

(21) Application No: 1319493.1
 (22) Date of Filing: 04.11.2013

(51) INT CL: H04N 19/132 (2014.01)
 (56) Documents Cited: WO 2013/129203 A1
 (58) Field of Search: INT CL H04N
 Other: Online: WPI, EPODOC, Inspec

(71) Applicant(s):
British Broadcasting Corporation
(Incorporated in the United Kingdom)
Broadcasting House, LONDON, W1A 1AA,
United Kingdom

(72) Inventor(s):
Marta Mrak
Saverio Giovanni Blasi
Ebroul Izquierdo

(74) Agent and/or Address for Service:
Reddie & Grose
16 Theobalds Road, LONDON, WC1X 8PL,
United Kingdom

(54) Title of the Invention: **An improved compression algorithm for video compression codecs**
 Abstract Title: **Generating or Modifying Coefficients in a Region of a Prediction Block According to Substitution**

(57) Encoding video data comprises generating or receiving a prediction block containing coefficients representative of an original picture block. A residual is calculated by subtracting the original picture block and the prediction block from each other. The prediction block is generated or modified by generating coefficients at a first set of locations (first region) which are preserved or generated according to usual intra or inter prediction coding while generating coefficients at a second set of locations (second region) according to substitution with other values. The prediction coefficients in locations with very low correlation with the original can therefore be discarded and replaced with more informative content. The selection of locations / regions may be performed using a mask, the optimal mask being selected from a plurality of shapes. The modified coefficients improve the efficiency of the compression process once the prediction block is subtracted from the original block to generate a residual signal for quantisation and encoding. The processes may operate in either of the frequency or the spatial domain. The prediction block may also be modified according to its size and/or shape.

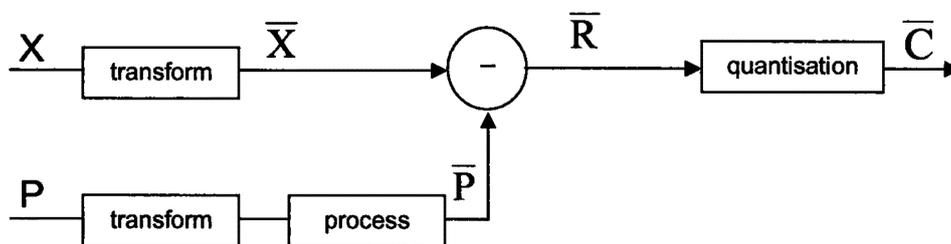


Figure 7

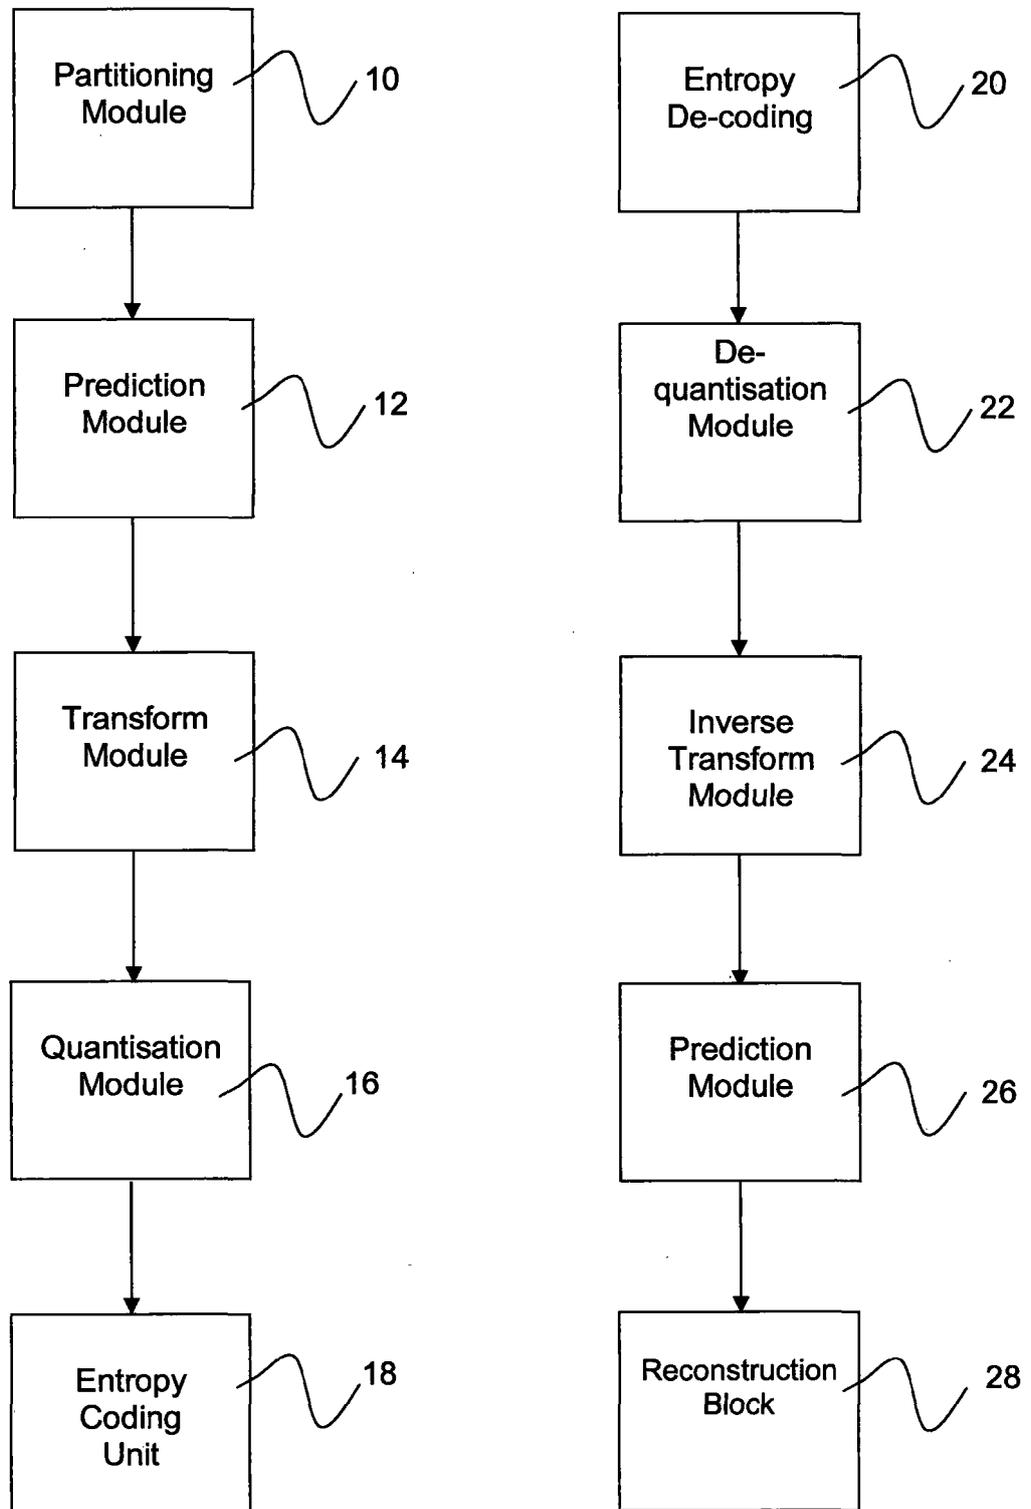


Figure 1

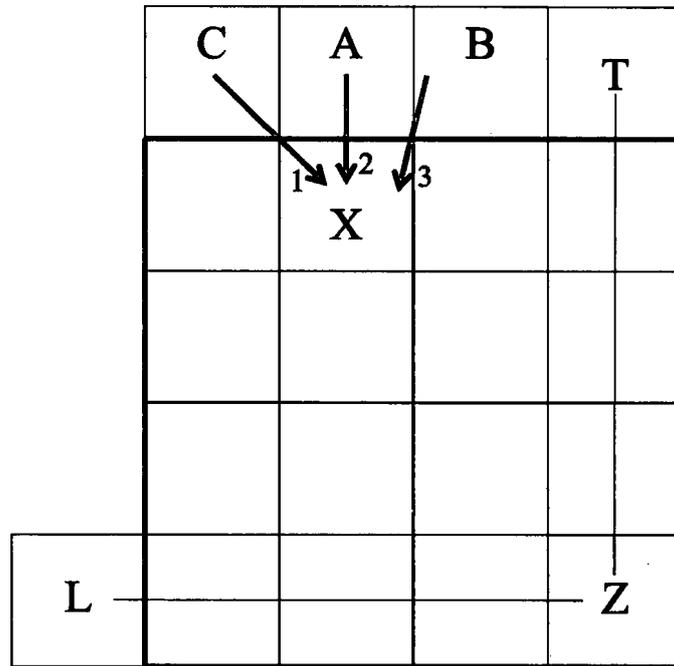


Figure 2

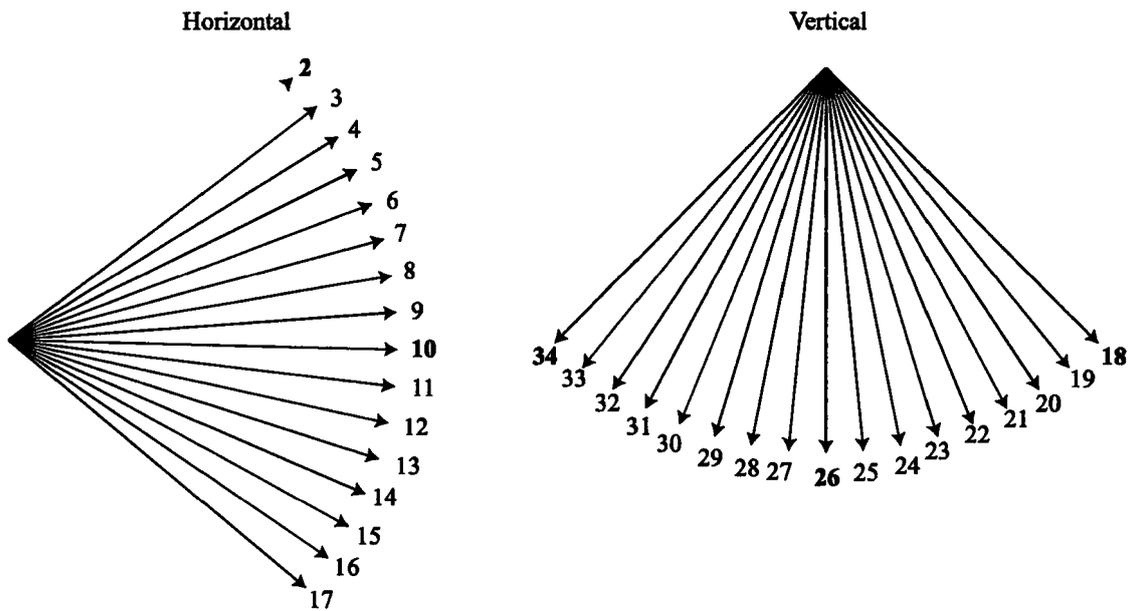


Figure 3

64	64	64	64	64	64	64	64
89	75	50	18	-18	-50	-75	-89
83	36	-36	-83	-83	-36	36	83
75	-18	-89	-50	50	89	18	-75
64	-64	-64	64	64	-64	-64	64
50	-89	18	75	-75	-18	89	-50
36	-83	83	-36	-36	83	-83	36
18	-50	75	-89	89	-75	50	-18

Figure 4

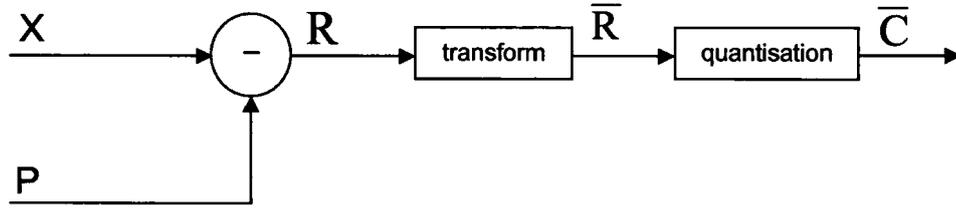


Figure 5

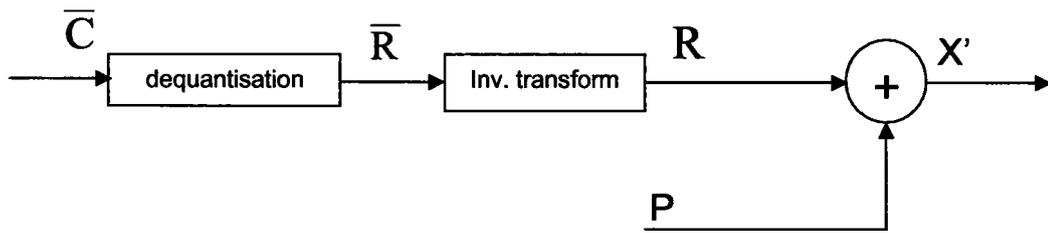


Figure 6

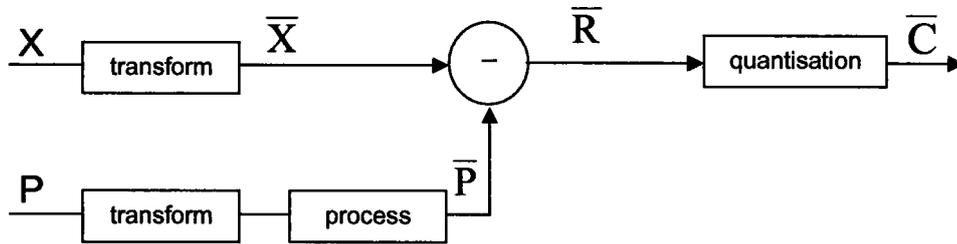


Figure 7

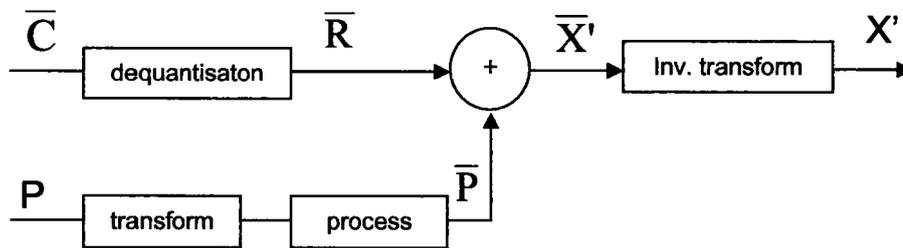
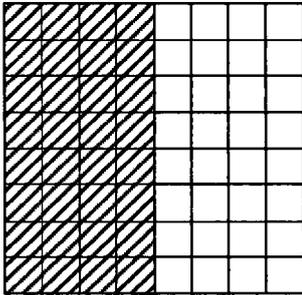
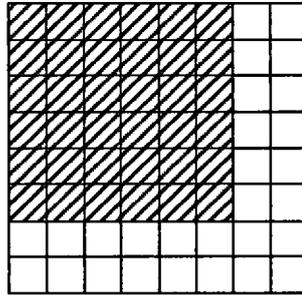


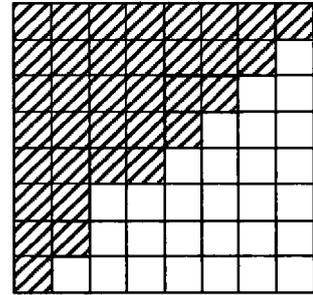
Figure 8



a)



b)



c)

Figure 9

Index	Value
0	N/A
1	1
2	-1
3	2
4	-2
5	4
6	-4
7	8
8	-8
9	16
10	-16
11	32
12	-32

Figure 10

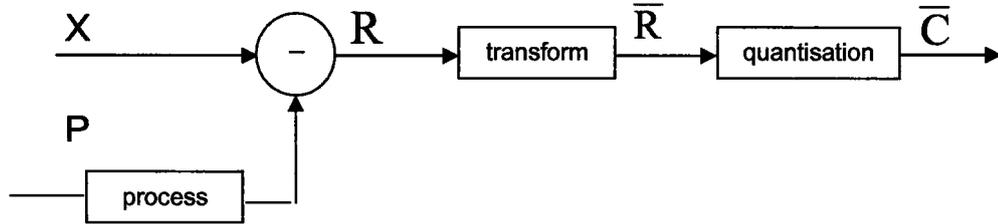
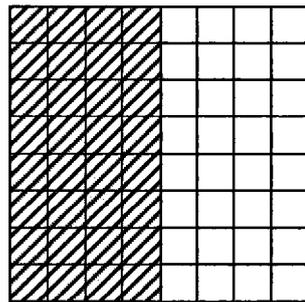
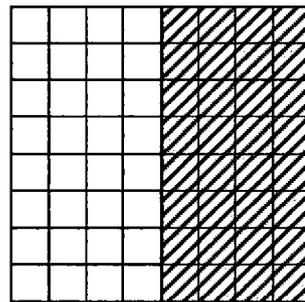


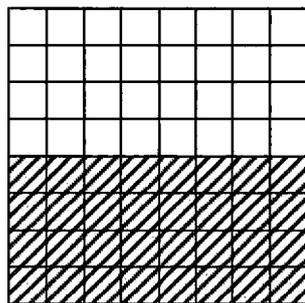
Figure 11



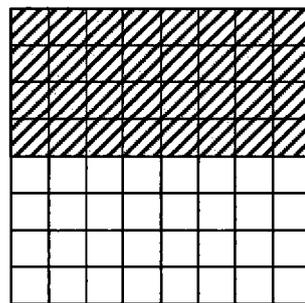
a)



b)



c)



d)

Figure 12

Index	Value
0	N/A
1	15
2	31
3	47
4	64
5	107
6	163
7	191
8	255

Figure 13

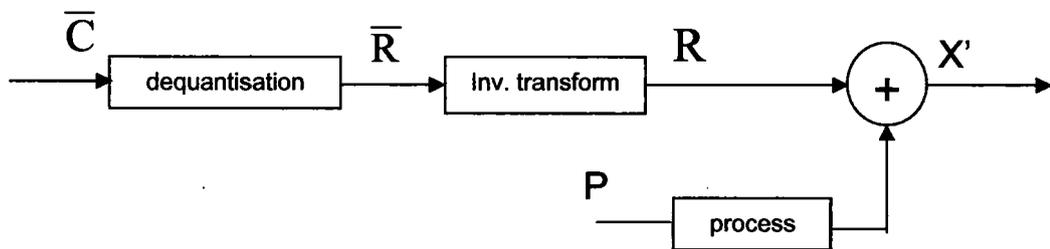


Figure 14

AN IMPROVED COMPRESSION ALGORITHM FOR VIDEO COMPRESSION CODECS

5 The invention relates to an improved compression algorithm for video compression codecs, and in particular to a method, apparatus and computer program for implementing the method.

 Known video coding standards are highly efficient when employed in applications where low to medium quality is acceptable, such coding for conventional internet or
10 computer based viewing, and most effort in the video coding community is dedicated to improving the efficiency of video coders at these levels of quality. The H.265/HEVC (High Efficiency Video Coding) standard for example is reported as achieving more than 50% higher efficiency than its predecessor H.264/AVC (Advanced Video Coding) for such applications.

15 Unfortunately while such levels of quality are acceptable for many purposes, there are many applications where higher levels of quality are necessary. In these cases the decoded video at the receiving end is required to be as faithful to the original video prior to the encoding as possible. Typical examples of such kind of applications can be found in
20 medical imaging applications, in the transmission of signals from cameras throughout the production chain, in-screen mirroring systems (when the content on the screen of a device is mirrored in real-time to a different screen), and so on. Moreover with the increasing diffusion of high-definition televisions capable of handling very high frame rates and of displaying content at high bit depths, the quality of decoded videos is becoming an extremely important issue even in the context of consumer applications. Users want new
25 broadcasting systems to deliver videos as accurately as possible, with the lowest visible errors. Under these quality constraints, the efficiency of HEVC with respect to its predecessor AVC has been found to decrease considerably.

 Moreover, conventional prediction methods, used in the above coding standards to provide compression, are based on the minimization of distortion metrics computed in the
30 spatial domain (such as the sum of absolute differences, SAD). Such metrics do not always take into account the accuracy of different prediction modes on the frequency components of the signal.

 We have therefore appreciated that it would be desirable to provide an approach for efficient high quality video coding that seeks to address these problems.

SUMMARY OF THE INVENTION

5 The invention is defined in the independent claims to which reference should now be made. Advantageous features are set forth in the dependent claims.

BRIEF DESCRIPTION OF THE DRAWINGS

10 Embodiments of the invention will now be described by way of example and with reference to the drawings in which:

 Figure 1 is a schematic illustration illustrating an encoding and decoding process;

 Figure 2 is a schematic illustration of a intra-angular and intra-planar prediction processes;

15 Figure 3 is an illustration of angular intra prediction modes available in HEVC (modes indicated in bold and dashed arrows are obtained by direct extrapolation, whereas all other modes are obtained by bilinear interpolation);

 Figure 4 is an illustration of a transform matrix, where each row represents a transform vector for an 8-point transform;

20 Figure 5 is a schematic illustration of a known encoder;

 Figure 6 is a schematic illustration of a known decoder;

 Figure 7 is a schematic illustration of a modified encoder according to a first example of the invention;

25 Figure 8 is a schematic illustration of a modified decoder according to a first example of the invention;

 Figure 9 is an illustration of different substitution patterns that can be applied to the prediction block in a first example of the invention;

 Figure 10 is an illustration of a look up table used in a coefficient substitution process;

30 Figure 11 is a schematic illustration of a modified encoder according to a further example of the invention;

 Figure 12 is an illustration of different substitution patterns that can be applied to the prediction block in a further example of the invention;

35 Figure 13 is an illustration of a look up table used in a coefficient substitution process; and

Figure 14 is a schematic illustration of a modified decoder according to a further example of the invention.

DETAILED DESCRIPTION OF THE EXAMPLE EMBODIMENTS

5 Most of the video compression standards available today, including HEVC and AVC, follow a block based hybrid scheme involving a number of successive stages. Referring to Figure 1, the current picture or frame in a sequence of pictures is passed to a partitioning module 10 where it is partitioned into blocks of a given size for sequential processing by the encoder. Each block is then input to a prediction module 12, which seeks
10 to discard temporal and spatial redundancies present in the sequence and obtain a prediction signal using previously coded content. Any information necessary to compute such prediction is encoded in the bit stream. The prediction signal is subtracted from the original signal to obtain a residual signal. This is then input to a transform module 14, which attempts to further reduce spatial redundancies within a block by using a more suitable
15 representation of the data. The resulting signal is then typically quantised by optional quantisation module 16, and finally the resulting data formed of the coefficients and the information necessary to compute the prediction for the current block is input to an entropy coding unit that 18 makes use of statistical redundancy to represent the signal in a compact form by means of short binary codes. Figures 2 and 3 illustrate in more detail example
20 prediction techniques that may be used in embodiments of the invention. Figure 4 illustrates by way of example a base matrix for an 8 x 8 transform unit.

The decoding process illustrated on the right hand side of Figure 1 is essentially the reverse process. A received bit stream is decoded in module 20 to obtain a block of coefficients and the information necessary to compute the prediction for the current block.
25 The block of coefficients is de-quantised in dequantisation module 22 and inverse transformed to the spatial domain by transform unit 24. A prediction signal is generated as before, from previously decoded samples from current or previous frames and using the information decoded from the bit stream, by prediction module 26. An original picture block is then reconstructed from the decoded residual signal and the calculated prediction block
30 in the reconstruction block 28.

The prediction module 12 of a video encoder provides the prediction signal for a given block. The way this prediction is computed depends on the current coding conditions, such as the temporal order of the current frame in the sequence or the coding configuration, and is generally based on rate-distortion decisions. There are a number of
35 different ways that the prediction signal may be generated, as is well known in the art.

Two prediction schemes are generally used in modern video coding standards. So called intra-prediction methods make use of content extracted from the same frame as the currently encoded block. Usually only information extracted from samples in the surroundings of the current block being encoded are used to compute the prediction. This kind of prediction is generally useful in the case of high spatial correlation within the current picture, or when the content in the frame is the only available information (e.g. while encoding the first frame in the sequence). Conversely, inter-prediction methods make use of the content of previously encoded frames (such as neighbouring frames in the sequence).

In general the content of these frames can be very similar to the content of the current frame. Most of the encoders available today perform inter-prediction by means of motion estimation and compensation: the current picture block is compared with candidate prediction blocks extracted at certain spatial displacements from a previously encoded reference frame; for each candidate prediction block a distortion value is computed to measure the similarity between the two blocks; finally the candidate at minimum distortion is extracted to form the prediction for the current block. The information needed by the decoder to select and extract the correct prediction is included in the bit stream transmitted from the encoder, referred to as motion information.

Once a prediction is computed for each picture block, this is subtracted from the original data to obtain a residual signal. This residual signal is input to the transform module 14 to obtain a more compact representation. In general this is achieved by transforming the signal to the frequency domain, using a discrete cosine transform (DCT) or other process, to condense most of the information in a small number of coefficients suitable for a compact representation. Quantisation is applied on transform coefficients which reduces their precision to achieve desired high compression. Finally the obtained residual coefficients are entropy coded in the bit stream by Entropy Coding Unit 16.

Conventional intra-prediction methods typically compute the prediction block using a number of spatial samples extracted from the boundaries of the currently encoded picture block. To ensure that the process can be repeated at the decoder side, only samples in blocks that have already been encoded can be used for this purpose. For instance an encoder that processes blocks in raster scan order can only use samples in the top-left portion of the frame.

Several methods can be used to obtain the prediction block. In particular three modes are usually defined at this purpose in recent standards such as AVC or HEVC: angular, DC or planar prediction.

When using angular prediction, samples are extrapolated from data in the boundary according to a certain specified angular direction. Different angular directions can be used to extrapolate samples at the boundary throughout the block. In the case of HEVC up to 33 directions (or modes) can be considered . Each sample is predicted using one or two samples in the boundary by means of linear interpolation, such as in the expression:

$$x = \frac{(32 - w)}{32} s_1 + \frac{w}{32} s_2$$

where s_1 and s_2 are the intensity values of the samples used for the prediction, and w is an integer from 0 to 32 which identifies the prediction direction.

The process of obtaining the angular prediction is illustrated in Figure 2. Suppose that the encoder is trying to compute the sample in the location denoted as X in the figure. For angular directions between the directions depicted by arrows 1 and 2, X is predicted using the samples located in A and C. The prediction at the direction depicted by arrow 1 (pure diagonal) happens by direct extrapolation of the sample located in C, or $w = 0$ in the above expression; similarly the prediction at the direction depicted by arrow 2 (pure vertical) happens by direct extrapolation of the sample located in A, or $w = 32$ in the above expression. Conversely in the case of the angular direction depicted by arrow 3, the samples located in A and B are used for the prediction; a weight $w = 8$ is used to obtain the predicted sample.

All angular intra-prediction modes available in HEVC are illustrated in Figure 3, using the same nomenclature as in the standard. The directions indicated in bold indicate the modes that are obtained by means of direct extrapolation (such as pure vertical, pure horizontal, or pure diagonal).

The DC intra-prediction mode (referred to as mode 1 in HEVC) consists in predicting all samples in the prediction block using the same value, usually obtained by averaging a number of samples extracted from the boundary. This mode is suitable to predict very smooth areas of static content in the frame, but fails in providing a good prediction on content presenting higher frequency components.

For this reason, planar intra-prediction mode can be used. It has already been introduced in the AVC standard and is also used in HEVC. Planar prediction is based on successive interpolations of samples. The version currently implemented in the HEVC standard (referred to as mode 0) can be summarised as follows, again with the help of Figure 2. First the sample in location Z (at the bottom-right position in the block) is

predicted as the arithmetic average of the sample in L (extracted from the boundary in the

top-right location), and in T (extracted from the boundary in the bottom-left location). Then the right-most column and bottom-most row are predicted using bilinear interpolation between samples in T and Z, or in L and Z respectively. All other samples are predicted by means of bilinear interpolation using a combination of previously predicted samples and boundary samples.

Due to the fact that a relatively large number of samples is predicted using a small amount of information strongly localized on a particular area in the frame, conventional intra-prediction methods might introduce unwanted prediction artefacts and in general might not provide sufficiently accurate predictions. In the case of angular prediction this is mostly evident when using modes with a strong directionality, such as pure vertical or pure horizontal modes. Consider for instance the case of a block predicted by means of pure horizontal mode: the entire data is predicted using exclusively the information in the samples in the column immediately to the left of the block. The original samples, particularly those in locations close to the right edge of the block, might be very different from these predicted samples, resulting in considerably high residuals localized in a particular area in the predicted block.

High residuals cannot be efficiently compressed and an attempt to reduce related bitrates typically results in blocking artifacts in the decoded frames. Other types of intra-prediction modes (such as DC prediction) also might produce similar artefacts.

For this reason most recent intra-prediction algorithms typically include a filtering algorithm usually applied to the samples in the boundary prior to performing intra-prediction. In particular HEVC makes use of a smoothing filter, consisting of a gradient-based smoothing (for vertical or horizontal prediction) or a two-tap filtering (for DC prediction).

When the smoothing filter is disabled, errors in the residual signal tend to be distributed towards the right edge of the block, while a more uniform distribution of the error is obtained when the smoothing filter is enabled. While the smoothing filters help in distributing the residual error throughout the block, it does not decrease the energy of such residuals. In fact relatively high residuals can be expected as a result of conventional intra-prediction methods even when using such filters. This is especially problematic when targeting high quality applications where such information cannot be discarded but needs to be compressed and encoded.

With reference to Figure 4, the concept of transform encoding will now be discussed. The main idea behind transform coding is that of compacting the information in the signal in a smaller number of parameters by transforming it to the frequency domain.

Ideally the transform should be able to decorrelate the signal as much as possible, obtaining a representation that is optimal for the purpose of data compression.

A common way of obtaining such representation for image and video coding methods is by means of the two-dimensional DCT. When using DCT, each block of $N \times N$ pixel intensities is expressed as a weighted sum of N^2 scaled cosinusoids of different amplitudes and frequencies, referred to as the DCT basis functions. The weights that must be applied to each basis function in order to return exactly the original signal are referred to as the transformed coefficients.

Consider a certain $N \times N$ block of samples X and referring to each sample as $x(i, j)$.
10 The DCT-II is defined as:

$$x(i, j) = \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} k_{m,n} X(m, n) \cos \frac{\pi(2i+1)m}{2N} \cos \frac{\pi(2j+1)n}{2N}$$

[1]

15 where the $k_{m,n}$ are predetermined scaling factors and the $X(m, n)$ are the transformed coefficients. When $m = 0$ and $n = 0$ the frequency of the two cosines in the function becomes zero and the corresponding basis function reduces to a constant; such frequency component is referred to as the DC component, and the associated weight $X(0, 0)$ at the top-right location in the transformed block is referred to as the DC coefficient. Each
20 following value of m and n corresponds to a higher frequency component $AC(m, n)$, as illustrated in Figure 4. Each transformed coefficient $X(m, n)$ can be determined in terms of an inverse expression as a function of the original elements $x(i, j)$. As a result, each block of $N \times N$ samples in the spatial domain is exactly represented by its corresponding block of $N \times N$ coefficients in the frequency domain.

25 Both computations of the forward DCT transform (to obtain the transformed coefficients) and inverse DCT transform (to compute the original samples) can be compacted in the form of simple matrix multiplications following from the definition of an appropriate $N \times N$ matrix Q , referred to as the transform base matrix. The elements of Q are then defined as:

30
$$q(m, n) = k_{m,n} \cos \frac{\pi(2m+1)n}{2N}$$

[2]

Due to the orthogonality of Q its inverse is the same as its transpose, and the forward and inverse DCT transforms can be expressed as:

$$\bar{\mathbf{X}} = \mathbf{Q}\mathbf{X}\mathbf{Q}^T$$

[3]

and

5

$$\mathbf{X} = \mathbf{Q}^T \bar{\mathbf{X}} \mathbf{Q}$$

[4]

In practical applications, due to the limited availability of resources, the elements in the base matrices used to compute the transform are obtained by means of rounded integer approximations of the values obtained using Equation 2 above. Also, the transform coding is most effective in blocks that contain a relatively small amount of texture changes. In these cases the majority of coefficients at higher frequencies would be equal or close to zero while most of the signal would be compacted in a few coefficients at lower frequencies. It is therefore crucial that a frame is partitioned into blocks specifically for transform coding in the most efficient possible way. In the case of HEVC, a recursive approach is used for this purpose in which square blocks of different sizes are considered, referred to as transform units (TU).

A single base matrix \mathbf{Q}_{32} is defined for the largest allowed TU size (set at 32 x 32 pixels), obtained by appropriately rounding to integer approximations the values in Equation. 2. The base matrices for smaller TU sizes are obtained by downsampling this largest base matrix. The matrix \mathbf{Q}_8 used for 8 x 8 TUs is illustrated in Figure 4.

The DCT has many desirable characteristics (such as ease of computation and very good performances for inter-predicted blocks), but it is not the optimal choice to decorrelate the residual signal especially in the case of intra-predicted blocks. This is due to the fact that its cosinusoidal basis functions are bad approximations of the behavior of the residual signal. Consider for instance the case of angular prediction. Samples in locations closer to the top or left boundaries are likely to be predicted more accurately than samples close to the bottom right corner of the block. Consequently, the residuals are likely to assume smaller values along the top row and right column in the block, and progressively larger values towards the left and bottom edges of the block.

A better decorrelation of the residual signal in intrapredicted blocks can be obtained by means of a discrete sine transform (DST) with appropriate frequency and phase components. DST was implemented already in selected intraprediction modes in AVC and is also used in HEVC for small 4 x 4 intra-predicted TUs. It is not used for larger blocks due to its generally higher computational complexity and the lack of fast algorithms for computing the transformed coefficients.

A detailed study on the optimality of different transforms in the case of HEVC angular intra-prediction has found that there exists a strong correlation between the optimal transform to be used (in either the horizontal or vertical steps), and the angular direction of the prediction. Similarly, another study on the optimality of the DCT transform making use of Gaussian Markov random field models concluded that DCT is indeed not optimal for
 5 intra-predicted residual signals.

Transform is followed by quantization. Each coefficient is quantized to a given step (depending on a parameter usually referred to as the quantization parameter, QP). The higher the QP, the coarser is the quantization. Coefficients that are closer to zero, usually
 10 corresponding to the higher frequency components, are completely discarded.

The above discussion illustrates the processes of partitioning, prediction and transforming blocks of picture data. Figure 5 illustrates how the processes are put together in a more mathematical sense in a known encoder. Consider that a certain $N \times N$ square block of original samples X is being encoded. Consider also that an equally sized block of
 15 samples P is being considered as a predictor or prediction block for X , obtained from one of the possible intra-prediction modes. We denote as Q the $N \times N$ transform base matrix. Typically in video coding standards this corresponds to an appropriately sized N -point DCT or DST transform matrix. In known video codecs the residual samples are then computed in the spatial domain as $R = X - P$, and are then transformed to the frequency domain as
 20 $\bar{R} = QRQ^T$. These steps are illustrated in the scheme in Figure 5.

Similarly at the decoder side in known video codecs, the coded coefficients C that are extracted from the bitstream are dequantised to give a residual signal \bar{R} , which is then inverse transformed into the spatial domain into the signal R , before being added to the prediction block P . This process is illustrated in Figure 6.

In one example embodiment of the invention, we have appreciated that it would be advantageous to also process the prediction block and the original block in the frequency domain, and in addition the residual signal,. At the encoder side the predictors and original signals are directly transformed to the frequency domain before the residual is calculated. The residual signal is then computed in the frequency domain as the difference between
 30 the transformed original and prediction signals. This is illustrated in Figure 7.

At the decoder on the other hand, in one example embodiment of the invention, the prediction samples are transformed to the frequency domain, added to the dequantised coefficients, and finally the reconstructed samples are inverse transformed as in Figure 8.

Operating in the frequency domain allows the example embodiment of the invention to more easily manipulate the frequency coefficients prior to encoding in a way that provides advantages over the prior art. This will be discussed in more detail below.

Otherwise, the frequency domain operation can be understood to be largely
 5 equivalent to the spatial domain operation of the prior art. If the same X and P are used as input to the two encoding schemes of Figures 5 and 7, the same residual signal should be obtained in the frequency domain. In practice, it is not always identical due to limits on computing resources. The linearity of the transform however can be easily demonstrated by noticing that transforming the original picture block X as in the scheme in (b) results in a
 10 block of coefficients $\bar{\mathbf{X}} = \mathbf{Q}\mathbf{X}\mathbf{Q}^T$. Equivalently the prediction P can be transformed in $\bar{\mathbf{P}}$. Finally the residuals are directly computed in the frequency domain as:

$$\begin{aligned} \bar{\mathbf{R}} &= \bar{\mathbf{X}} - \bar{\mathbf{P}} = \mathbf{Q}\mathbf{X}\mathbf{Q}^T - \mathbf{Q}\mathbf{P}\mathbf{Q}^T \\ &= \mathbf{Q}(\mathbf{X} - \mathbf{P})\mathbf{Q}^T = \mathbf{Q}\mathbf{R}\mathbf{Q}^T \end{aligned}$$

15

demonstrating the equivalence.

In practical applications, due to the integer approximations and limitations on the coefficient buffer size, the transforms are not linear. Variables during the stages of
 20 transform computation might be truncated to limited precisions. For this reason the two schemes of Figures 5 and 7 (ignoring the process block) can result in slightly different transformed residuals. In the case of HEVC, and denoting with the letter B the input/output bitdepth of the representation, and with N the width or height of the current blocks, the first (horizontal) stage of DCT transform in Figure 7 may be followed by a binary shift of:

$$S_{\text{hor}} = \log_2(N) - 1 + (B - 8)$$

25

The shifts needed in HEVC in the case of DCT for 8-bit input data are shown in Table I, for the two stages of transform.

30 The example embodiment of the invention seeks to provide a more efficient prediction block for the encoding process as will be described below. In general by providing a more accurate prediction of the current block, a better encoder performance can be expected (due to the smaller residual samples which require less bits to be coded). While common distortion metrics in the spatial domain such as the sum of squared
 35 differences (SSD) can be used to estimate the accuracy of a prediction, these types of

metrics can fail in measuring the impact of prediction methods on the residual coefficients at different frequency components. It is instead reasonable to expect particular effects of certain prediction modes on specific frequency components of the residual signal. These effects might be captured and analysed to formulate appropriate processing methods to improve the coding efficiency.

An effective measure of the similarity between prediction and original blocks in the frequency domain can be obtained by means of the per-coefficient correlation. This is the normalized cross-correlation between the time series of prediction coefficients and corresponding original coefficients at each specific location in the block. To estimate these correlations, the approach was implemented in the context of HEVC intraprediction and a few sequences were encoded to collect test data.

The normalized cross-correlation can be defined as:

$$R_{N,s}[\bar{p}[m,n], \bar{x}[m,n]] = R_{N,s}[m,n] = \frac{\left[\sum_{i=0}^{K_{N,s}-1} (\bar{p}_i[m,n] - E\{\bar{p}[m,n]\})(\bar{x}_i[m,n] - E\{\bar{x}[m,n]\}) \right]}{K_{N,s} \sigma_{\bar{p}[m,n]} \sigma_{\bar{x}[m,n]}}$$

15

where the expected values $E\{\cdot\}$ and standard deviations σ are estimated from the samples. $K_{N,s}$ denotes the number of $N \times N$ blocks tested using the intra-prediction mode s , the two time series of prediction and original coefficients are referred to as $\bar{p}[m,n]$ and $\bar{x}[m,n]$ respectively; namely the elements $\bar{p}_i[m,n]$ for $i = 0, \dots, (K_{N,s} - 1)$ are the prediction coefficients extracted at location $[m, n]$ from all the $N \times N$ blocks predicted using the intra-prediction mode s , and similarly $\bar{x}_i[m,n]$ are the original coefficients at the same location.

20

Values of $R_{N,s}[m,n]$ close to +1 indicate that the intra-prediction mode s is good at predicting the coefficient located in $[m, n]$ when the TU size is $N \times N$. Values of the cross-correlation close to zero indicate instead that the predicted samples in $[m, n]$ carry almost no information on the original samples.

25

The correlation results for TUs of different sizes (4×4 , 8×8 and 16×16) are shown in Tables 2 for the planar intra-prediction mode; We have found that the TU size has a strong impact on the correlation values especially at higher frequencies (i.e. towards the bottom-right corner of the blocks). Relatively high correlation values are reported in 4×4 blocks at all locations (minimum correlation of 0.3), whereas very low correlation values were found in larger block sizes, showing that the prediction coefficients at these locations carry almost no information about the original coefficients. Similarly the correlation values

30

are strongly influenced by the intra-prediction mode, as illustrated in Table 3. The correlation values for three angular modes are shown in the case of 8 x 8 blocks. Notice in particular that pure horizontal angular prediction (corresponding to $s = 10$) results in high correlations in the left region of the block, and very low correlation values elsewhere; similarly pure vertical angular prediction ($s = 26$) results in high correlations in the top region, and very low values elsewhere.

Conventional video coding architectures such as those illustrated in Figure 5 and 6 do not take into consideration these effects. By computing independently the transform of the prediction and original blocks, however, the approach taken in the example embodiment of the invention allows instead the encoder to appropriately filter and replace selected frequency components of the prediction, as will now be illustrated.

The correlation analysis performed using the HEVC codec and described above shows that in many cases the prediction coefficients obtained by conventional intra-prediction methods carry very little information on the original coefficients.

This seems to be less evident for small transform sizes (such as 4 x 4 TUs in HEVC), which generally resulted in higher correlation values less influenced by the prediction mode. The values of the correlation reported for blocks equal or larger than 8 x 8 show instead a clear relationship with the intra-prediction mode being used, especially in the case of angular direction where they closely follow the angle of prediction. In the frequency domain this results in high correlation values between the prediction and original coefficients in the first column, in slightly lower values in the next column, and very low values elsewhere.

In the example embodiment of the invention, therefore the prediction coefficients in the locations with very low correlation with the original signal are discarded and replaced with more informative content generated by applying a number of value substitution processes. Some high frequency components can therefore be removed from the residual signal providing higher compression efficiency.

The process of selecting particular coefficients in the transformed block is easily formalised through the definition of a set of masking matrices, referred to as masks or patterns. Each mask is a matrix of binary elements that can be applied to a block of coefficients; the value of a binary element in a certain location determines whether the corresponding coefficient in the block is preserved or discarded. To illustrate and validate the method the approach was implemented again in the context of the HEVC codec, but the method can be implemented in any video codec making use of block-based hybrid scheme.

We refer again to each element in the transformed prediction block as $\bar{p}[m, n]$, and assume a TU size of $N \times N$. We then assume that the coefficients that are discarded are replaced with known values $r[m, n]$ (the process of obtaining such values will be discussed below). Lastly, we define also a mask size L , where L is an integer such that $0 \leq L \leq N$.

5 Although more complex options are possible, only four classes of pattern will be described here.

1) Vertical rectangular patterns, referred to as vr , consisting of L consecutive rows of preserved coefficients at the top of the pattern.

10 2) Horizontal rectangular patterns, referred to as hr , consisting of L consecutive columns of preserved coefficients in the left region of the pattern.

3) square patterns, referred to as sq , consisting of $L \times L$ preserved coefficients in the top-left corner of the pattern.

4) Triangular patterns, referred to as tr and consisting of a region of preserved coefficients at the top-left of the pattern.

15 Three values of L are considered for illustration for the classes vr , hr and sq , specifically $L = N/4$, $L = N/2$ and $L = 3N/4$ although many values are possible. These are illustrated in Figure 9, as masks a, b, and c. The size of the mask within the block can be varied as appropriate. Also, it will be appreciated that other shapes may be possible, depending on the coding circumstances. The masks shown in Figure 9 should therefore be
20 taken as illustrations only. In Figure 9, the shaded areas correspond to a first region of the prediction block in which the prediction coefficients are left unchanged. The white, or unshaded areas, correspond to a second region of the prediction block in which the prediction coefficients are substituted or replaced.

We have found that particular masks or patterns work consistently well with
25 particular prediction modes, and prediction mode angular directions. In order to determine which patterns produce the best results for a respective intra-prediction mode, the method was first implemented assuming that all the elements in the prediction block that are discarded are replaced with zero-valued coefficients, or, $r[m, n] = 0$ in expressions 1 to 4 above. Following from the previous results the method is used on TUs larger or equal than
30 8×8 , with 4×4 TUs being conventionally coded. Each prediction block in the transform domain is processed using a certain pattern. The processed transformed prediction block is then compared with the transformed original block. In particular, a distortion measure is computed, such as the sum of squared differences (SSD), or the sum of absolute differences (SAD). Other techniques may be used as will be known to those skilled in the
35 art. This is repeated for all available patterns. The pattern that results in the lowest distortion is selected as the optimal pattern and used to process the current TU. In a first

embodiment, once the encoder determines the optimal mask or pattern for a particular size of transform unit, the encoder continues to select this pattern when ever a transform unit of that size is encountered. This technique results in an improvement in encoding efficiency, and means that there is no need to continuously signal to the decoder which masks or patterns are used. Either the encoder can signal the preferred masks or patterns to the decoder once it has performed its initial analysis, and the decoder can store the required information for prediction block processing, or the decoder can also perform the same analysis as the encoder and store the results for future use.

For example, the method was tested on a few sequences (encoded at very high quality, QP = 5), with the results illustrated in Table 4. The table reports the most frequently selected patterns, according to the TU size and intra-prediction mode being used. For most cases, patterns with a very high number of discarded coefficients (for instance with $L = N/4$) were selected, especially for large 32 x 32 TUs. Also, as expected, patterns in the hr class were mostly selected in angular horizontal modes, and conversely patterns in the vr class were mostly selected in vertical angular modes. The square pattern with $L = N/4$ was often selected regardless of the directionality of the prediction. The triangular pattern was mostly chosen when using planar mode, and rarely selected otherwise.

While the previous results are useful to determine a relationship between the characteristics of a certain TU and the best pattern that can be used to process the prediction coefficients, zeroing-out the prediction coefficients is not expected to be an optimal choice in terms of compression efficiency of the proposed method. Although, it may still be used to provide improvements in performance. In the example embodiment of the invention, selected frequency components of the prediction signal are therefore replaced with coefficients that are more correlated with the corresponding components in the original signal.

Depending on the type of data in the original blocks, the original signal can contain large amount of details or textured areas, resulting in many non-zero components at high frequencies. On the other hand typical conventional intra-prediction modes result in many zero-valued components at high frequencies (for example, the DC prediction mode returns a single non-zero coefficient, and similarly pure horizontal or vertical angular predictions return a single column or row of non-zero valued coefficients), and consequently high frequency components are often left in the residual signal. These sparse non-zero residual coefficients at high frequencies are extremely expensive to encode using conventional entropy coding methods. When targeting medium quality applications such coefficients are quantised and discarded, but in the case of high quality video coding this is not allowed.

Higher coding efficiency is therefore achieved by removing these coefficients from the residual signal prior to quantization.

5 A very good prediction can be obtained in the transformed domain using the best mask or pattern for each particular block, and replacing the coefficients that are discarded with values that are as close as possible to those in the same locations in the original transformed block. The use of a mask or pattern and/or the appropriate substitution of coefficients in a region of the prediction block not covered by the mask will be generally termed a modification process. The encoder according to the example embodiment of the invention may be implemented to test all possible candidates of modification process
10 (varying both masks and coefficient substitution) for a particular transformed prediction block and select the optimal candidate. In practice, however, allowing both the choice of the pattern and the values to be used in place of the discarded coefficients can lead to delays in encoding, and might not result in efficient compression due to the large amount of information that would need to be processed and transmitted for each block.

15 For this reason, in a first example embodiment of the invention, the pattern to use on a block is not transmitted, but is instead fixed depending on characteristics of the block such as size and intra-prediction mode. The encoder and the decoder may then simply look up the appropriate mask to use according to the coding circumstances using identical look up tables.

20 The choice of values to use for $r[m, n]$ in expressions 4 to 7 is therefore important and in a more complex embodiment of the invention can be optimised for a given mask or pattern. Ideally, the choice of values to use in the coefficient substitution process should require as few bits as possible to be signalled in the bitstream, and at the same time they should be able to predict the high frequency components in the original block in order to
25 remove such components from the residual signal. We have found that both of these requirements can be satisfied replacing selected components with a constant single values.

In a second example embodiment of the invention these values are extracted from a look-up table that has been generated to contain suitable candidate values. By keeping the number of elements in the look-up tables as low as possible, very few bits are needed
30 to extract the correct element for each block. At the same time testing several candidates for the values of $r[m, n]$ allow the encoder to flexibly choose the candidate that better predicts the original signal in each particular case. While the proposed method can be used in any video codec making use of the block-based hybrid approach, testing and validation of the approach was obtained again implementing the method on top of conventional
35 HEVC. In particular, the patterns in Table 4 (i.e. the most frequently selected patterns obtained when zeroing-out components) are used on a TU depending on its size and intra-

prediction mode. The elements in the look-up tables were similarly derived following from extensive empirical analysis.

In a simple example, a look up table was generated with all possible candidate values that could be employed, for example, positive and negative integers ranging from 0 to 64. These look-up tables were then tested against actual TUs and masks and filtered to remove candidate values that were rarely found to give the optimal coding. In most cases, a look up table containing between 8 and 16 values was found to be more than sufficient to provide suitable candidate values. An example look up table is included as Figure 10 for illustration. It will be appreciated that the look up table is merely one way of retaining for a given prediction mode and TU size a set of candidate values for coefficient substitution. Other ways are performing coefficient substitution are possible. For instance, instead of using a fixed set of constant values, in an alternative embodiment, we might formulate and compute a fixed set of matrices of values. These matrices would be formed of pre-computed, coefficients which replace the existing coefficients in the prediction signal. Consider for instance the case in which we are coding a 8 x 8 block of samples, and using a pattern such as horizontal rectangular with size half. The area that we modified is formed of 4x8 samples. In this case we might produce a set of pre-computed 4x8 matrices. Each would be associated to an index as in a look-up table. The encoder would test all the available matrices before selecting the optimal one. The corresponding index would be transmitted to the decoder. The set of matrices should be made available at both the encoder and decoder side.

Other more complex substitutions can be formulated. In Figure 10, a default substitution of a zero value may be used, requiring no additional information to be transmitted in the bit stream.

Different look-up tables are therefore derived and used depending on TU size and intra-prediction mode. The number of elements in the tables depends on the TU size, with fewer values allowed for tables used in case of 8 x 8 TUs, and increasingly more values in tables used for larger transform sizes. Once derived, such tables are then made available at both the encoder and decoder side.

Although in the first example embodiment of the invention, the encoder does not need to transmit information to the decoder specifying which mask or patterns and/or which substitution technique for the prediction coefficients are used, in more complex embodiments the encoder may calculate for every individual transform unit partitioned from an original block, the optimal mask or pattern, and the optimal coefficient values for substitution. In this case, it is useful if the encoder calculates the cost of transmitting the

necessary information to the decoder to indicate which pattern or mask, and/or which coefficient substitution technique was selected, and factor this into its selection decision.

For example, any information transmitted to the decoder to indicate the selected compression scheme will count as extra transmission overhead and so will impact the amount of picture data that can be transmitted and/or the quality of the transmitted data. As noted above, this means that it is preferable to carefully manage the number of variables involved in the modification process and where possible keep these to an optimal, minimal number. In more complex embodiments of the invention, it is therefore desirable to factor in the cost of transmitting the data describing the modification process when calculating which mask or pattern, and/or coefficient substitution process is to be used. In this regard, rather than just selecting the best pattern or mask based on the distortion error, such as the sum of square differences (SSD), the sum of absolute differences (SAD), or the normalised cross-correlation, the encoder may use a rate distortion optimisation technique to determine which modification process and signalling scheme to indicate the selected modification process is optimal.

Rate distortion techniques combine a distortion metric, such as SSD, SAD or normalised cross-correlation, with an estimated bitrate necessary to signal to the decoder the required information to describe the modification process. Estimates can be produced by way of known fast estimation methods which return a figure for the number of bits needed to encode the index associated with the current pattern, or alternately, the encoder could in fact perform the actual coding method it would use to transmit the picture information and signalling information to the decoder, and then feedback the exact number of bits required to the decision process. The optimisation is then carried out using a Lagrangian multiplier and approximated Lagrangian optimisation methods for the expression:

25

$$\text{Rate Distortion Cost (RD Cost)} = \text{Distortion} + \lambda (\text{Estimated Bit rate})$$

The encoder would compute the RD cost for each modification process extracted from the list of considered patterns, and select the pattern at minimum cost.

Thus in an encoder according to a further example embodiment of the invention, each TU is encoded using the proposed method and testing all of the elements in the appropriate look-up table. The same TU is also encoded using conventional HEVC, and finally the best solution in a RD sense is selected and signalled in the bitstream. The algorithm can be summarised as follows.

1) Conventional HEVC coding is performed: the residuals are computed in the spatial domain, transformed and quantised, and entropy coded; the reconstruction block is

computed using the inverse process and compared with the original block to obtain a distortion. The RD cost is computed for this solution and used as the current minimum cost. conventional HEVC is considered as the temporary optimal solution.

2) Prediction and original blocks are independently transformed (using adjusted
5 binary shifts). A specific pattern and look-up table are considered according to current TU size and intra-prediction mode. The first element in the look-up table is extracted.

3) The prediction block is processed using the selected pattern and current element
10 in the look-up table. The residual signal is computed in the frequency domain, and the coefficients are successively quantised and entropy coded. An index to signal the current element in the look-up table is also entropy coded in the bitstream

The reconstruction block is computed using the inverse process and compared with the original block to obtain a distortion. The RD cost is computed for this solution.

4) The RD cost is compared with the current minimum cost. If the RD cost is lower
15 than the current minimum cost, this becomes the new minimum, and correspondingly the temporary optimal solution is updated to the current solution. If there are elements left in the lookup table, the next element is extracted and step 3 is repeated.

5) Otherwise if there are no other elements in the look-up table the algorithm
outputs the optimal solution for the current TU and exits.

A novel approach for efficiently coding video sequences in high quality conditions
20 has therefore been presented based on frequency domain prediction methods. Conventional intra-prediction techniques often fail in providing good prediction of the original signal, especially at higher frequencies and when using large transform sizes. For high quality applications it is crucial that these high frequency components are not quantised and discarded, and for this reason conventional video coding methods result in
25 very high bitrates. The proposed method allows instead to adaptively replace the high frequency components in the prediction signal, possibly resulting in smaller residual coefficients without any loss of data.

In order to do so, the approach is based on a modified encoder scheme where
30 original and prediction signal are independently transformed to the frequency domain; this is opposite to conventional schemes where the residual block is computed in the spatial domain and successively transformed. The proposed encoder scheme allows for an additional stage of processing, introduced at the encoder side after the prediction is transformed, and prior to the residual calculation. The processing is based on a set of
35 masking patterns applied to the transformed prediction block. Each pattern identifies which coefficients in the prediction signal are discarded, and which coefficients are instead preserved. Coefficients that are discarded are replaced with constant values extracted from

lookup tables, appropriately derived depending on the transform size and intra-prediction mode.

The approach is shown achieving consistent gains against conventional HEVC under high quality conditions. Up to -4.6% BD-rate reductions are achieved in the all-intra profile, with up to -4.3% reductions achieved in the low delay profile.

Moreover, while the approach is considerably more complex than conventional methods at the encoder side, it has a very little impact on the decoding complexity.

Although the above description assumes that the prediction block and the original block are transformed to the frequency domain before modifying the prediction block and calculating the residual signal, it is also possible to apply the technique of simplifying the prediction block, using a mask or pattern, in the spatial domain.

For example, referring now to Figure 11, a scheme in which the processing of the prediction block is carried out in the spatial domain will now be described. Figure 11 illustrates the process of calculating the residual signal, transforming this, and quantising this for transmission in a bit stream. The arrangement is identical to the prior art situation of Figure 5, except that the prediction block P is processed in the spatial domain before being subtracted from the original block. As with the frequency domain processing, in the processing block, a number of possible patterns or masks are applied to the prediction block. Example mask or patterns for the processing step are illustrated in Figure 12 as masks, a, b, c and d. The size of the mask within the block can be varied as appropriate. Also, it will be appreciated that other shapes may be possible, depending on the coding conditions. The masks shown in Figure 12 should therefore be taken as illustrations only. In Figure 12, the shaded areas correspond to a first region of the prediction block in which the prediction coefficients are left unchanged. The white, or unshaded areas, correspond to a second region of the prediction block in which the prediction coefficients are substituted or replaced. In this example, the coefficients to be substituted are spatial domain samples corresponding to pixel intensities.

As with the frequency domain case, the modification of the samples in the prediction block can be carried out in a number of different ways, all with the purpose of ensuring that the coding of the resulting residual signal is carried out more efficiently. For example, one way is by using a substitution mechanism similar to the one used in the frequency domain and illustrated in Figure 10. For example, look-up tables might be used which replace the selected samples in the prediction block with synthetic samples, and associate each of these elements with an index that can be signalled to the decoder side. In this case, the samples in the prediction block are representative of pixel intensities, and the term synthetic samples is used to suggest that the values substituted are scaled to

match values in the range of pixel intensities in the prediction block, but are otherwise selected by the modification process. In the same way as for Figure 10, an example of such a table is illustrated in Figure 13. The decoding process for the spatial domain is shown in Figure 14, and is similar to that of Figure 6 for the frequency domain.

5 The process of selecting either the pattern and/or the elements in the look-up table would be identical to the one in the frequency domain. An index to identify the pattern and/or an index to identify the element in the look-up table might need to be encoded in the bitstream. Look up tables are not themselves essential but merely provide a convenient way that substitution values may be stored and referred to.

10 In both of the techniques described above, the coefficients in the second region of the mask that are substituted or replaced, may be substituted or replaced with values that are a combination of the initial prediction coefficient and the predetermined substitution value. For example a weighted combination of the initial prediction value and the predetermined substitution value could be used.

15 The processing of the prediction block carried out in both the frequency and the spatial domains is generally a separate step carried out after the prediction block has been calculated in the normal way by the encoder or the decoder. In this sense, the prior art encoder and decoder schemes of Figures 1 and 2 can be thought of as having a modification processing block situated after the prediction module that receives the formed or generated prediction block and applies the modification process accordingly. This is in fact illustrated in Figures 7, 8, 11 and 14. Alternatively, the prediction module itself may be modified to generate a prediction block having the necessary modifications made to improve the efficiency of the encoding. The principle would be the same as described above, with generation of a prediction block having two regions, a first region in which the coefficients are standard coefficients generated from known prediction modes and methods, and a second region in which the coefficients are generated by an algorithm taking into account a value substitution, either substitution in the case of replacement with predetermined values or a combination of values. In embodiments of the invention, the first and second set of predetermined locations are exclusive of one another. In other
20
25
30
embodiments, the regions may overlap, such that the substituted values are combinations of predicted values and predetermined values for substitution.

35 It will be appreciated that the logical blocks or modules of Figures 1 and 2, 7, 8, 11 and 14 can be implemented in software and/or in hardware. Furthermore, although there are illustrated as separate blocks, the logical separate functions of the blocks or modules may be combined in single software or hardware modules as appropriate.

The above description is intended to be illustrative in nature, and not to limit the scope of the invention defined by the claims. To the extent that features of the invention are described with respect to separate example embodiments, it will be appreciated that these are contemplated as being combined with the features of other embodiments.

5

Table 1: Binary Shifts in Conventional HEVC for 8 bit data representation

TU Size	Bin. Shift
First DCT stage	
4 x 4	$\gg 1$
8 x 8	$\gg 2$
16 x 16	$\gg 3$
32 x 32	$\gg 4$
Second DCT stage	
4 x 4	$\gg 8$
8 x 8	$\gg 9$
16 x 16	$\gg 10$
32 x 32	$\gg 11$

10

Table 2: Selected values of the cross-correlation for blocks of different sizes intra-predicted user planar mode. For 8 x 8 blocks only the DC, AC(0,2) AC(0,4), AC(0,6) AC(2,0) ... components are shown. Similarly, for 16 x 16 blocks only the DC, AC (0,4), AC(0,8), AC(0,12), AC(4,0) ... components are shown.

5

4 x 4	0.956	0.899	0.790	0.675
	0.905	0.807	0.670	0.540
	0.827	0.697	0.527	0.427
	0.719	0.621	0.448	0.368
8 x 8	0.960	0.502	0.363	0.300
	0.455	0.011	0.004	0.002
	0.353	0.004	0.003	0.009
	0.198	0.011	0.003	0.000
16 x 16	0.949	0.374	0.280	0.318
	0.415	0.002	0.010	0.007
	0.177	0.002	0.017	0.008
	0.097	0.004	0.001	0.000

Table 3: Selected values of the cross-correlation for 8 x 8 blocks predicted using different intra-prediction modes. For 8 x 8 blocks only the DC, AC(0,2) AC(0,4), AC(0,6) AC(2,0) components are shown. Similarly, for 16 x 16 blocks only the DC, AC(0,4), AC(0,8), AC(0,12), AC(4,0) ... components are shown.

5

Mode 7	<u>0.963</u>	0.037	0.026	0.020
	<u>0.681</u>	0.153	0.109	0.081
	<u>0.503</u>	0.202	0.105	0.092
	0.088	0.207	0.072	0.005
Mode 10	<u>0.969</u>	0.351	0.234	0.129
	<u>0.939</u>	0.301	0.252	0.182
	<u>0.860</u>	0.201	0.187	0.183
	<u>0.856</u>	0.182	0.192	0.200
Mode 26	<u>0.952</u>	<u>0.755</u>	<u>0.514</u>	<u>0.571</u>
	0.372	0.341	0.225	0.199
	0.342	0.278	0.169	0.210
	0.348	0.219	0.178	0.153

Table 4: Patterns At Minimum Distortion According To Intra-Prediction Mode And TU Size.

Mode	8 x 8	16 x 16	32 x 32
0	tr	tr	tr
1	sq, L = N/4	sq, L = N/4	sq, L = N/4
2	sq, L = N/4	hr, L = N/4	hr, L = N/4
3	sq, L = N/4	hr, L = N/4	hr, L = N/4
4	sq, L = N/4	hr, L = N/4	hr, L = N/4
5	hr, L = N/4	hr, L = N/4	hr, L = N/4
6	hr, L = N/4	hr, L = N/4	hr, L = N/4
7	hr, L = N/4	hr, L = N/4	hr, L = N/4
8	hr, L = N/2	hr, L = N/4	hr, L = N/4
9	hr, L = N/2	hr, L = N/2	hr, L = N/4
10	vr, L = N/4	vr, L = N/4	vr, L = N/4
11	vr, L = N/4	tr	vr, L = N/4
12	hr, L = N/2	hr, L = N/4	hr, L = N/4
13	hr, L = N/2	hr, L = N/4	hr, L = N/4
14	hr, L = N/2	hr, L = N/4	hr, L = N/4
15	sq, L = N/4	hr, L = N/4	hr, L = N/4
16	sq, L = N/4	hr, L = N/4	hr, L = N/4
17	sq, L = N/4	sq, L = N/4	hr, L = N/4
18	sq, L = N/4	sq, L = N/4	sq, L = N/4
19	sq, L = N/4	sq, L = N/4	sq, L = N/4
20	sq, L = N/4	sq, L = N/4	sq, L = N/4
21	sq, L = N/4	vr, L = N/4	vr, L = N/4
22	vr, L = N/4	vr, L = N/4	vr, L = N/4
23	vr, L = N/4	vr, L = N/4	vr, L = N/4
24	vr, L = N/2	vr, L = N/4	vr, L = N/4
25	vr, L = N/2	vr, L = N/4	vr, L = N/4
26	vr, L = N/4	vr, L = N/4	vr, L = N/4
27	vr, L = N/4	sq, L = N/4	tr
28	vr, L = N/2	vr, L = N/2	vr, L = N/4
29	vr, L = N/2	vr, L = N/2	vr, L = N/4
30	vr, L = N/2	vr, L = N/2	vr, L = N/4
31	sq, L = N/4	sq, L = N/4	sq, L = N/4
32	sq, L = N/4	vr, L = N/4	vr, L = N/4
33	sq, L = N/4	vr, L = N/4	sq, L = N/4

CLAIMS

1. A method of encoding video data with compression:
 - 5 a) receiving an original picture block containing coefficients representative of a picture region;
 - b) receiving a prediction block containing coefficients representative of a predicted picture region;
 - 10 c) forming a modified prediction block by substituting the values of coefficients in the received prediction block for other predetermined values according to a modification process, the predetermined modification process including applying a mask to the received prediction block, the mask defining at least a first region of the prediction block where the values of the coefficients are preserved and at least a second region of the prediction block where the values of coefficients are substituted for other predetermined values; and
 - 15 d) calculating a residual signal for transmission, by subtracting the original picture block and the modified prediction block from each other.

2. The method of claim 1 comprising:
 - 20 transforming the original picture block and the prediction block into the frequency domain; and
 - forming the modified prediction block and calculating the residual signal in the frequency domain.

3. The method of claim 1 or 2, wherein the modification process is selected from a
25 plurality of predetermined modification processes.

4. The method of claim 3, wherein the modification process selected from the plurality of predetermined modification processes depends on one or more of the shape and/or size of the prediction block, and the prediction encoding scheme used to generate the prediction
30 block.

5. The method of claim 2, 3 or 4, wherein the predetermined plurality of modification processes include a predetermined plurality of differently shaped masks for application to the prediction block and/or a predetermined plurality of different processes for determining
35 said values to be substituted.

6. The method of claim 5, wherein the predetermined plurality of processes for determining said values to be substituted include one or more of:

i) replacing the value of all coefficients in the at least a second region with zero;

5 ii) replacing the value of all coefficients in the at least a second region with a predetermined non-zero value;

iii) replacing the value of all coefficients in the at least a second region with a value calculated according to a predetermined algorithm.

7. The method of claim 3, comprising:

10 e) selecting a candidate modification process from the plurality of predetermined modification processes;

f) forming a modified prediction block according to the candidate modification process;

15 g) calculating the prediction error between the modified prediction block and the original picture block;

h) repeating steps e) to g) for each modification process in the plurality of predetermined modification processes and for each modification process calculating a prediction error between the modified prediction block and the original picture block;

20 i) selecting a modification process to use for forming the modified prediction block for use in the calculation of the residual signal from the plurality of predetermined modification processes;

j) quantising and encoding the residual signal for transmission in a bitstream, wherein the selection of the modification process to use for calculating the residual signal is based on the prediction error.

25

8. The method of claim 7, wherein the predetermined plurality of processes for determining said values to be substituted includes consulting a look-up table containing a plurality of candidate values, wherein different look-up tables are consulted depending on the candidate modification process, and wherein the candidate values in the look table
30 have been selected based on previous analysis of values that have provided favourable prediction errors.

9. The method of claim 8, comprising:

35 k) assigning an index to each of the candidate modification processes and encoding the index as additional bits in the bitstream for transmission;

l) calculating for each candidate modification process a rate distortion cost based on the prediction error and the estimated cost of transmitting the index;

m) in step j, the modification process selected for calculating the residual signal is based on the prediction error and the estimated cost of transmitting the index.

5

10. The method of claim 9, wherein step l) comprises reconstructing at the encoder the original picture block using the inverse decoding process and comparing the reconstructed block to the original picture block to give a measure of the rate distortion.

10 11. A method of decoding video data with compression:

a) receiving an encoded residual signal containing coefficients representative of a coded region of a picture signal;

b) receiving a prediction block containing coefficients representative of a predicted picture region;

15 c) forming a modified prediction block by substituting the values of coefficients in the received prediction block for other predetermined values according to a modification process, the predetermined modification process including applying a mask to the received prediction block, the mask defining at least a first region of the prediction block where the values of the coefficients are preserved and at least a second region of the prediction block
20 where the values of coefficients are substituted for other predetermined values; and

d) calculating an original picture block containing coefficients representative of a picture region by adding the residual signal and the modified prediction block to each other.

12. The method of claim 11 comprising:

25 forming the modified prediction block in the frequency domain; and
calculating the original block in the frequency domain.

13. The method of claim 11 or 12, wherein the modification process is selected from a plurality of predetermined modification processes.

30

14. The method of claim 13, wherein the modification process selected from the plurality of predetermined modification processes depends on one or more of the shape and/or size of the prediction block, and the prediction encoding scheme used to generate the prediction block.

35

15. The method of claim 12, 13 or 14, wherein the predetermined plurality of modification processes include a predetermined plurality of differently shaped masks for application to the prediction block and/or a predetermined plurality of different processes for determining said values to be substituted.

5

16. The method of claim 15, wherein the predetermined plurality of processes for determining said values to be substituted include one or more of:

i) replacing the value of all coefficients in the at least a second region with zero;

ii) replacing the value of all coefficients in the at least a second region with a

10 predetermined non-zero value;

iii) replacing the value of all coefficients in the at least a second region with a value calculated according to a predetermined algorithm.

17. The method of any of claims 11 to 16, comprising:

15 receiving as a bitstream, the encoded residual signal and an indication indicating a modification processes to modify the prediction block; and

selecting a modification process to form the modified prediction block based on the received index.

20 18. A method of encoding video data with compression:

a) receiving an original picture block containing coefficients representative of a picture region;

b) generating a prediction block containing coefficients representative of a predicted picture region;

25 c) calculating a residual signal for transmission, by subtracting the original picture block and the prediction block from each other,

wherein generating the prediction block further comprises:

b1) generating coefficients at a first set of predetermined locations in the prediction block according to an intra-prediction or inter-prediction coding scheme;

30 and

b2) generating coefficients at a second set of predetermined locations in the prediction block according to a coefficient substitution process, wherein the coefficients in the second set of predetermined locations in the prediction block are not representative of coefficients representative of a predicted picture region.

35

19. The method of claim 18, wherein the coefficient substitution process includes one or more of:

i) setting a coefficient to have a value of zero;

ii) setting a coefficient to have a predetermined non zero-value;

5 iii) setting the value of all coefficients in the at least a second region with a value calculated according to a predetermined algorithm.

20. The method of claim 18 or 19, wherein the locations in the prediction block that are in the first and second sets of predetermined locations, and/or the coefficient substitution
10 process depend, on one or more of the shape and/or size of the prediction block, and the prediction encoding scheme used to generate the prediction block.

21. The method of any of claims 18 to 20, comprising:

15 transforming the original picture block and the prediction block into the frequency domain; and

 forming the modified prediction block and calculating the residual signal in the frequency domain.

22. The method of any of claims 18 to 21, comprising:

20 g) calculating the prediction error between the prediction block and the original picture block;

 h) generating a second and subsequent prediction block containing coefficients representative of a predicted picture region, wherein in the generation of the second prediction block, at least one of the coefficient substitution process and/or the locations of
25 the prediction block in the first and second sets of predetermined locations is varied, to generate a plurality of different candidate modified prediction blocks;

 i) for each of the plurality of different candidate modified prediction blocks, calculating the prediction error between the respective candidate modified prediction block and the original picture block; and

30 j) selecting one of the plurality of different candidate modified prediction blocks for use in the calculation of the residual signal based on the prediction error.

23. The method of claim 22, comprising in step h) setting the coefficients in the second set of predetermined locations from a look up table containing a plurality of candidate
35 values, wherein different look-up tables are consulted depending on the shape and/or size

of the prediction block, and the prediction encoding scheme used to generate the prediction block.

24. The method of claim 23, comprising:

5 k) assigning an index to each of the plurality of different candidate modified prediction blocks and encoding the index as additional bits in the bitstream for transmission;

10 l) calculating for each of the plurality of different candidate modified prediction blocks candidate a rate distortion cost based on the prediction error and the estimated cost of transmitting the index;

m) in step j, the selection of the prediction block for calculating the residual signal is based on the prediction error and the estimated cost of transmitting the index.

25. A method of decoding video data with compression:

15 a) receiving an encoded residual signal containing coefficients representative of a coded region of a picture signal;

b) generating a prediction block containing coefficients representative of a predicted picture region;

20 c) calculating an original picture block containing coefficients representative of a picture region by subtracting the residual signal and the prediction block from each other. wherein generating the prediction block further comprises:

b1) generating coefficients at a first set of predetermined locations in the prediction block according to an intra-prediction or inter-prediction coding scheme; and

25 b2) generating coefficients at a second set of predetermined locations in the prediction block according to a coefficient substitution process, wherein the coefficients in the second set of predetermined locations in the prediction block are not representative of coefficients representative of a predicted picture region.

30 26. The method of claim 25, wherein the coefficient substitution process includes one or more of:

i) setting a coefficient to have a value of zero;

ii) setting a coefficient to have a predetermined non zero-value;

35 iii) setting the value of all coefficients in the at least a second region with a value calculated according to a predetermined algorithm.

27. The method of claim 25 or 26, wherein the locations in the prediction block that are in the first and second sets of predetermined locations, and/or the coefficient substitution process depend, on one or more of the shape and/or size of the prediction block, and the prediction encoding scheme used to generate the prediction block.

5

28. The method of any of claims 25 to 27, comprising:
generating the prediction block in the frequency domain and calculating the reconstructed block in the frequency domain.

10 29. The method of any of claims 25 to 28 comprising:
receiving as a bitstream, the encoded residual signal and an index indicating a modification processes used at the encoder to modify the prediction block; and
selecting one or more of the locations in the prediction block that are in the first and second sets of predetermined locations, and/or the coefficient substitution process based
15 on the received index.

30. A computer readable medium having computer code stored thereon which when executed by a computer processor causes the computer processor to perform the steps of any of method claims 1 to 29.

20

31. A video encoder comprising:
a) a partition unit for outputting an original picture block containing coefficients representative of a picture region;
b) a prediction unit for generating a prediction block containing coefficients
25 representative of a predicted picture region;
c) a modified prediction block generating unit operable to form a modified prediction block by substituting the values of coefficients in the generated prediction block for other predetermined values according to a modification process, the predetermined modification process including applying a mask to the received prediction block, the mask defining at
30 least a first region of the prediction block where the values of the coefficients are preserved and at least a second region of the prediction block where the values of coefficients are substituted for other predetermined values; and
d) a calculation unit for calculating a residual signal for transmission, by subtracting the original picture block and the modified prediction block from each other.

35

32. The apparatus of claim 31 comprising:

a transform block for transforming the original picture block and the prediction block into the frequency domain; and

wherein the calculation block is coupled to the transform block for forming the modified prediction block and calculating the residual signal in the frequency domain.

5

33. The apparatus of claim 31 or 32, wherein the modified prediction block generating unit is operable to select the modification process from a plurality of predetermined modification processes.

10

34. The apparatus of claim 33, wherein the modified prediction block generating unit is operable to select the modification process from the plurality of predetermined modification processes based on one or more of the shape and/or size of the prediction block, and the prediction encoding scheme used to generate the prediction block.

15

35. The apparatus of claim 32, 33 or 34, wherein the predetermined plurality of modification processes include a predetermined plurality of differently shaped masks for application to the prediction block and/or a predetermined plurality of different processes for determining said values to be substituted.

20

36. The apparatus of claim 35, wherein the predetermined plurality of processes for determining said values to be substituted include one or more of:

i) replacing the value of all coefficients in the at least a second region with zero;

ii) replacing the value of all coefficients in the at least a second region with a predetermined non-zero value;

25

iii) replacing the value of all coefficients in the at least a second region with a value calculated according to a predetermined algorithm.

37. The apparatus of claim 33, wherein the modified prediction block generating unit is operable to:

30

e) select a candidate modification process from the plurality of predetermined modification processes;

f) form a modified prediction block according to the candidate modification process;

g) calculate the prediction error between the modified prediction block and the original picture block;

h) repeat steps e) to g) for each modification process in the plurality of predetermined modification processes and for each modification process calculate a prediction error between the modified prediction block and the original picture block;

5 i) select a modification process to use for forming the modified prediction block for use in the calculation of the residual signal from the plurality of predetermined modification processes;

j) quantise and encode the residual signal for transmission in a bitstream, wherein the selection of the modification process to use for calculating the residual signal is based on the prediction error.

10

38. The apparatus of claim 37, wherein the predetermined plurality of processes for determining said values to be substituted includes consulting a look-up table containing a plurality of candidate values, wherein different look-up tables are consulted depending on the candidate modification process, and wherein the candidate values in the look table
15 have been selected based on previous analysis of values that have provided favourable prediction errors.

39. The apparatus of claim 38, wherein the modified prediction block generating unit is operable to:

20

k) assign an index to each of the candidate modification processes and encoding the index as additional bits in the bitstream for transmission;

l) calculate for each candidate modification process a rate distortion cost based on the prediction error and the estimated cost of transmitting the index;

25

m) in step j, selected the modification process for calculating the residual signal based on the prediction error and the estimated cost of transmitting the index.

40. The apparatus of claim 39, wherein step l) comprises reconstructing at the encoder the original picture block using the inverse decoding process and comparing the reconstructed block to the original picture block to give a measure of the rate distortion.

30

41. A video decoder comprising:

a) a receiver for receiving an encoded residual signal containing coefficients representative of a coded region of a picture signal;

35

b) a prediction unit for generating a prediction block containing coefficients representative of a predicted picture region block,

c) a modified prediction block generating unit operable to form a modified prediction block by substituting the values of coefficients in the generated prediction block for other predetermined values according to a modification process, the predetermined modification process including applying a mask to the received prediction block, the mask defining at least a first region of the prediction block where the values of the coefficients are preserved and at least a second region of the prediction block where the values of coefficients are substituted for other predetermined values; and

d) a calculating unit for calculating an original picture block containing coefficients representative of a picture region by subtracting the residual signal and the modified prediction block from each other.

42. The apparatus of claim 41 comprising:

a transform block for transforming the original picture block to the spatial domain, wherein the transform block receives the original block from the calculating unit.

43. The apparatus of claim 41 or 42, wherein the modified prediction block generating unit selects the modification process from a plurality of predetermined modification processes.

44. The apparatus of claim 43, wherein the modified prediction block generating unit selects the modification process from the plurality of predetermined modification processes based on one or more of the shape and/or size of the prediction block, and the prediction encoding scheme used to generate the prediction block.

45. The apparatus of claim 42, 43 or 44, wherein the predetermined plurality of modification processes include a predetermined plurality of differently shaped masks for application to the prediction block and/or a predetermined plurality of different processes for determining said values to be substituted.

46. The apparatus of claim 45, wherein the predetermined plurality of processes for determining said values to be substituted include one or more of:

i) replacing the value of all coefficients in the at least a second region with zero;

ii) replacing the value of all coefficients in the at least a second region with a predetermined non-zero value;

iii) replacing the value of all coefficients in the at least a second region with a value calculated according to a predetermined algorithm.

47. The apparatus of any of claims 41 to 46, wherein the modified prediction block generating unit is operable to select a modification process to form the modified prediction block based on a indication received in the bitstream indicating a modification processes to
5 modify the prediction block.
48. A video encoder comprising method of encoding video data with compression:
a) a partition unit for outputting an original picture block containing coefficients representative of a picture region;
10 b) a prediction unit for generating a prediction block containing coefficients representative of a predicted picture region;
c) a calculation block for calculating a residual signal for transmission, by subtracting the original picture block and the prediction block from each other, wherein the prediction unit is operable to:
15 b1) generate coefficients at a first set of predetermined locations in the prediction block according to an intra-prediction or inter-prediction coding scheme; and
b2) generate coefficients at a second set of predetermined locations in the prediction block according to a coefficient substitution process, wherein the
20 coefficients in the second set of predetermined locations in the prediction block are not representative of coefficients representative of a predicted picture region.
49. The apparatus of claim 48, wherein the coefficient substitution process includes one or more of:
25 i) setting a coefficient to have a value of zero;
ii) setting a coefficient to have a predetermined non zero-value;
iii) setting the value of all coefficients in the at least a second region with a value calculated according to a predetermined algorithm.
- 30 50. The apparatus of claim 48 or 49, wherein the locations in the prediction block that are in the first and second sets of predetermined locations, and/or the coefficient substitution process depend, on one or more of the shape and/or size of the prediction block, and the prediction encoding scheme used to generate the prediction block.
- 35 51. The apparatus of any of claims 48 to 50, comprising:

a transform block for transforming the original picture block and the prediction block into the frequency domain; and

wherein the calculation block is coupled to the transform block for forming the modified prediction block and calculating the residual signal in the frequency domain.

5

52. The apparatus of any of claims 48 to 51, wherein the prediction unit is operable to:

g) calculate the prediction error between the prediction block and the original picture block;

10 h) generate a second and subsequent prediction block containing coefficients representative of a predicted picture region, wherein in the generation of the second prediction block, at least one of the coefficient substitution process and/or the locations of the prediction block in the first and second sets of predetermined locations is varied, to generate a plurality of different candidate modified prediction blocks;

15 i) for each of the plurality of different candidate modified prediction blocks, calculate the prediction error between the respective candidate modified prediction block and the original picture block; and

j) select one of the plurality of different candidate modified prediction blocks for use in the calculation of the residual signal based on the prediction error.

20 53. The apparatus of claim 52, wherein the prediction unit is operable to in step h) to set the coefficients in the second set of predetermined locations from a look up table containing a plurality of candidate values, wherein different look-up tables are consulted depending on the shape and/or size of the prediction block, and the prediction encoding scheme used to generate the prediction block.

25

54. The apparatus of claim 53, wherein the prediction unit is operable to:

k) assign an index to each of the plurality of different candidate modified prediction blocks and encoding the index as additional bits in the bitstream for transmission;

30 l) calculate for each of the plurality of different candidate modified prediction blocks candidate a rate distortion cost based on the prediction error and the estimated cost of transmitting the index;

m) in step j, the selection of the prediction block for calculating the residual signal is based on the prediction error and the estimated cost of transmitting the index.

35 55. A video decoder comprising:

a) a receiver for receiving an encoded residual signal containing coefficients representative of a coded region of a picture signal;

b) a prediction unit for generating a prediction block containing coefficients representative of a predicted picture region;

5 c) a calculation unit for calculating an original picture block containing coefficients representative of a picture region by subtracting the residual signal and the prediction block from each other.

wherein the prediction unit is operable to:

10 b1) generate coefficients at a first set of predetermined locations in the prediction block according to an intra-prediction or inter-prediction coding scheme; and

b2) generate coefficients at a second set of predetermined locations in the prediction block according to a coefficient substitution process, wherein the coefficients in the second set of predetermined locations in the prediction block are
15 not representative of coefficients representative of a predicted picture region.

56. The apparatus of claim 55, wherein the coefficient substitution process includes one or more of:

i) setting a coefficient to have a value of zero;
20 ii) setting a coefficient to have a predetermined non zero-value;
iii) setting the value of all coefficients in the at least a second region with a value calculated according to a predetermined algorithm.

57. The apparatus of claim 55 or 56, wherein the locations in the prediction block that
25 are in the first and second sets of predetermined locations, and/or the coefficient substitution process depend, on one or more of the shape and/or size of the prediction block, and the prediction encoding scheme used to generate the prediction block.

58. The apparatus of any of claims 55 to 57, comprising:
30 a transform block for transforming the original picture block to the spatial domain, wherein the transform block receives the original block from the calculating unit.

59. The apparatus of any of claims 45 to 58 wherein the prediction unit is operable to
35 select one or more of the locations in the prediction block that are in the first and second sets of predetermined locations, and/or the coefficient substitution process based on an indication received in the bitstream indicating a modification processes.

60. A video encoding method substantially as described herein and with respect to the drawings.

5 61. A video decoding method substantially as described herein and with respect to the drawings.

62. A video encoder substantially as described herein and with respect to the drawings.

63. A video decoder substantially as described herein and with respect to the drawings.

