.

For example for the 8x8 4:2:0 luma PU, interpolation is $\frac{1}{4}$ pixel, and so an 8-tap x4 filter is applied horizontally first, and then the same 8-tap x4 filter is applied vertically, so that the luma PU is effectively stretched 4 times in each direction, to form an interpolated array 1320 as shown in Figure 25. Meanwhile the corresponding 4x4 4:2:0 chroma PU is $\frac{1}{8}$ pixel interpolated to generate the same eventual resolution, and so a 4-tap x8 filter is applied horizontally first, then the same 4-tap x8 filter is applied vertically, so that the 4:2:0 chroma PUs are effectively stretched 8 times in each direction to form an array 1330, as shown in Figure 26.

4:2:2 inter-prediction

A similar arrangement for 4:2:2 will now be described with reference to Figures 27 and 28, which illustrate a luma PU 1350 and a pair of corresponding chroma PUs 1360.

Referring to Figure 28, as noted previously, in the 4:2:2 scheme the chroma PU 1360 can be non-square, and for the case of an 8x8 4:2:2 luma PU, will typically be a 4 wide x 8 high 4:2:2 Chroma PU for each of the Cb and Cr channels. Note that the chroma PU is drawn, for

the purposes of Figure 28, as a square shaped array of non-square pixels, but in general terms it is noted that the PUs 1360 are 4 (horizontal) x 8 (vertical) pixel arrays.

Whilst it may be possible therefore to use the existing 8-tap x4 luma filter vertically on the chroma PU, in an embodiment of the present disclosure it has been appreciated that the existing 4-tap x8 chroma filter would suffice for vertical interpolation as in practice one is only interested in the even fractional locations of the interpolated chroma PU.

Hence Figure 27 shows the 8x8 4:2:2 luma PU 1350 interpolated as before with an 8-tap x4 filter, and the 4x8 4:2:2 chroma PUs 1360 interpolated with the existing 4-tap x8 chroma filter in the horizontal and vertical direction, but only with the even fractional results used for forming the interpolated image in the vertical direction.

These techniques are applicable to video coding or decoding methods, apparatus or programs using inter-image prediction to encode input video data in which each chrominance component has $1/M$ th of the horizontal resolution of the luminance component and $1/N$ th of the vertical resolution of the luminance component, where M and N are integers equal to 1 or more, For example, For 4:2:2, M=2, N=1. For 4:2:0, M=2, N=2.

The frame store 570 is operable to store one or more images preceding a current image.

The interpolation filter 580 is operable to interpolate a higher resolution version of prediction units of the stored images so that the luminance component of an interpolated prediction unit has a horizontal resolution P times that of the corresponding portion of the stored image and a vertical resolution Q times that of the corresponding portion of the stored image, where P and Q are integers greater than 1. In the current examples, $P = Q = 4$ so that the interpolation filter 580 is operable to generate an interpolated image at $1/4$ sample resolution.

The motion estimator 550 is operable to detect inter-image motion between a current image and the one or more interpolated stored images so as to generate motion vectors between a prediction unit of the current image and areas of the one or more preceding images.

The motion compensated predictor 540 is operable to generate a motion compensated prediction of the prediction unit of the current image with respect to an area of an interpolated stored image pointed to by a respective motion vector.

Returning to a discussion of the operation of the interpolation filter 580, embodiments of this filter are operable to apply a xR horizontal and xS vertical interpolation filter to the chrominance components of a stored image to generate an interpolated chrominance prediction unit, where R is equal to $(U \times M \times P)$ and S is equal to $(V \times N \times Q)$, U and V being integers equal to 1 or more; and to subsample the interpolated chrominance prediction unit, such that its horizontal resolution is divided by a factor of U and its vertical resolution is divided by a factor of V, thereby resulting in a block of MP x NQ samples.

So, in the case of 4:2:2, the interpolation filter 580 applies a x8 interpolation in the horizontal and vertical directions, but then vertically subsamples by a factor of 2, for example by

using every 2nd sample in the interpolated output. One way to achieve this is by doubling an index value into the array of samples. So, consider an array direction (such as the vertical direction in this example) in the array of samples which is 8 samples across, indexed as 0..7. A required sample in the subsampled range is indexed in the range 0..3. So, doubling this index will give values of 0, 2, 4 and 6 which can then be used to access samples in the original array so that every alternate sample is used. This is an example of selecting a subset of samples of the interpolated chrominance prediction unit.

This technique therefore allows the same (for example, x8) filter to be used in respect of 4:2:0 and 4:2:2, but with a further step of subsampling where needed with 4:2:2.

In embodiments, as discussed, the interpolated chrominance prediction unit has a height in samples twice that of a 4:2:0 format prediction unit interpolated using the same xR and xS interpolation filters.

The need to provide different filters can be avoided or alleviated using these techniques, and in particular by using the same xR horizontal and xS vertical interpolation filters, in respect of 4:2:0 input video data and 4:2:2 input video data.

As discussed, the step of subsampling the interpolated chrominance prediction unit comprises using every Vth sample of the interpolated chrominance prediction unit in the vertical direction, and/or using every Uth sample of the interpolated chrominance prediction unit in the vertical direction. More generally, the subsampling can comprise selecting a subset of samples of the interpolated chrominance prediction unit.

Embodiments can involve deriving a luminance motion vector for a prediction unit; and independently deriving one or more chrominance motion vectors for that prediction unit.

In some embodiments, at least one of R and S is equal to 2 or more, and in some embodiments the xR horizontal and xS vertical interpolation filters are also applied to the luminance components of the stored image.

4:4:4 inter-prediction variants

By extension, the same principle of only using the even fractional results for the existing 4-tap x8 chroma filter can be applied both vertically and horizontally for the 8x8 4:4:4 chroma PUs.

Further to these examples, the x8 chroma filter may be used for all interpolation, including luma.

Further inter-prediction variants

In one implementation of motion vector (MV) derivation, one vector is produced for a PU in a P-slice (and two vectors for a PU in a B-slice (where a P-slice takes predictions from a preceding frame, and a B-slice takes predictions from a preceding and following frame, in a similar manner to MPEG P and B frames). Notably, in this implementation in the 4:2:0 scheme the vectors are common to all channels, and moreover, the chroma data need not be used to

calculate the motion vectors. In other words, all the channels use a motion vector based on the luma data.

In an embodiment, in the 4:2:2 scheme the chroma vector could be derived so as to be independent from luma (i.e. a single vector for the Cb and Cr channels could be derived separately), and in the 4:4:4 scheme chroma vectors could further be independent for each of the Cb and Cr channels.

Transforms

In HEVC, most images are encoded as motion vectors from previously encoded/decoded frames, with the motion vectors telling the decoder where, in these other decoded frames, to copy good approximations of the current image from. The result is an approximate version of the current image. HEVC then encodes the so-called residual, which is the error between that approximate version and the correct image. This residual requires much less information than specifying the actual image directly. However, it is still generally preferable to compress this residual information to reduce the overall bitrate further.

In many encoding methods including HEVC, such data is transformed into the spatial frequency domain using an integer cosine transform (ICT), and typically some compression is then achieved by retaining low spatial frequency data and discarding higher spatial frequency data according to the level of compression desired.

4:2:0 transforms

The spatial frequency transforms used in HEVC are conventionally ones that generate coefficients in powers of 4 (for example 64 frequency coefficients) as this is particularly amenable to common quantisation/compression methods. The square TUs in the 4:2:0 scheme are all powers of 4 and hence this is straightforward to achieve.

If the NSQT options are enabled, some non-square transforms are available for non-square TUs, such as 4x16, but again notably these result in 64 coefficients, i.e. again a power of 4.

4:2:2 and 4:4:4 transform variants

The 4:2:2 scheme can result in non-square TUs that are not powers of 4; for example a 4x8 TU has 32 pixels, and 32 is not a power of 4.

In an embodiment therefore, a non-square transform for a non-power of 4 number of coefficients may be used, acknowledging that modifications may be required to the subsequent quantisation process.

Alternatively, in an embodiment non-square TUs are split into square blocks having a power of 4 area for transformation, and then the resulting coefficients can be interleaved.

For example, for 4x8 blocks odd/even vertical samples can be split into two square blocks. Alternatively, for 4x8 blocks the top 4x4 pixels and the bottom 4x4 pixels could form two

square blocks. Alternatively again, for 4x8 blocks a Haar wavelet decomposition can be used to form a lower and an upper frequency 4x4 block.

Any of these options may be made available, and the selection of a particular alternative may be signalled to or derived by the decoder.

5 Other transform modes

In the 4:2:0 scheme there is a proposed flag (the so-called 'qprime_y_zero_transquant_bypass_flag') allowing the residual data to be included in the bit stream losslessly (i.e. without being transformed, quantised or further filtered). In the 4:2:0 scheme the flag applies to all channels.

10 Accordingly, such embodiments represent a video coding or decoding method, apparatus or program in which luminance and chrominance samples are predicted and the difference between the samples and the respective predicted samples is encoded, making use of an indicator configured to indicate whether luminance difference data is to be included in an output bitstream losslessly; and to independently indicate whether chrominance difference data
15 is to be included in the bitstream losslessly. Such a flag or flags (or indicator or indicators respectively) may be inserted by the controller 343, for example.

In an embodiment, it is proposed that the flag for the luma channel is separate to the chroma channels. Hence for the 4:2:2 scheme, such flags should be provided separately for the luma channel and for the chroma channels, and for the 4:4:4 scheme, such flags should be
20 provided either separately for the luma and chroma channels, or one flag is provided for each of the three channels. This recognises the increased chroma data rates associated with the 4:2:2 and 4:4:4 schemes, and enables, for example, lossless luma data together with compressed chroma data.

For intra-prediction coding, mode-dependent directional transform (MDDT) allows the
25 horizontal or vertical ICT (or both ICTs) for a TU to be replaced with an Integer Sine Transform depending upon the intra-prediction direction. In the 4:2:0 scheme this is not applied to chroma TUs. However in an embodiment it is proposed to apply it to 4:2:2 and 4:4:4 chroma TUs, noting that the IST is only currently defined for a 4 sample transform dimensions (either horizontally or vertically), and therefore cannot currently be applied vertically to a 4x8 chroma TU.

30 In methods of video coding, the various embodiments can be arranged so as to indicate whether luminance difference data is to be included in an output bitstream losslessly; and independently to indicate whether chrominance difference data is to be included in the bitstream losslessly, and to encode or include the relevant data in the form defined by such indications.

Quantisation

35 In the 4:2:0 scheme, the quantisation calculation is the same for chrominance as for luminance. Only the quantisation parameters (QPs) differ.

QPs for chrominance are calculated from the luminance QPs as follows:

$$Qp_{Cb} = scalingTable[Qp_{luminance} + chroma_qp_index_offset]$$

$$Qp_{Cr} = scalingTable[Qp_{luminance} + second_chroma_qp_index_offset]$$

where the scaling table is defined as seen in Figure 29a or 29b (for 4:2:0 and 4:2:2 respectively), and “chroma_qp_index_offset” and “second_chroma_qp_index_offset” are defined in the picture parameter set and may be the same or different for Cr and Cb. In other words, the value in square brackets defines in each case an “index” into the scaling table (Figure 29a and b) and the scaling table then gives a revised value of Qp (“value”).

Note that “chroma_qp_index_offset” and “second_chroma_qp_index_offset” may instead be referred to as cb_qp_offset and cr_qp_offset respectively.

Chrominance channels typically contain less information than luminance and hence have smaller-magnitude coefficients; this limitation on the chrominance QP may prevent all chrominance detail being lost at heavy quantisation levels.

The QP-divisor relationship in the 4:2:0 is a logarithmic one such that an increase of 6 in the QP is equivalent to a doubling of the divisor (the quantisation step size discussed elsewhere in this description, though noting that it may be further modified by Qmatrices before use). Hence the largest difference in the scaling table of 51-39=12 represents a factor-of-4 change in the divisor.

However, in an embodiment, for the 4:2:2 scheme, which potentially contains twice as much chroma information as the 4:2:0 scheme, the maximum chrominance QP value in the scaling table may be raised to 45 (i.e. halving the divisor). Similarly for the 4:4:4 scheme, the maximum chrominance QP value in the scaling table may be raised to 51 (i.e. the same divisor). In this case the scaling table is in effect redundant, but may be retained simply for operational efficiency (i.e. so that the system works by reference to a table in the same way for each scheme). Hence more generally in an embodiment the chroma QP divisor is modified responsive to the amount of information in the coding scheme relative to the 4:2:0 scheme. Note that in other embodiments, the QP association may be such that a chrominance quantisation step size is less than a corresponding luminance quantisation step size for the 4:2:2 format but equal to the corresponding luminance quantisation step size for the 4:4:4 format over a predetermined range of QP values.

Figure 29c schematically illustrates a table mapping an input QP to a chroma QP. The input QP is based upon the luma QP:

$$inputQP = lumaQP + chromaQPoffset$$

The chromaQPoffset is fixed for the slice; lumaQP can vary at the CU level.

As additional background relating to the use of this table, note that the lumaQP can go as low as ‘-qpBdOffset’, where qpBdOffset is defined as (bitDepthLuma-8)*6. (bitDepth is the number of bits of the video data, representing the resolution of each sample in bits). However,

since bitDepth can in principle be different for the luma and chroma channels (though this is not an expected mode of operation of the system), the value used to form the inputQP is clipped to -qpBdOffsetChroma, so actually

```
    if (lumaQP < -qpBDOffsetChroma)
5      then inputQP = -qpBDOffsetChroma
      else inputQP = lumaQP + chromaQPOffset
```

Also note that the values used during the quantisation are lumaQP+qpBdOffsetLuma and chromaQP+qpBdOffsetChroma.

It can be seen that for QP values less than 28, the output QP equals the input QP.

10 For at least a predetermined range of QP values, the association of output QP to input QP is such that a chrominance quantisation step size is less than a corresponding luminance quantisation step size for the 4:2:2 format but equal to the corresponding luminance quantisation step size for the 4:4:4 format. That range as shown in Figure 29c extends to an upper limit corresponding to a 4:2:2 QP of 50.

15 The same mappings are used at the encoder and decoder sides.

Accordingly, embodiments apply to a video coding or decoding method operable to quantise blocks of frequency-transformed luminance and chrominance component video data in a 4:4:4 or a 4:2:2 format according to a selected quantisation parameter which defines a quantisation step size. A quantisation parameter association (such as, for example, the
20 appropriate table in Figure 29a, 29b or 29c) is defined between luminance and chrominance quantisation parameters, where the association is such that a maximum chrominance quantisation step size is less than a maximum luminance quantisation step size for the 4:2:2 format (for example, 45) but equal to the maximum luminance quantisation step size for the 4:4:4 format (for example, 51). Another example is where the association is such that a
25 chrominance quantisation step size is less than a corresponding luminance quantisation step size for the 4:2:2 format but equal to the corresponding luminance quantisation step size for the 4:4:4 format over a predetermined range of QP values. The quantisation process operates in that each component of the frequency-transformed data is divided by a respective value derived from the respective quantisation step size, and the result is rounded to an integer value, to
30 generate a corresponding block of quantised spatial frequency data.

It will be appreciated that the dividing and rounding steps are indicative examples of a generic quantising stage, according to the respective quantisation step size (or data derived from it, for example by the application of Qmatrices). Note that in this regard, the quantiser described above can act as an example of a divider and a rounder. The controller 343 can act
35 as an example of a controller.

At the corresponding decoder operations (dequantisation), the rounding step is not necessarily used, and the dividing step is replaced by a multiplying operation by a respective

value derived from the respective quantisation step size. In this regard, the dequantiser discussed above can act as an example of a multiplier. The controller 343 can act as an example of a controller.

The encoder can communicate the various data associated with the QP selection, such as first and second offsets in association with the encoded video data. The decoder side can receive such data in association with the encoded video data.

Embodiments include the step of selecting a quantisation parameter or index (QP for luminance) for quantising the spatial frequency coefficients, the quantisation parameter acting as a reference to a respective one of a set of quantisation step sizes according to the QP tables applicable to luminance data. The process of defining the quantisation parameter association can then comprise: for chrominance components, referencing a table of modified quantisation parameters (such as the table of Figures 29a, 29b or 29c) according to the selected quantisation parameter, which in turn can involve (i) for the first chrominance component, adding a first offset (such as `chroma_qp_index_offset`) to the quantisation parameter and selecting the modified quantisation index corresponding to the entry, in the table, for the quantisation index plus the first offset; and (ii) for the second chrominance component, adding a second offset (such as `second_chroma_qp_index_offset`) to the quantisation parameter and selecting the modified quantisation index corresponding to the entry, in the table, for the quantisation index plus the second offset; and referencing a respective quantisation step size in the set according to the quantisation parameter for the luminance data and the first and second modified quantisation indices for the first and second chrominance components. Viewed in a different way, this is an example of a process involving selecting a quantisation parameter for quantising the spatial frequency coefficients, the quantisation parameter acting as a reference to a respective one of a set of quantisation step sizes; and in which the defining step comprises: for chrominance components, referencing a table of modified quantisation parameters according to the selected quantisation parameter, the referencing step comprising: for each chrominance component, adding a respective offset to the quantisation parameter and selecting the modified quantisation parameter corresponding to the entry, in the table, for the quantisation parameter plus the respective offset; and referencing a respective quantisation step size in the set according to the quantisation parameter for the luminance data and the first and second modified quantisation parameters for the first and second chrominance components.

The techniques are particularly applicable to arrangements in which successive values of the quantisation step sizes in the set are related logarithmically, so that a change in quantisation parameter of m (where m is an integer) represents a change in quantisation step size by a factor of p (where p is an integer greater than 1). In the present embodiments, $m = 6$ and $p = 2$.

In embodiments, as discussed above, in at least some embodiments, a maximum luminance quantisation parameter is 51; a maximum chrominance quantisation parameter is 45 for the 4:2:2 format; and a maximum chrominance quantisation parameter is 51 for the 4:4:4 format.

5 In embodiments, the first and second offsets can be communicated in association with the encoded video data.

In 4:2:0 the transform matrices A are initially created (by the transform unit 340) from those of a true normalised $N \times N$ DCT A' using:

$$10 \quad A_{ij} = \text{int}(64 \times \sqrt{N} \times A'_{ij})$$

where i and j indicate a position within the matrix. This scaling with respect to a normalised transform matrix provides an increase in precision, avoids the need for fractional calculations and increases the internal precision.

15 Ignoring differences due to rounding of A_{ij} , since X is multiplied by both A and A^T (the transposition of the matrix A) the resulting coefficients differ from those of a true normalised $M \times N$ (M =height; N =width) DCT by a common scaling factor of:

$$(64 \times \sqrt{N})(64 \times \sqrt{M}) = 4096 \sqrt{N} \sqrt{M}$$

20 Note that the common scaling factor could be different to this example. Note also that matrix multiplying by both A and A^T can be carried out in various ways, such as the so-called Butterfly method. The significant fact is whether the operation that is carried out is equivalent to a traditional matrix multiplication, not whether it is performed in a particular traditional order of operations.

25 This scaling factor is equivalent to a binary left-shift bitwise operation by a number of bits *transformShift*, since in HEVC this results in a power of 2:

$$\text{transformShift} = (12 + 0.5 \log_2(N) + 0.5 \log_2(M))$$

To reduce the requirement on internal bit-precision, the coefficients are right-shifted (using positive rounding) twice during the transform process:

$$30 \quad \begin{aligned} \text{shift1} &= \log_2(N) + \text{bitDepth} - 9 \\ \text{shift2} &= \log_2(M) + 6 \end{aligned}$$

As a result, the coefficients as they leave the forward transform process and enter the quantiser are effectively left-shifted by:

$$35 \quad \begin{aligned} \text{resultingShift} &= (12 + 0.5 \log_2(NM)) - (\text{shift1} + \text{shift2}) \\ &= (12 + 0.5 \log_2(N) + 0.5 \log_2(M)) - (\log_2(N) + \text{bitDepth} - 9 + \log_2(M) + 6) \\ &= 15 - (0.5 \log_2(N) + 0.5 \log_2(M) + \text{bitDepth}) \end{aligned}$$

In 4:2:0, the frequency separated (for example, DCT) coefficients generated by the frequency transform are a factor of ($2^{\text{resultingShift}}$) larger than those that a normalised DCT would produce.

5 In some embodiments, the blocks are either square or rectangular with a 2:1 aspect ratio. Therefore, for a block size of $N \times M$, either:

$N=M$, in which case, *resultingShift* is an integer and $S=N=M=\text{sqrt}(NM)$; or

$0.5N=2M$ or $2N=0.5M$, in which case *resultingShift* is still an integer and $S=\text{sqrt}(NM)$

$$10 \quad \text{resultingShift} = 15 - (0.5\log_2(N) + 0.5\log_2(M) + \text{bitDepth}) = 15 - (\log_2(S) + \text{bitDepth})$$

The coefficients are subsequently quantised, where the quantising divisor is derived according to the quantisation parameter QP.

Note that *resultingShift* is equivalent to an integer, so the common scaling factor is an integer power of 2, the overall left-shift '*resultingShift*' of the transform process is also
15 accounted for at this stage by applying an equal but opposite right-shift, '*quantTransformRightShift*'.

This bit-shift operation is possible because *resultingShift* is an integer.

Also note that the divisor-QP (quantisation parameter or index) relationship follows a base-2 power curve, as mentioned above, in that an increase in QP by 6 has the effect of
20 doubling the divisor whereas an increase in QP by 3 has the effect of increasing the divisor by a factor of $\text{sqrt}(2)$ (square root of 2).

Due to the chroma format in 4:2:2, there are more TU width:height ($N:M$) ratios:

$N=M$ (from before) where $S=N=M=\text{sqrt}(NM)$ (*resultingShift* is an integer)

25 $0.5N=2M$ and $2N=0.5M$, (from before), where $S=\text{sqrt}(NM)$ (*resultingShift* is an integer)

$N=2M$ where $S=\text{sqrt}(NM)$

$2M=N$ where $S=\text{sqrt}(NM)$

30 $4N=0.5M$ where $S=\text{sqrt}(NM)$

$$\text{resultingShift} = 15 - (\log_2(S) + \text{bitDepth})$$

In these latter three situations, *resultingShift* is not an integer. For example, this may apply where at least some of the blocks of video data samples comprise $M \times N$ samples, where
35 the square root of N/M is not equal to an integer power of 2. Such block sizes can occur in respect of chroma samples in some of the present embodiments.

Accordingly, in such instances, the following techniques are relevant, that is to say, in video coding or decoding methods, apparatus or programs operable to generate blocks of quantised spatial frequency data by performing frequency-transformation on blocks of video data samples using a transform matrix comprising an array of integer values which are each scaled with respect to respective values of a normalized transform matrix by an amount dependent upon a dimension of the transform matrix, and to quantise the spatial frequency data according to a selected quantisation step size, having the step of frequency-transforming a block of video data samples by matrix-multiplying the block by the transform matrix and the transposition of the transform matrix to generate a block of scaled spatial frequency coefficients which are each larger, by a common scaling factor (for example, *resultingShift*), than the spatial frequency coefficients which would result from a normalized frequency-transformation of that block of video data samples.

Therefore at the quantisation stage, an appropriate bit-shift operation cannot be used to cancel out the operation in a simple manner.

A solution to this is proposed as follows:

At the quantiser stage, apply a right shift:

$$\text{quantTransformRightShift} = 15 - \log_2(S') - \text{bitDepth}$$

Where the value S' is derived such that

$$\text{resultingShift} - \text{quantTransformRightShift} = +\frac{1}{2}$$

$\text{quantTransformRightShift}$ is an integer

The difference between shifts of $\frac{1}{2}$ is equivalent to multiplication by $\sqrt{2}$, i.e. at this point the coefficients are $\sqrt{2}$ times larger than they should be, making the bit shift an integer bit shift.

For the quantisation process, apply a quantisation parameter of $(QP + 3)$, meaning that the quantising divisor is effectively increased by a factor of $\sqrt{2}$, thus cancelling out the $\sqrt{2}$ scale factor from the previous step.

Accordingly, these steps can be summarised (in the context of a video coding or decoding method (or corresponding apparatus or program) operable to generate blocks of quantised spatial frequency data by performing frequency-transformation on blocks of video data samples using a transform matrix comprising an array of integer values which are each scaled with respect to respective values of a normalized transform matrix, and to quantise the spatial frequency data according to a selected quantisation step size, involving frequency-transforming a block of video data samples by matrix-multiplying the block by the transform matrix and the transposition of the transform matrix to generate a block of scaled spatial frequency coefficients which are each larger, by a common scaling factor, than the spatial frequency coefficients which would result from a normalized frequency-transformation of that block of video data samples) as follows: selecting a quantisation step size for quantising the

spatial frequency coefficients; applying an n-bit shift (for example, *quantTransformRightShift*) to divide each of the scaled spatial frequency coefficients by a factor of 2^n , where n is an integer; and detecting a residual scaling factor (for example, *resultingShift- quantTransformRightShift*), being the common scaling factor divided by 2^n . For example, in the situation discussed above,

5 the quantisation step size is then according to the residual scaling factor to generate a modified quantisation step size; and each of the scaled spatial frequency coefficients in the block is divided by a value dependent upon the modified quantisation step size and rounding the result to an integer value, to generate the block of quantised spatial frequency data. As discussed, the modification of the quantisation step size can be carried out simply by adding an offset to

10 QP so as to select a different quantisation step size when QP is mapped into the table of quantisation step sizes.

The coefficients are now of the correct magnitude for the original QP.

The transform matrix can comprise an array of integer values which are each scaled with respect to respective values of a normalized transform matrix by an amount dependent upon a

15 dimension of the transform matrix.

It follows that the required value for S' can always be derived as follows:

$$S' = \sqrt{2 * M * N}$$

As an alternative proposal, S' could be derived such that:

$$\text{resultingShift- quantTransformRightShift} = -\frac{1}{2}$$

20 In this case, $S' = \sqrt{\frac{1}{2} * M * N}$, and the applied quantisation parameter is $(QP - 3)$

In either of these cases, (adding 3 to QP or subtracting 3 from QP), the step of selecting the quantisation step size comprises selecting a quantisation index (for example, QP), the quantisation index defining a respective entry in a table of quantisation step sizes, and the modifying step comprises changing the quantisation index so as to select a different

25 quantisation step size, such that the ratio of the different quantisation step size to the originally selected quantisation step size is substantially equal to the residual scaling factor.

This works particularly well where, as in the present embodiments, successive values of the quantisation step sizes in the table are related logarithmically, so that a change in quantisation index (for example, QP) of m (where m is an integer) represents a change in

30 quantisation step size by a factor of p (where p is an integer greater than 1). In the present embodiments, $m = 6$ and $p = 2$, so that an increase of 6 in QP represents a doubling of the applied quantisation step size, and a decrease in QP of 6 represents a halving of the resulting quantisation step size.

As discussed above, the modification can be carried out by selecting a quantisation

35 index (for example, a base QP) in respect of luminance samples; generating a quantisation index offset, relative to the quantisation index selected for the luminance samples, for samples of each or both chrominance components; changing the quantisation index offset according to

the residual scaling factor; and communicating the quantisation index offset in association with the coded video data. In embodiments of HEVC, QP offsets for the two chroma channels are sent in the bit stream. These steps correspond to a system in which the QP offset (to account for the residual scaling factor) of ± 3 could be incorporated into these offsets, or they could be

5 incremented/decremented when they are used to derive the chroma QP.

Note that the QP offset does not have to be ± 3 if differently shaped blocks were used; it is just that ± 3 represents an offset applicable to the block shapes and aspect ratios discussed above in respect of 4:2:2 video, for example.

In some embodiments, n (the bit shift as applied) is selected so that 2^n is greater than or equal to the common scaling factor. In other embodiments, n is selected so that 2^n is less than or equal to the common scaling factor. In embodiments (using either of these arrangements), a

10 bit shift n can be selected so as to be the next nearest (in either direction) to the common scaling factor, so that the residual scaling factor represents a factor having a magnitude of less than 2.

In other embodiments, the modification of the quantisation step size can simply be performed by multiplying the quantisation step size by a factor dependent upon the residual scaling factor. That is to say, the modification need not involve modifying the index QP.

Note also that the quantisation step size as discussed is not necessarily the actual quantisation step size by which a transformed sample is divided. The quantisation step size

20 derived in this way can be further modified. For example, in some arrangements, the quantisation step size is further modified by respective entries in a matrix of values (Qmatrix) so that different final quantisation step sizes are used at different coefficient positions in a quantised block of coefficients.

It is also notable that in the 4:2:0 scheme, the largest chroma TU is 16x16, whereas for

25 the 4:2:2 scheme 16x32 TUs are possible, and for the 4:4:4 scheme, 32x32 chroma TUs are possible. Consequently in an embodiment quantisation matrices (Qmatrices) for 32x32 chroma TUs are proposed. Similarly, Qmatrices should be defined for non-square TUs such as the 16x32 TU, with one embodiment being the subsampling of a larger square Q matrix

Qmatrices could be defined by any one of the following:

30 values in a grid (as for 4x4 and 8x8 Qmatrices);

interpolated spatially from respective smaller or larger matrices;

- in HEVC larger Qmatrices can be derived from respective groups of coefficients of smaller reference ones, or smaller matrices can be sub-sampled from larger matrices. Note that this interpolation or subsampling can be carried out within a channel ratio – for example, a

35 larger matrix for a channel ratio can be interpolated from a smaller one for that channel ratio.

relative to other Qmatrices (i.e. difference values, or deltas);

- hence only the deltas need to be sent.

Taking a small example just for illustrative purposes, a particular matrix for one channel ratio could be defined, such as a 4 x 4 matrix in respect of 4:2:0

$$\begin{pmatrix} a & b \end{pmatrix}$$

$$\begin{pmatrix} c & d \end{pmatrix}$$

5 where a, b, c and d are respective coefficients. This acts as a reference matrix.

Embodiments of the disclosure could then define a set of difference values for a similar-sized matrix in respect of another channel ratio:

$$\begin{pmatrix} \text{diff1} & \text{diff2} \end{pmatrix}$$

$$\begin{pmatrix} \text{diff3} & \text{diff4} \end{pmatrix}$$

10 so that in order to generate the Qmatrix for the other channel ratio, the matrix of differences is matrix-added to the reference matrix.

Instead of differences, a matrix of multiplicative factors could be defined for the other channel ratio, such that either (i) the matrix of multiplicative factors is matrix-multiplied with the reference matrix to generate the Qmatrix for the other channel ratio, or (ii) each coefficient in the reference matrix is individually multiplied by a respective factor to generate the Qmatrix for the other channel ratio.

as a function of another Qmatrix;

- e.g. a scaling ratio relative to another matrix (so that each of a, b, c and d in the above example is multiplied by the same factor, or has the same difference added to it). This reduces the data requirements for transmitting the difference or factor data.

- hence only the coefficients of the functions need to be sent (such as the scaling ratio),

as an equation/function (e.g. piece-wise linear curve, exponential, polynomial);

- hence only the coefficients of the equations need to be sent to derive the matrix, or any combination of the above. For example, each of a, b, c and d could in fact be defined by a function which could include a dependence upon the coefficient position (i,j) within the matrix. (i, j) could represent, for example, the coefficient position from left to right followed by the coefficient position from top to bottom of the matrix. An example is:

$$\text{coefficient}_{ij} = 3i + 2j$$

30 Note that Qmatrices can be referred to as Scaling Lists within the HEVC environment. In embodiments in which the quantisation is applied after the scanning process, the scanned data may be a linear stream of successive data samples. In such instances, the concept of a Qmatrix still applies, but the matrix (or Scanning List) may be considered as a 1xN matrix, such that the order of the N data values within the 1xN matrix corresponds to the order of scanned samples to which the respective Qmatrix value is to be applied. In other words, there is a 1:1 relationship between data order in the scanned data, spatial frequency according to the scan pattern, and data order in the 1xN Qmatrix.

Note that it is possible, in some implementations, to bypass or omit the DCT (frequency separation) stage, but to retain the quantisation stage.

Other useful information includes an optional indicator of to which other matrix the values are related, i.e. the previous channel or the first (primary) channel; for example the matrix for Cr could be a scaled factor of a matrix for Y, or for Cb, as indicated.

Accordingly, embodiments of the disclosure can provide a video coding or decoding method (and a corresponding apparatus or computer program) operable to generate blocks of quantised spatial frequency data by (optionally) performing frequency-transformation on blocks of video data samples and quantising the video data (such as the spatial frequency data) according to a selected quantisation step size and a matrix of data modifying the quantisation step size for use at different respective block positions within an ordered block of samples (such as an ordered block of frequency-transformed samples), the method being operable with respect to at least two different chrominance subsampling formats.

For at least one of the chrominance subsampling formats, one or more quantisation matrices are defined as one or more predetermined modifications with respect to one or more reference quantisation matrices defined for a reference one of the chrominance subsampling formats.

In embodiments of the disclosure, the defining step comprises defining one or more quantisation matrices as a matrix of values each interpolated from a respective plurality of values of a reference quantisation matrix. In other embodiments, the defining step comprises defining one or more quantisation matrices as a matrix of values each subsampled from values of a reference quantisation matrix.

In embodiments of the disclosure, the defining step comprises defining one or more quantisation matrices as a matrix of differences with respect to corresponding values of a reference quantisation matrix.

In embodiments of the disclosure, the defining step comprises defining one or more quantisation matrices as a predetermined function of values of a reference quantisation matrix. In such instances, the predetermined function may be a polynomial function.

In embodiments of the disclosure, one or both of the following is provided, for example as part of or in association with the coded video data: (i) reference-indicator data to indicate, with respect to encoded video data, the reference quantisation matrix; and (ii) modification-indicator data to indicate, with respect to encoded data values, the one or more predetermined modifications.

These techniques are particularly applicable where two of the chrominance subsampling formats are 4:4:4 and 4:2:2 formats.

The number of Q Matrices in HEVC 4:2:0 is currently 6 for each transform size: 3 for the corresponding channels, and one set for intra and for inter. In the case of a 4:4:4 GBR scheme,

it will be appreciated that either one set of quantisation matrices could be used for all channels, or three respective sets of quantisation matrices could be used.

In embodiments of the disclosure, at least one of the matrices is a $1 \times N$ matrix. This would be the case in (as described here) one or more of the matrices is in fact a Scaling List or the like, being a linear $1 \times N$ ordered array of coefficients.

The proposed solutions involve incrementing or decrementing the applied QP. However this could be achieved in a number of ways:

In HEVC, QP offsets for the two chroma channels are sent in the bit stream. The ± 3 could be incorporated into these offsets, or they could be incremented/decremented when they are used to derive the chroma QP.

As discussed, above, in HEVC, (luma QP + chroma offset) is used as an index to a table in order to derive the chroma QP. This table could be modified to incorporate the ± 3 (i.e. by incrementing/decrementing the values of the original table by 3)

After the chroma QP has been derived, as per the normal HEVC process, the results could then be incremented (or decremented) by 3.

As an alternative to modifying the QP, a factor of $\sqrt{2}$ or $1/\sqrt{2}$ can be used to modify the quantisation coefficients.

For forward/inverse quantisation, the division/multiplication processes are implemented by using $(QP \% 6)$ as an index to a table to obtain a quantisation coefficient or quantisation step size, $\text{inverseQStep}/\text{scaledQStep}$. (Here, $QP \% 6$ signifies QP modulo 6). Note that, as discussed above, this may not represent the final quantisation step size which is applied to the transformed data; it may be further modified by the Qmatrices before use.

The default tables in HEVC are of length 6, covering an octave (a doubling) of values. This is simply a means of reducing storage requirements; the tables are extended for actual use by selecting an entry in the table according to the modulus of QP (mod 6) and then multiplying or dividing by an appropriate power of 2, dependent upon the difference of $(QP - QP \text{ modulus } 6)$ from a predetermined base value.

This arrangement could be varied to allow for the offset of ± 3 in the QP value. The offset can be applied in the table look-up process, or the modulus process discussed above could instead be carried out using the modified QP. Assuming the offset is applied at the table look-up, however, additional entries in the table can be provided as follows:

One alternative is to extend the tables by 3 entries, where the new entries are as follows (for the index values of 6-8).

The example table shown in Figure 30 would be indexed by $[(QP \% 6) + 3]$ (a "QP increment method"), where the notation $QP \% 6$ signifies "QP modulus 6".

The example table shown in Figure 31 would be indexed by $[(QP \% 6) - 3]$ (a "QP decrement method"), having extra entries for the index values of -1 to -3:

Entropy encoding

Basic entropy encoding comprises assigning codewords to input data symbols, where the shortest available codewords are assigned to the most probable symbols in the input data. On average the result is a lossless but much smaller representation of the input data.

5 This basic scheme can be improved upon further by recognising that symbol probability is often conditional on recent prior data, and consequently making the assignment process context adaptive.

In such a scheme, context variables (CVs) are used to determine the choice of respective probability models, and such CVs are provided for in the HEVC 4:2:0 scheme.

10 To extend entropy encoding to the 4:2:2 scheme, which for example will use 4x8 chroma TUs rather than 4x4 TUs for an 8x8 luma TU, optionally the context variables can be provided for by simply vertically repeating the equivalent CV selections.

However, in an embodiment of the present disclosure the CV selections are not repeated for the top-left coefficients (i.e. the high-energy, DC and/or low spatial frequency coefficients), and instead new CVs are derived. In this case, for example, a mapping may be
15 derived from the luma map. This approach may also be used for the 4:4:4 scheme.

During coding, in the 4:2:0 scheme, a so-called zig-scan scans through the coefficients in order from high to low frequencies. However, again it is noted that the chroma TUs in the 4:2:2 scheme can be non-square, and so in an embodiment of the present disclosure a different
20 chroma scan is proposed with the angle of the scan be tilted to make it more horizontal, or more generally, responsive to the aspect ratio of the TU.

Similarly, the neighbourhood for significance map CV selection and the c1/c2 system for greater-than-one and greater-than-two CV selection may be adapted accordingly.

Likewise, in an embodiment of the present disclosure the last significant coefficient
25 position (which becomes the start point during decoding) could also be adjusted for the 4:4:4 scheme, with last-significant positions for chroma TUs being coded differentially from the last-significant position in the co-located luma TU.

The coefficient scanning can also be made prediction mode dependent for certain TU sizes. Hence a different scan order can be used for some TU sizes dependent on the intra-
30 prediction mode.

In the 4:2:0 scheme, mode dependent coefficient scanning (MDCS) is only applied for 4x4/8x8 luma TUs and 4x4 chroma TUs for intra prediction. MDCS is used dependent on the intra-prediction mode, with angles +/-4 from the horizontal and vertical being considered.

In an embodiment of the present disclosure, it is proposed that in the 4:2:2 scheme
35 MDCS is applied to 4x8 and 8x4 chroma TUs for intra prediction. Similarly, it is proposed that in the 4:4:4 scheme MDCS is applied to 8x8 and 4x4 chroma TUs. MDCS for 4:2:2 may only be

done in the horizontal or vertical directions, and that the angle ranges may differ for 4:4:4 chroma vs. 4:4:4 luma vs. 4:2:2 chroma vs. 4:2:2 luma vs. 4:2:0 luma.

In-loop filters

Deblocking

5 Deblocking is applied to all CU, PU and TU boundaries, and the CU/PU/TU shape is not taken into account. The filter strength and size is dependent on local statistics, and deblocking has a granularity of 8x8 Luma pixels.

Consequently it is anticipated that the current deblocking applied for the 4:2:0 scheme should also be applicable for the 4:2:2 and 4:4:4 schemes.

10 Sample adaptive offsetting

In sample adaptive offsetting (SAO) each channel is completely independent. SAO splits the image data for each channel using a quad-tree, and the resulting blocks are at least one LCU in size. The leaf blocks are aligned to LCU boundaries and each leaf can run in one of three modes, as determined by the encoder ("Central band offset", "Side band offset" or "Edge offset"). Each leaf categorises its pixels, and the encoder derives an offset value for each of the 16 categories by comparing the SAO input data to the source data. These offsets are sent to the decoder. The offset for a decoded pixel's category is added to its value to minimise the deviation from the source.

15 In addition, SAO is enabled or disabled at picture level; if enabled for luma, it can also be enabled separately for each chroma channel. SAO will therefore be applied to chroma only if it is applied to luma.

Consequently the process is largely transparent to the underlying block scheme and it is anticipated that the current SAO applied for the 4:2:0 scheme should also be applicable for the 4:2:2 and 4:4:4 schemes.

25 Adaptive loop filtering

In the 4:2:0 scheme, adaptive loop filtering (ALF) is disabled by default. However, in principle (i.e. if allowed) then ALF would be applied to the entire picture for chroma.

In ALF, luma samples may be sorted into one of a number of categories, as determined by the HEVC documents; each category uses a different Wiener-based filter.

30 By contrast, in 4:2:0 chroma samples are not categorised – there is just one Wiener-based filter for Cb, and one for Cr.

Hence in an embodiment of the present disclosure, in light of the increased chroma information in the 4:2:2 and 4:4:4 schemes, it is proposed that the chroma samples are categorised; for example with K categories for 4:2:2 and J categories for 4:4:4.

35 Whilst in the 4:2:0 scheme ALF can be disabled for luma on a per-CU basis using an ALF control flag (down to the CU-level specified by the ALF control depth), it can only be

disabled for chroma on a per-picture basis. Note that in HEVC, this depth is currently limited to the LCU level only.

Consequently in an embodiment of the present disclosure, the 4:2:2 and 4:4:4 schemes are provided with one or two channel specific ALF control flags for chroma.

5 Syntax

In HEVC, syntax is already present to indicate 4:2:0, 4:2:2 or 4:4:4 schemes, and is indicated at the sequence level. However, in an embodiment of the present disclosure it is proposed to also indicate 4:4:4 GBR coding at this level.

Data Signals

10 It will be appreciated that data signals generated by the variants of coding apparatus discussed above, and storage or transmission media carrying such signals, are considered to represent embodiments of the present disclosure.

CLAIMS

1. A video decoding apparatus operable to dequantise blocks of frequency-transformed luminance and chrominance component video data in a 4:4:4 or a 4:2:2 format according to a selected quantisation parameter which defines a quantisation step size, the apparatus comprising:

a controller configured to define a quantisation parameter association between luminance and chrominance quantisation parameters, where the association is such that a chrominance quantisation step size is less than a corresponding luminance quantisation step size for the 4:2:2 format but equal to the corresponding luminance quantisation step size for the 4:4:4 format, over at least a predetermined range of quantization parameters; and

a multiplier configured to multiply each component of the frequency-transformed data by a respective value derived from the respective quantisation step size.

2. A video coding apparatus operable to quantise blocks of frequency-transformed luminance and chrominance component video data in a 4:4:4 or a 4:2:2 format according to a selected quantisation parameter which defines a quantisation step size, the apparatus comprising:

a controller configured to define a quantisation parameter association between luminance and chrominance quantisation parameters, where the association is such that a chrominance quantisation step size is less than a corresponding luminance quantisation step size for the 4:2:2 format but equal to the corresponding luminance quantisation step size for the 4:4:4 format, over at least a predetermined range of quantization parameters;

a divider configured to divide each component of the frequency-transformed data by a respective value derived from the respective quantisation step size; and

a rounder configured to round the result to an integer value, to generate a corresponding block of quantised spatial frequency data.

3. Apparatus according to claim 1 or claim 2, in which the controller is configured, for chrominance components, to reference a table of modified quantisation parameters according to a selected quantisation parameter, the quantisation parameter acting as a reference to a respective one of a set of quantisation step sizes, the referencing comprising:

for each chrominance component, adding a respective offset to the quantisation parameter and selecting the modified quantisation parameter corresponding to the entry, in the table, for the quantisation parameter plus the respective offset; and

referencing a respective quantisation step size in the set according to the quantisation parameter for the luminance data and the first and second modified quantisation parameters for the first and second chrominance components.

4. Apparatus according to claim 3, in which successive values of the quantisation step sizes in the set are related logarithmically, so that a change in quantisation parameter of m (where m is an integer) represents a change in quantisation step size by a factor of p (where p is an integer greater than 1).

5

5. Apparatus according to claim 4, in which $m = 6$ and $p = 2$.

6. Apparatus according to any one of the preceding claims, in which the association is such that a maximum chrominance quantisation step size is less than a maximum luminance quantisation step size for the 4:2:2 format but equal to the maximum luminance quantisation step size for the 4:4:4 format.

10

7. Apparatus according to claim 6, in which:

a maximum luminance quantisation parameter is 51;

15

a maximum chrominance quantisation parameter is 45 for the 4:2:2 format; and

a maximum chrominance quantisation parameter is 51 for the 4:4:4 format.

8. Apparatus according to any one of claims 1 to 5, in which the predetermined range has an upper limit corresponding to a 4:2:2 chrominance quantisation parameter of 50.

20

9. Apparatus according to any one of claims 3 to 8, the apparatus being configured to communicate or to receive the first and second offsets in association with the encoded video data.

10. A video decoding method operable to dequantise blocks of frequency-transformed luminance and chrominance component video data in a 4:4:4 or a 4:2:2 format according to a selected quantisation parameter which defines a quantisation step size, the method comprising:

25

defining a quantisation parameter association between luminance and chrominance quantisation parameters, where the association is such that a chrominance quantisation step size is less than a corresponding luminance quantisation step size for the 4:2:2 format but equal to the corresponding luminance quantisation step size for the 4:4:4 format, over at least a predetermined range of quantization parameters; and

30

multiplying each component of the frequency-transformed data by a respective value derived from the respective quantisation step size.

11. A video coding method operable to quantise blocks of frequency-transformed luminance and chrominance component video data in a 4:4:4 or a 4:2:2 format according to a selected quantisation parameter which defines a quantisation step size, the method comprising:

35

defining a quantisation parameter association between luminance and chrominance quantisation parameters, where the association is such that a chrominance quantisation step size is less than a corresponding luminance quantisation step size for the 4:2:2 format but equal to the corresponding luminance quantisation step size for the 4:4:4 format, over at least a predetermined range of quantization parameters;

quantising each component of the frequency-transformed data according to a respective value derived from the respective quantisation step size; and

rounding the result to an integer value, to generate a corresponding block of quantised spatial frequency data.

12. A method according to claim 10 or claim 11, in which the defining step comprises:

for chrominance components, referencing a table of modified quantisation parameters according to the selected quantisation parameter, the quantisation parameter acting as a reference to a respective one of a set of quantisation step sizes, the referencing step comprising:

for each chrominance component, adding a respective offset to the quantisation parameter and selecting the modified quantisation parameter corresponding to the entry, in the table, for the quantisation parameter plus the respective offset; and

referencing a respective quantisation step size in the set according to the quantisation parameter for the luminance data and the first and second modified quantisation parameters for the first and second chrominance components.

13. A method according to claim 12, in which successive values of the quantisation step sizes in the set are related logarithmically, so that a change in quantisation parameter of m (where m is an integer) represents a change in quantisation step size by a factor of p (where p is an integer greater than 1).

14. A method according to claim 13, in which $m = 6$ and $p = 2$.

15. A method according to any one of claims 10 to 14, in which the association is such that a maximum chrominance quantisation step size is less than a maximum luminance quantisation step size for the 4:2:2 format but equal to the maximum luminance quantisation step size for the 4:4:4 format.

16. A method according to claim 16, in which:

a maximum luminance quantisation parameter is 51;

a maximum chrominance quantisation parameter is 45 for the 4:2:2 format; and

a maximum chrominance quantisation parameter is 51 for the 4:4:4 format.

17. A method according to any one of claims 10 to 14, in which the predetermined range has an upper limit corresponding to a 4:2:2 chrominance quantisation parameter of 50.
18. A method according to any one of claims 10 to 17, comprising the step of communicating or
5 receiving the first and second offsets in association with the encoded video data.
19. Computer software which, when executed by a computer, causes the computer to carry out a method according to any one of claims 10 to 18.
- 10 20. A machine-readable non-transitory storage medium which stores software according to claim 19.
21. A data signal comprising coded data generated according to the method of any one of claims 10 to 18.
15
22. Video storage, capture, display, transmission and/or reception apparatus comprising apparatus according to any one of claims 1 to 9.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8




FIG. 9




FIG. 10

PU	PU	PU							
PU	PU	PU							
PU (MC)	PU (MC)	PU		PU		PU			
				PU	PU	PU			
				PU	PU				
		PU	PU	PU	PU	PU	PU	PU	PU
				PU	PU	PU	PU	PU	PU
		PU	PU	PU	PU	PU		PU	PU
PU	PU					PU	PU		

FIG. 11

					TU		SDIP-P	SDIP-P	SDIP-P	SDIP-P
					TU	TU				
					TU					
TU					TU					
TU					TU					TU
TU					TU	TU				
TU										
TU										

FIG. 12




FIG. 13




FIG. 16




FIG. 14




FIG. 15





FIG. 17

1140		1150		
CU Size	PU options for given CU size All TUs within a PU have the same prediction modes	TU options, given PU and CU size TUs can be recursively split down to 2 levels		
		LEVEL 0	LEVEL 1	LEVEL 2
64x64 Y	64x64 Y			
		NSQT enabled: (use 64x64 case when disabled)		
		NSQT enabled: (use 64x64 case when disabled)		
32x32 Y	32x32 Y			
		NSQT Enabled: (use 32x32 case when disabled)		
		NSQT enabled: (use 32x32 case when disabled)		

Chroma channel block sizes cover the same area as the corresponding Luma block sizes (unless shown otherwise), but the number of samples depends on the format (4:2:2, 4:2:0)

For 4:4:4, the Chroma channels have the same block structure as the Luma structure shown.
For 4:2:2, the rectangular TUs may be split further into square TUs prior to transform/quantisation.

FIG. 18A



Chroma channel block sizes cover the same area as the corresponding Luma block sizes (unless shown otherwise), but the number of samples depends on the format (4:2:2, 4:2:0)

For 4:4:4, the Chroma channels have the same block structure as the Luma structure shown.

For 4:2:2, the rectangular TUs may be split further into square TUs prior to transform/quantisation.

FIG. 18B




FIG. 19

FIG. 20




FIG. 21




FIG. 22




FIG. 23




FIG. 24

FIG. 25




FIG. 26




FIG. 27





FIG. 28



Index	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
Value	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25

Index	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51
Value	26	27	28	29	29	30	31	32	32	33	34	34	35	35	36	36	37	37	37	38	38	38	39	39	39	39

(4:2:0)

FIG. 29A

Index	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
Value	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25

Index	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51
Value	26	27	28	29	30	31	32	33	33	34	35	36	37	37	38	39	40	40	41	42	42	43	44	44	45	45

(4:2:2)

FIG. 29B

Index	0	1	2	3	4	5	6	7	8
inverseQStep	26214	23302	20560	18396	16384	14564	12945	11523	10280
scaledQStep	40	45	51	57	64	72	81	91	102

FIG. 30

Index	-3	-2	-1	0	1	2	3	4	5
inverseQStep	36158	32768	29127	26214	23302	20560	18396	16384	14564
scaledQStep	29	32	36	40	45	51	57	64	72

FIG. 31

Input QP	Output QP			
	4:2:0	4:2:2	4:4:4	
Q	Q	Q	Q	for any input QP, Q, less than 28
28	28	28	28	
29	29	29	29	
30	29	30	30	*start of non-linear section
31	30	31	31	
32	31	32	32	
33	32	32	33	
34	33	33	34	
35	33	34	35	
36	34	35	36	
37	34	36	37	
38	35	37	38	
39	35	37	39	
40	36	38	40	
41	36	39	41	
42	37	40	42	
43	37	40	43	*start of linear section for 4:2:0
44	38	41	44	
45	39	42	45	
46	40	42	46	
47	41	43	47	
48	42	44	48	
49	43	44	49	
50	44	45	50	
51	45	45	51	*start of linear section for 4:2:2
52	46	46	51	
53	47	47	51	
54	48	48	51	
55	49	49	51	
56	50	50	51	
57	51	51	51	
>57	51	51	51	

FIG. 29C

INTERNATIONAL SEARCH REPORT

International application No
PCT/GB2013/051074

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04N7/26
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04N

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

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

EPO-Internal, WPI Data, INSPEC, COMPENDEX

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	LU: "Quantization Weighting Proposal for PExt", 11. JVT MEETING; 68. MPEG MEETING; 15-03-2004 - 19-03-2004; MUNICH,DE; (JOINT VIDEO TEAM OF ISO/IEC JTC1/SC29/WG11 AND ITU-T SG.16),, no. JVT-K029r1, 19 March 2004 (2004-03-19) , XP030005846, ISSN: 0000-0421 Sections 1-4 ----- -/--	1-22



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

12 July 2013

Date of mailing of the international search report

06/08/2013

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Morbee, Marleen

INTERNATIONAL SEARCH REPORT

International application No

PCT/GB2013/051074

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	SULLIVAN: "Proposed Draft Errata List with Revision-Marked Corrections for H.264/AVC", 10. JVT MEETING; 67. MPEG MEETING; 08-12-2003 - 12-12-2003; WAIKOLOA,US; (JOINT VIDEO TEAM OF ISO/IEC JTC1/SC29/WG11 AND ITU-T SG.16),, no. JVT-J010d7, 12 December 2003 (2003-12-12), XP030005787, ISSN: 0000-0422 Section 8.5.5 -----	1-22
A	US 2005/271288 A1 (SUZUKI TERUHIKO [JP] ET AL) 8 December 2005 (2005-12-08) paragraph [0022] - paragraph [0081] paragraph [0104] - paragraph [0241] -----	1-22
A	XU (SONY) J ET AL: "Chroma QP extension and signalling enhancement", 9. JCT-VC MEETING; 100. MPEG MEETING; 27-4-2012 - 7-5-2012; GENEVA; (JOINT COLLABORATIVE TEAM ON VIDEO CODING OF ISO/IEC JTC1/SC29/WG11 AND ITU-T SG.16); URL: HTTP://WFTP3.ITU.INT/AV-ARCH/JCTVC-SITE/,, no. JCTVC-I0265, 17 April 2012 (2012-04-17), XP030112028, abstract Sections 1, 2 -----	1-22
X,P	SILCOCK P ET AL: "AHG12: Extension of HM7 to Support Additional Chroma Formats", 10. JCT-VC MEETING; 101. MPEG MEETING; 11-7-2012 - 20-7-2012; STOCKHOLM; (JOINT COLLABORATIVE TEAM ON VIDEO CODING OF ISO/IEC JTC1/SC29/WG11 AND ITU-T SG.16); URL: HTTP://WFTP3.ITU.INT/AV-ARCH/JCTVC-SITE/,, no. JCTVC-J0191, 2 July 2012 (2012-07-02), XP030112553, Sections 1-4; the whole document -----	1-22

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/GB2013/051074

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2005271288	A1	08-12-2005	
		BR PI0406183 A	19-07-2005
		CN 1701616 A	23-11-2005
		CN 101646085 A	10-02-2010
		EP 1648180 A1	19-04-2006
		EP 2506574 A2	03-10-2012
		EP 2506575 A2	03-10-2012
		EP 2506576 A2	03-10-2012
		EP 2506577 A2	03-10-2012
		EP 2506578 A2	03-10-2012
		EP 2506579 A2	03-10-2012
		JP 4617644 B2	26-01-2011
		JP 2005039743 A	10-02-2005
		KR 20060121651 A	29-11-2006
		MX PA05002671 A	05-05-2005
		RU 2350041 C2	20-03-2009
		US 2005271288 A1	08-12-2005
		US 2009190829 A1	30-07-2009
		US 2011122947 A1	26-05-2011
		US 2011123103 A1	26-05-2011
		US 2011123104 A1	26-05-2011
		US 2011123105 A1	26-05-2011
		US 2011123106 A1	26-05-2011
		US 2011123107 A1	26-05-2011
		US 2011123108 A1	26-05-2011
		US 2011123109 A1	26-05-2011
		WO 2005009050 A1	27-01-2005