PREDICTION OF TRANSFORM COEFFICIENTS FOR IMAGE COMPRESSION

Copyright Authorization

A portion of the disclosure of this patent document contains material that is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

Background

Block Transform-Based Coding

Transform coding is a compression technique used in many audio, image and video compression systems. Uncompressed digital image and video is typically represented or captured as samples of picture elements or colors at locations in an image or video frame arranged in a two-dimensional (2D) grid. This is referred to as a spatial-domain representation of the image or video. For example, a typical format for images consists of a stream of 24-bit color picture element samples arranged as a grid. Each sample is a number representing color components at a pixel location in the grid within a color space, such as RGB, or YIQ, among others. Various image and video systems may use various different color, spatial and time resolutions of sampling. Similarly, digital audio is typically represented as time-sampled audio signal stream. For example, a typical audio format consists of a stream of 16-bit amplitude samples of an audio signal taken at regular time intervals.

Traditionally, compression of video is performed by compressing the first image frame, and compressing differences between successive frames. This process is repeated periodically across the video sequence. Therefore, the compression of video is closely related to the compression of "still" images.

Uncompressed digital audio, image and video signals can consume considerable storage and transmission capacity. Transform coding reduces the size of digital audio, images and video by transforming the spatial-domain representation of the signal into a frequency-domain (or other like transform

domain) representation, and then reducing resolution of certain generally less perceptible frequency components of the transform-domain representation. This generally produces much less perceptible degradation of the digital signal compared to reducing color or spatial resolution of images or video in the spatial domain, or of audio in the time domain.

More specifically, a typical block transform-based codec 100 shown in Figure 1 divides the uncompressed digital image's pixels into fixed-size two dimensional blocks (X_1, ... X_n), each block possibly overlapping with other blocks. A linear transform 120-121 that does spatial-frequency analysis is applied to each block, which converts the spaced samples within the block to a set of frequency (or transform) coefficients generally representing the strength of the digital signal in corresponding frequency bands over the block interval. For compression, the transform coefficients may be selectively quantized 130 (i.e., reduced in resolution, such as by dropping least significant bits of the coefficient values or otherwise mapping values in a higher resolution number set to a lower resolution), and also entropy or variable-length coded 130 into a compressed data stream. At decoding, the transform coefficients will inversely transform 170-171 to nearly reconstruct the original color/spatial sampled image/video signal (reconstructed blocks $\hat{X}_1$, ... $\hat{X}_n$).

The block transform 120-121 can be defined as a mathematical operation on a vector x of size N. Most often, the operation is a linear multiplication, producing the transform domain output $y = M x$, M being the transform matrix. When the input data is arbitrarily long, it is segmented into N sized vectors and a block transform is applied to each segment. For the purpose of data compression, reversible block transforms are chosen. In other words, the matrix M is invertible. In multiple dimensions (e.g., for image and video), block transforms are typically implemented as separable operations. The matrix multiplication is applied separably along each dimension of the data (i.e., both rows and columns).

For compression, the transform coefficients (components of vector y) may be selectively quantized (i.e., reduced in resolution, such as by dropping least significant bits of the coefficient values or otherwise mapping values in a higher

resolution number set to a lower resolution), and also entropy or variable-length coded into a compressed data stream.

At decoding in the decoder 150, the inverse of these operations (dequantization/entropy decoding 160 and inverse block transform 170-171) are applied on the decoder 150 side, as show in Fig. 1. While reconstructing the data, the inverse matrix M^{-1} (inverse transform 170-171) is applied as a multiplier to the transform domain data. When applied to the transform domain data, the inverse transform nearly reconstructs the original time-domain or spatial-domain digital media.

In many block transform-based coding applications, the transform is desirably reversible to support both lossy and lossless compression depending on the quantization factor. With no quantization (generally represented as a quantization factor of 1) for example, a codec utilizing a reversible transform can exactly reproduce the input data at decoding. However, the requirement of reversibility in these applications constrains the choice of transforms upon which the codec can be designed.

Many image and video compression systems, such as MPEG and Windows Media, among others, utilize transforms based on the Discrete Cosine Transform (DCT). The DCT is known to have favorable energy compaction properties that result in near-optimal data compression. In these compression systems, the inverse DCT (IDCT) is employed in the reconstruction loops in both the encoder and the decoder of the compression system for reconstructing individual image blocks.

Transform Coefficient Prediction

As just noted, block transforms commonly use the discrete cosine transform (DCT) or variants. At high levels of loss, block transforms suffer from visual artifacts due to annoying block discontinuities. A "lapped transform" technique, in which the transformation windows overlap, can be used to smooth reconstructions even under loss.

In both block and lapped transforms, long linear features oriented along the horizontal or vertical directions cause high transform values along the left or top edges of transform domain blocks. The left and top edges are often referred to as DCAC values. This name is because these are the coefficients that are DC in one

direction and AC in the other. The top left position is called the DC value (DC in both directions).

Block transforms often show a correlation between blocks. It can be easily appreciated that the DC coefficients of adjacent blocks are correlated and tend to be close in a probabilistic sense. Less evident is the correlation between the corresponding DCAC coefficients of adjacent blocks. Notably, if a certain area of an image shows strong horizontal features (such as line or patterns), the transform coefficients which are DC in the horizontal direction and AC in the vertical direction show inter block numerical correlation as well.

The process of exploiting inter-block DC and DCAC continuity by forming a prediction for the DC and DCAC terms from neighboring blocks, and encoding prediction differences is commonly referred to as "DCAC prediction". This term also covers the decoder side processes of recovering the original (or approximate) DC & DCAC transform coefficients. The DCAC terms being predicted may be a subset of all DCAC terms, determined by the direction of prediction.

Summary

A digital media coding and decoding technique and realization of the technique in a digital media codec described herein uses a transform coefficient prediction that takes into account a dominant directionality of the digital media data (e.g., an image with strong horizontal or vertical features), and further operates compatibly with a two-stage transform.

For DC and DCAC coefficients from an inner stage transform of a macroblock, the codec calculates and compares directionality metrics based on inner stage transform DC coefficients of neighboring macroblocks to determine dominant directionality. The determination of directional dominance can also take into account information from other channels (e.g., color or chrominance). In the absence of directional dominance, the DC coefficient of the macroblock is predicted from an average of DC coefficients of preceding neighbor macroblocks. Directional prediction may be skipped for the DCAC coefficients of the macroblock if the neighboring macroblock has a different quantizer. Alternatively, directional prediction of DCAC coefficients may be restricted to macroblock(s) that has(have) the same quantizer.

For DCAC coefficients from an outer stage transform of blocks within the macroblock, the codec calculates and compares directionality metrics based on the inner stage transform DCAC coefficients of the macroblock to detect dominant directionality. The determination of directional dominance can also take into account information from other channels (e.g., color or chrominance). This way the directional prediction determination for outer stage DCAC coefficients can be made solely based on information in the macroblock. If directional dominance is found, the outer stage DCAC coefficients of the blocks within the macroblock are predicted unidirectionally from the dominant direction.

This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

Brief Description Of The Drawings

Figure 1 is a block diagram of a conventional block transform-based codec in the prior art.

Figure 2 is a flow diagram of a representative encoder incorporating a predictive transform coefficient coding.

Figure 3 is a flow diagram of a representative decoder incorporating the predictive transform coefficient coding.

Figure 4 is a diagram of a 4x4 transform block structure illustrating DC and DCAC coefficients.

Figure 5 is a diagram of a 422 chroma low pass block structure of coefficients produced in an inner stage transform of the encoder of Figure 3 for a YUV 4:2:2 color format.

Figure 6 is a diagram of a 420 chroma low pass block structure of coefficients produced in an inner stage transform of the encoder of Figure 3 for a YUV 4:2:0 color format image.

Figure 7 is a diagram illustrating an example of DCAC prediction.

Figure 8 is a diagram identifying the direction of prediction in the example DCAC prediction of Figure 7.

Figure 9 is a diagram illustrating the macroblocks used in determining the DC prediction mode of a macroblock.

Figure 10 is a pseudo-code listing of a DC prediction mode determination in the predictive transform coefficient coding in the encoder of Figure 3 and decoder of Figure 4.

Figure 11 is a pseudo-code listing of a low pass DCAC prediction mode determination in the predictive transform coefficient coding in the encoder of Figure 3 and decoder of Figure 4.

Figure 12 is a pseudo-code listing of a high pass DCAC prediction mode determination in the predictive transform coefficient coding in the encoder of Figure 3 and decoder of Figure 4.

Figure 13 is a diagram illustrating left high pass DCAC prediction of a macroblock.

Figure 14 is a diagram illustrating top high pass DCAC prediction of a macroblock.

Figure 15 is a block diagram of a suitable computing environment for implementing the adaptive coding of wide range coefficients of Figure 4.

Detailed Description

The following description relates to coding and decoding techniques that provides an efficient coding/decoding of transform coefficients of a block transform-based codec based on coefficient prediction (referred to herein as "Predictive Transform Coefficient Coding"). The following description describes an example implementation of the technique in the context of a digital media compression system or codec. The digital media system codes digital media data in a compressed form for transmission or storage, and decodes the data for playback or other processing. For purposes of illustration, this exemplary compression system incorporating this predictive transform coefficient coding is an image or video compression system. Alternatively, the technique also can be incorporated into compression systems or codecs for other 2D data. The predictive transform coefficient coding technique does not require that the digital media compression system encodes the compressed digital media data in a particular coding format.

1. Encoder/Decoder

Figures 2 and 3 are a generalized diagram of the processes employed in a representative 2-dimensional (2D) data encoder 200 and decoder 300. The diagrams present a generalized or simplified illustration of a compression system incorporating the 2D data encoder and decoder that implement the block pattern coding. In alternative compression systems using the block pattern coding, additional or fewer processes than those illustrated in this representative encoder and decoder can be used for the 2D data compression. For example, some encoders/decoders may also include color conversion, color formats, scalable coding, lossless coding, macroblock modes, etc. The compression system (encoder and decoder) can provide lossless and/or lossy compression of the 2D data, depending on the quantization which may be based on a quantization parameter varying from lossless to lossy.

The 2D data encoder 200 produces a compressed bitstream 220 that is a more compact representation (for typical input) of 2D data 210 presented as input to the encoder. For example, the 2D data input can be an image, a frame of a video sequence, or other data having two dimensions. The 2D data encoder tiles 230 the input data into macroblocks, which are 16×16 pixels in size in this representative encoder. The 2D data encoder further tiles each macroblock into 4×4 blocks. A "forward overlap" operator 240 is applied to each edge between blocks, after which each 4×4 block is transformed using a block transform 250. This block transform 250 can be the reversible, scale-free 2D transform described by Srinivasan, U.S. Patent Application No. 11/015,707, entitled, "Reversible Transform For Lossy And Lossless 2-D Data Compression," filed December 17, 2004. The overlap operator 240 can be the reversible overlap operator described by Tu et al., U.S. Patent Application No. 11/015,148, entitled, "Reversible Overlap Operator for Efficient Lossless Data Compression," filed December 17, 2004; and by Tu et al., U.S. Patent Application No. 11/035,991, entitled, "Reversible 2-Dimensional Pre-/Post-Filtering For Lapped Biorthogonal Transform," filed January 14, 2005. Alternatively, the discrete cosine transform or other block transforms and overlap operators can be used. Subsequent to the transform, the DC coefficient 260 of each 4×4 transform block is subject to a similar processing chain (tiling, forward

overlap, followed by 4×4 block transform). The resulting DC transform coefficients and the AC transform coefficients are quantized 270, entropy coded 280 and packetized 290.

The decoder performs the reverse process. On the decoder side, the transform coefficient bits are extracted 310 from their respective packets, from which the coefficients are themselves decoded 320 and dequantized 330. The DC coefficients 340 are regenerated by applying an inverse transform, and the plane of DC coefficients is "inverse overlapped" using a suitable smoothing operator applied across the DC block edges. Subsequently, the entire data is regenerated by applying the 4×4 inverse transform 350 to the DC coefficients, and the AC coefficients 342 decoded from the bitstream. Finally, the block edges in the resulting image planes are inverse overlap filtered 360. This produces a reconstructed 2D data output.

In an exemplary implementation, the encoder 200 (Figure 2) compresses an input image into the compressed bitstream 220 (e.g., a file), and the decoder 300 (Figure 3) reconstructs the original input or an approximation thereof, based on whether lossless or lossy coding is employed. The process of encoding involves the application of a forward lapped transform (LT) discussed below, which is implemented with reversible 2-dimensional pre-/post-filtering also described more fully below. The decoding process involves the application of the inverse lapped transform (ILT) using the reversible 2-dimensional pre-/post-filtering.

The illustrated LT and the ILT are inverses of each other, in an exact sense, and therefore can be collectively referred to as a reversible lapped transform. As a reversible transform, the LT/ILT pair can be used for lossless image compression.

The input data 210 compressed by the illustrated encoder 200/decoder 300 can be images of various color formats (e.g., RGB/YUV4:4:4, YUV4:2:2 or YUV4:2:0 color image formats). Typically, the input image always has a luminance (Y) component. If it is a RGB/YUV4:4:4, YUV4:2:2 or YUV4:2:0 image, the image also has chrominance components, such as a U component and a V component. The separate color planes or components of the image can have different spatial resolutions. In case of an input image in the YUV 4:2:0 color

format for example, the U and V components have half of the width and height of the Y component.

As discussed above, the encoder 200 tiles the input image or picture into macroblocks. In an exemplary implementation, the encoder 200 tiles the input image into 16x16 macroblocks in the Y channel (which may be 16x16, 16x8 or 8x8 areas in the U and V channels depending on the color format). Each macroblock color plane is tiled into 4x4 regions or blocks. Therefore, a macroblock is composed for the various color formats in the following manner for this exemplary encoder implementation:

1. For a grayscale image, each macroblock contains 16 4x4 luminance (Y) blocks.
2. For a YUV4:2:0 format color image, each macroblock contains 16 4x4 Y blocks, and 4 each 4x4 chrominance (U and V) blocks.
3. For a YUV4:2:2 format color image, each macroblock contains 16 4x4 Y blocks, and 8 each 4x4 chrominance (U and V) blocks.
4. For a RGB or YUV4:4:4 color image, each macroblock contains 16 blocks each of Y, U and V channels.

Figures 4, 5 and 6 illustrate examples of various transform blocks in the representative encoder/decoder. Figure 4 illustrates a 4x4 transform block for blocks in the luminance channel of the various formats, as well as the chrominance channels of a YUV 4:4:4 color format image. Figure 5 shows a 422 chroma low pass block format, which contains the transform block coefficients from the inner stage transform of the chroma channel of a YUV 422 color format image. Figure 6 shows a 420 chroma low pass block containing the coefficients from the inner stage transform of the chroma channel of a YUV 4:2:0 color format image.

With reference to Figure 4, the following description uses a block transform convention or notation that is in line with matrix notation, with ordered basis functions. In the transform coefficient block 400 shown in Figure 4 for example, the coefficient representing DC frequency is the first row/column of the transform coefficient block (labeled coefficient '0'), and that of the highest AC frequency is the last row/column (labeled coefficient '15'). Further, DCAC coefficients in the top row of the block (coefficients labeled '1,' '2' and '3' in Figure 4) correspond to

vertical line patterns, and DCAC values in the left column (coefficients labeled '4,' '8' and '12' in Figure 4) correspond to horizontal line patterns. In actual practical implementation, the transform block may be transposed, in which case appropriate coefficient reindexing must be taken into account.

2. Predictive Transform Coefficient Coding Overview

As discussed briefly in the Background section above, DCAC prediction is a process for exploiting inter-block DC and DCAC continuity by forming a prediction for the DC and DCAC terms from neighboring blocks, and encoding prediction differences (as well as the inverse process on the decoder).

With reference to Figures 7 and 8, an example of a DCAC prediction is shown. In this example, the DC coefficient ('0') and the DCAC coefficient set corresponding to either vertical patterns ('1,' '2,' and '3') or horizontal patterns ('4,' '8,' and '12') of each transform block are predicted from those of a preceding neighboring block. In other words, the corresponding coefficients of the neighboring block are taken as a "predictor" for those of the current block. The direction of prediction for each block in this example is identified in the diagram in Figure 8 (with the notation 'N' indicating no prediction, 'L' indicating left prediction and 'T' indicating top prediction), as well as being illustrated by the arrows in the diagram of Figure 7. When the direction of prediction is from the top, the coefficients '1', '2' and '3' are predicted from the corresponding coefficients in the block to the top. When the direction of prediction is from the left, the coefficients '4', '8' and '12' are predicted from the corresponding coefficients in the block to the left. Specifically, no prediction is used for the top left block. The second and third blocks of the top row are predicted from the block to their left. The blocks in the second row are predicted from above, left and above, respectively. With the DCAC prediction, the encoder encodes a coefficient as the difference from its predictor. If the prediction is reasonably accurate, the differences from the predictor will be zero with a high probability, which can be efficiently encoded with variable length entropy coding techniques (e.g., run length coding) at the entropy coding 280 (Figure 2). For example, the variable length entropy coding technique described by Srinivasan, U.S. Patent Application No.

TBD, entitled "Efficient Encoding and Decoding of Transform Blocks," filed August 12, 2005.

The above representative encoder 200 (Figure 2) and decoder 300 (Figure 3) present certain challenges to use of DCAC prediction. First, the representative encoder and decoder use a 2-stage transform. These can be referred to as "outer" and "inner" stages, which correspond to the highpass and lowpass bands respectively. In particular, the DC coefficients 260 (Figure 20 of the blocks are subjected to a second stage transform (the "inner" stage). As a result, the DC coefficients are not available to the decoder for decoding the other coefficients without first performing the inverse transform of that stage.

Second, the representative encoder and decoder use a lapped transform. By its design and construction, the lapped transform already implicitly extracts continuity across adjacent blocks. For this reason, the further application of DCAC prediction could end up degrading the coding performance (as compared to without DCAC prediction).

Third, for the purpose of minimizing memory footprint and complexity, it is desirable to keep inter-macroblock DCAC prediction to a minimum.

Finally, the representative encoder and decoder may apply different quantizers in different macroblocks, which further complicates the DCAC prediction.

The predictive transform coefficient coding techniques described herein provide various enhancements to the conventional DCAC prediction. In particular, an exemplary implementation of the predictive transform coefficient coding techniques described below addresses the above-listed issues using a unique set of prediction rules. In summary, these rules include:

1. The DC coefficients (of the inner transform) are predicted based on DC coefficients from causal neighboring blocks based at least in part on color information when available.
2. The DCAC coefficients of the inner transform use a prediction direction derived from that of the DC coefficients, and also rely on out-of-macroblock information.

3. The prediction of DCAC coefficients of the inner transform is skipped (i.e. 0 is used as the predictor) when the current and predicting macroblocks have different quantizers.
4. The prediction of DCAC coefficients of the outer transform is performed purely within the macroblock.
5. The prediction direction of the outer transform DCAC coefficients is derived from DCAC coefficients of the inner transform of the same macroblock.

In the description below, the predictive transform coefficient coding is described as being performed on the quantized transform coefficients (e.g., the transform coefficients after quantization 270 in the encoder 200 of Figure 2, and before de-quantization 330 in the decoder 300 of Figure 3). However, it should be understood that alternative implementations of the predictive transform coefficient coding technique could be performed on the unquantized transform coefficients (or on dequantized values on the decoder side).

The predictive transform coefficient coding in the representative encoder/decoder includes three levels of prediction that are detailed in the following sections. These include:

1. DC prediction, which is the prediction of DC coefficients of the inner transform,
2. Low pass DCAC prediction, which is the prediction of DCAC coefficients of the inner transform, and
3. High pass DCAC prediction, which is the prediction of DCAC coefficients of the outer transform.

2.1 DC Prediction

As described above, the representative encoder 200 (Figure 2) tiles an image into macroblocks of 16x16 pixels, and further tiles the macroblocks into transform blocks of 4x4 pixels each. The outer stage transform is applied to the transform blocks, producing blocks containing 4x4 transform coefficients as shown in Figure 4 (for the luminance channel, as well as the chrominance channels of a YUV 4:4:4 color format image). The DC coefficients in these 16 transform coefficient blocks of the macroblock are separated (forming a 4x4 block), and an inner stage of the

transform is applied to this block. The resulting inner stage transform block again has 4x4 coefficients. The DC coefficient (labeled '0') of this inner stage transform block is referred to herein as the DC coefficient of the macroblock. The DCAC coefficients ('1,' '2,' '3,' '4,' '8,' and '12') of the inner stage transform block are referred to herein as the low pass DCAC coefficients of the macroblock. This macroblock structure can be varied in alternative encoders and decoder employing the predictive transform coefficient coding.

When encoding the coefficient of a macroblock, the predictive transform coefficient coding chooses from four modes for the prediction of the DC coefficient of the macroblock. These modes are:

1. Predict from left (i.e., the predictor for the DC coefficient of the macroblock is the DC coefficient of the macroblock to its left, or *predictor* = *DC* [*left_MB*]).
2. Predict from top (i.e., the predictor is the DC coefficient of the macroblock above it, *predictor* = *DC* [*top_MB*])
3. Predict from left and top (i.e., the predictor is an average of the DC coefficients of the macroblocks to its left and above, *predictor* = (*DC* [*left_MB*] + *DC* [*top_MB*]) / 2)
4. Null predict (i.e., no prediction, *predictor* = 0)

The encoder determines which prediction mode to use for a macroblock according to the procedure illustrated by the pseudo-code listing 1000 in Figure 10. In the listing 1000, the value $[m_x,m_y]$ is an index of the current macroblock in an image (or an image tile, if tiling is used) in terms of an offset number of macroblocks in horizontal (x) and vertical directions (y) starting from a top, left macroblock at [0,0].

In the illustrated procedure, the encoder determines which DC prediction mode is used based on the position of the macroblock within the image, as well as a consideration of the DC coefficients of the macroblocks to the left, top and top-left of the macroblock. More particularly, the encoder chooses the no prediction mode for the top-left macroblock of the image (i.e., the current macroblock $[m_x,m_y]=[0,0]$). The encoder chooses top prediction for macroblocks along a left edge in the image (i.e., where index (m_x) is 0). The encoder chooses left prediction

for macroblocks in a top row of macroblocks in the image (i.e., where index (m_y) is 0).

For all other (e.g., those in the interior of the image) macroblocks, the encoder determines which prediction mode to use based on a metric of dominant directionality. For example, an area of an image with horizontal stripes has a horizontal dominant directionality. In other words, the horizontal neighbor of the current macroblock is a better predictor of its DC coefficient, than its vertical neighbor. As illustrated in Figure 9, this metric is calculated based on directional differences in DC coefficients of the current (X) macroblock's left (L) and top (T) neighbors from the current macroblock's diagonal neighbor (D). If the image has color or chrominance channels (e.g., for YUV color format images, but not grayscale images), the metric may further consider the DC coefficients of the corresponding macroblocks of the chrominance channels. In effect, the procedure determines that there is a vertical dominant directionality if the diagonal neighbor's DC coefficient is substantially closer to that of the left neighbor. On the other hand, when the diagonal neighbor DC coefficient is substantially closer to that of the top neighbor, then the metric will show a horizontal dominant directionality. The metric determines the "dominancy" based on a weighting factor (orient_weight). In the illustrated procedure, this weighting factor is four. In other words, a direction is considered dominant when its directional difference is more than four times that of the other direction. However, the weighting factor is a parameter that can be set to a different value in alternative implementations.

In instances where there is neither horizontal or vertical dominancy, then the encoder picks the combined left and top prediction mode (# 3 in the above list), where the predictor is an average of the DC coefficients of top and left macroblocks.

2.2 Low pass DCAC Prediction

With reference again to Figures 2 and 4, the representative encoder 200 includes a second or inner stage transform, which is applied to the DC coefficients 260 (Figure 2) from the outer stage transform. The DCAC coefficients (at locations '1,' '2,' '3,' '4,' '8' and '12') of the resulting coefficient block 400 (Figure 4) are herein referred to as the low pass DCAC coefficients.

The representative encoder and decoder use three prediction modes for the prediction of the low pass DCAC coefficients from the inner stage transform of a macroblock. These modes include:

1. Predict from left (i.e., the predictor for the low pass DCAC coefficients of the macroblock is the corresponding DCAC coefficient of the macroblock to its left, or *predictor* = *DCAC* [*left_MB*]) – in which case coefficients marked '4', '8' and '12' alone are predicted;
2. Predict from top (i.e., the predictor is the corresponding low pass DCAC coefficient of the macroblock above it, *predictor* = *DCAC* [*top_MB*]) – in which case coefficients marked '1', '2' and '3' alone are predicted; and
3. Null predict (no predication, or *predictor* = 0).

The encoder determines which prediction mode to use for a macroblock according to the procedure illustrated by the pseudo-code listing 1100 in Figure 11. In the illustrated procedure, the encoder determines which low pass DCAC prediction mode is used based on the DC prediction mode of the macroblock, together with the quantizer indices of the current macroblock and macroblock that is the DC predictor. This rule ensures that the prediction of inner transform DCAC coefficients does not take place across macroblocks with different quantizers. Further, DCAC is predicted only if one direction is dominant, which is derived per the DC prediction mode procedure detailed above (i.e., the DC prediction mode is "predict from left" or "predict from top" when horizontal or vertical dominance is found).

With reference to Figure 5, the DCAC prediction procedure for YUV 4:2:2 color format images has a special case for the low pass DCAC coefficient at position labeled '5' in the block. In the special case, the coefficient labeled '5' is predicted from the coefficient '1' when the DC prediction mode is "predict from top" regardless of the DCAC prediction mode.

2.2 High pass DCAC Prediction

With reference again to Figure 2, the prediction for high pass DCAC coefficients is complicated in the representative encoder/decoder due to the two stage transform. The high pass DCAC coefficients are those of the highpass

transform coefficients 262 produced from the outer stage transform 250. Further, as previously noted, the prediction is performed after quantization 270 in the representative encoder, such that the prediction in the representative decoder also takes place at decoding 320 (Figure 3) prior to the dequantization 330 (Figure 3). At this point in the decoding process, the inverse inner stage transform has not yet been performed to reconstruct the DC coefficients of the outer stage transform blocks. Likewise, at entropy encoding 280, the representative encoder has already applied the inner stage transform to these DC coefficients 260 of the outer stage transform blocks, such that the outer stage DC coefficients 260 are no longer available without also performing the inverse inner stage transform. It is therefore undesirable to base high pass DCAC prediction on the DC coefficients 260 of the outer transform. Otherwise, the decoders would be constrained to have the same numerical accuracy. The encoders would be constrained to run a decoder loop (i.e., the inverse inner stage transform), which is not feasible for a codec operating on an "open-loop" basis. For these reasons, the representative encoder/decoder bases the high pass DCAC prediction coding only on the decoded quantized values prior to inverse transform.

The two stage transform of the representative encoder/decoder is advantageous in that some information about the nature of the current macroblock can already be known from the low pass transform coefficients, even before the high pass transform coefficients are decoded. This information is used to determine the directionality of prediction. The representative encoder/decoder uses a simple yet effective metric to determine the dominant directionality of the macroblock based only on the inner transform coefficients.

The representative encoder and decoder use three prediction modes for the prediction of the high pass DCAC coefficients from the outer stage transform of a macroblock. These modes include:

1. Predict from left (i.e., the predictor for the high pass DCAC coefficient of the block is the correspondingly-located high pass DCAC coefficient of the block to its left, or *predictor* = *DCAC* [*left_MB*], as illustrated for left-predicted DCAC coefficients of macroblock 1300 in Figure 13);

2. Predict from top (i.e., the predictor is the corresponding location, high pass DCAC coefficient of the block above it, *predictor* = *DCAC* [*top_MB*], as illustrated for top-predicted DCAC coefficients of macroblock 1400 in Figure 14); and
3. Null predict (no predication, or *predictor* = 0).

In the representative encoder/decoder, the same mode is applied to all blocks within a macroblock for which in-macroblock prediction is possible (but, alternative implementations need not apply the same mode to all blocks in a macroblock). In other words, no prediction is made for high pass DCAC coefficients of blocks that have no valid reference within the macroblock, even though the "predict from left" or "predict from top" mode is selected for the macroblock.

The encoder determines which prediction mode to use for a macroblock according to the procedure illustrated by the pseudo-code listing 1200 in Figure 12. In this procedure, a weighting factor is again used to test directional dominance (similar to the DC prediction procedure 1000 of Figure 10). The representative encoder and decoder use the value of the weighting factor for both procedures, but the same value parameter need not be used in both procedures in alternative implementations. The directional differences in the high pass DCAC prediction procedure 1200 are calculated based on the low pass DCAC coefficients from the inner stage transform of the macroblock, as well as the low pass DCAC coefficients from the chrominance channels (when present). In the pseudo-code listing 1200, the values "lowpass[*i*]" is the low pass DCAC coefficient at the corresponding index (*i*) position labeled in left-to-right, top-to-bottom order as in the block structure shown in Figure 4. The values "lowpass_U[*i*]" and "lowpass_V[*i*]" are the low pass DCAC coefficients from the inner stage transform blocks of the U and V chrominance channels of a YUV color format image. These chrominance channels have coefficient indices as in the block structures 400, 500 and 600 shown in Figures 4, 5 and 6 for YUV 4:4:4, YUV 4:2:2 and YUV 4:2:0, respectively.

In the illustrated high pass DCAC prediction procedure, the only information that needs to be kept available from a macroblock for use in the

prediction of neighboring macroblocks is one DC coefficient and 6 DCAC coefficients per macroblock channel (fewer for the chrominance channels of YUV 4:2:0 and YUV 4:2:2 color format images). This is at most 21 coefficients (in the YUV 4:4:4 case) that need to be cached per macroblock. Further, the coefficients used for prediction from the left can be discarded after the next macroblock is encoded/decoded. For YUV 4:4:4, it is therefore only necessary to cache 12 coefficients per macroblock for use in the next row of macroblocks.

When tiling is used by the encoder/decoder, each tile is deemed to be a separate image for the purpose of the predictive transform coefficient coding. This assures independent decoding of image tiles.

3. Computing Environment

The above described encoder 200 (Figure 2) and decoder 300 (Figure 3) and techniques for predictive transform coefficient coding can be performed on any of a variety of devices in which digital media signal processing is performed, including among other examples, computers; image and video recording, transmission and receiving equipment; portable video players; video conferencing; and etc. The digital media coding techniques can be implemented in hardware circuitry, as well as in digital media processing software executing within a computer or other computing environment, such as shown in Figure 15.

Figure 15 illustrates a generalized example of a suitable computing environment (1300) in which described embodiments may be implemented. The computing environment (1300) is not intended to suggest any limitation as to scope of use or functionality of the invention, as the present invention may be implemented in diverse general-purpose or special-purpose computing environments.

With reference to Figure 15, the computing environment (1300) includes at least one processing unit (1510) and memory (1520). In Figure 15, this most basic configuration (1530) is included within a dashed line. The processing unit (1510) executes computer-executable instructions and may be a real or a virtual processor. In a multi-processing system, multiple processing units execute computer-executable instructions to increase processing power. The memory (1520) may be volatile memory (e.g., registers, cache, RAM), non-volatile memory (e.g., ROM,

EEPROM, flash memory, etc.), or some combination of the two. The memory (1520) stores software (1580) implementing the described predictive transform coefficient coding techniques.

A computing environment may have additional features. For example, the computing environment (1500) includes storage (1540), one or more input devices (1550), one or more output devices (1560), and one or more communication connections (1570). An interconnection mechanism (not shown) such as a bus, controller, or network interconnects the components of the computing environment (1500). Typically, operating system software (not shown) provides an operating environment for other software executing in the computing environment (1500), and coordinates activities of the components of the computing environment (1500).

The storage (1540) may be removable or non-removable, and includes magnetic disks, magnetic tapes or cassettes, CD-ROMs, CD-RWs, DVDs, or any other medium which can be used to store information and which can be accessed within the computing environment (1500). The storage (1540) stores instructions for the software (1580) implementing the described encoder/decoder and predictive transform coefficient coding techniques.

The input device(s) (1550) may be a touch input device such as a keyboard, mouse, pen, or trackball, a voice input device, a scanning device, or another device that provides input to the computing environment (1500). For audio, the input device(s) (1550) may be a sound card or similar device that accepts audio input in analog or digital form, or a CD-ROM reader that provides audio samples to the computing environment. The output device(s) (1560) may be a display, printer, speaker, CD-writer, or another device that provides output from the computing environment (1500).

The communication connection(s) (1570) enable communication over a communication medium to another computing entity. The communication medium conveys information such as computer-executable instructions, compressed audio or video information, or other data in a modulated data signal. A modulated data signal is a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not

limitation, communication media include wired or wireless techniques implemented with an electrical, optical, RF, infrared, acoustic, or other carrier.

The digital media processing techniques herein can be described in the general context of computer-readable media. Computer-readable media are any available media that can be accessed within a computing environment. By way of example, and not limitation, with the computing environment (1500), computer-readable media include memory (1520), storage (1540), communication media, and combinations of any of the above.

The digital media processing techniques herein can be described in the general context of computer-executable instructions, such as those included in program modules, being executed in a computing environment on a target real or virtual processor. Generally, program modules include routines, programs, libraries, objects, classes, components, data structures, etc. that perform particular tasks or implement particular abstract data types. The functionality of the program modules may be combined or split between program modules as desired in various embodiments. Computer-executable instructions for program modules may be executed within a local or distributed computing environment.

For the sake of presentation, the detailed description uses terms like "determine," "generate," "adjust," and "apply" to describe computer operations in a computing environment. These terms are high-level abstractions for operations performed by a computer, and should not be confused with acts performed by a human being. The actual computer operations corresponding to these terms vary depending on implementation.

In view of the many possible variations of the subject matter described herein, we claim as our invention all such embodiments as may come within the scope of the following claims and equivalents thereto.

We claim:

1. A method of encoding digital media data using a predictive coding of transform coefficients, the method comprising:

applying a transform (250) to blocks of the digital media data to produce a set of transform coefficients (400, 500, 600) for the respective blocks, wherein the set of transform coefficients of a respective block comprises a DC coefficient for the respective block;

determining (1000) whether one of at least two directionalities of the digital media data in a locality of a current block is dominant;

when one of the at least two directionalities is determined to be dominant in the digital media locality of the current block, selecting (1000) a predictor of the DC coefficient of the current block according to the determined dominant directionality;

encoding (280) the DC coefficient of the current block relative to its predictor.

2. The method of claim 1 comprising:

in a case that none of the at least two directionalities is determined to be dominant in the digital media locality of the current block, selecting the predictor of the DC coefficient of the current block according to a combination of the at least two directionalities.

3. The method of claim 1 wherein the directionalities include at least horizontal and vertical.

4. The method of claim 1 wherein said determining whether one of at least two directionalities of the digital media data in a locality of a current block is dominant comprises:

calculating directional difference metrics for the at least two directionalities as a function of DC coefficients of preceding blocks in the at least two directionalities; and

determining whether one of the at least two directionalities dominates based on the calculated directional differences.

5. The method of claim 4 wherein said calculating directional difference metrics comprises:

calculating the directional difference metrics for the at least two directionalities as a function of the DC coefficients of preceding blocks from plural channels of the digital media data in the at least two directionalities.

6. The method of claim 4 wherein the plural channels comprise a luminance channel and at least one chrominance channel.

7. The method of claim 1 wherein the transform is a second, inner stage transform applied to a block of DC coefficients from a first, outer stage transform applied to blocks in a macroblock of the digital media data.

8. The method of claim 7 wherein the set of transform coefficients of a respective block comprises plural DCAC coefficients for the respective block, the method further comprising:

when one of the at least two directionalities is determined to be dominant in the digital media locality of the current block, selecting a predictor of the DCAC coefficients of the current block according to the determined dominant directionality; and

otherwise, encoding the DCAC coefficients of the current block with no predictor.

9. The method of claim 8 further comprising, when selecting the predictor of the DCAC coefficients of the current block when one of the at least two directionalities is determined to be dominant:

determining whether the block containing the predictor selected according to the determined dominant directionality has a same quantizer as the current block; and

if not, encoding the DCAC coefficients of the current block with no predictor.

10. The method of claim 7 wherein the first, outer stage transform produces sets of transform coefficients for the blocks in a current macroblock, the sets of transform coefficients of a respective block in the current macroblock comprise plural DCAC coefficients from the first, outer stage transform, the method further comprising:

calculating directional difference metrics for the at least two directionalities as a function of DCAC coefficients produced from the second, inner stage transform on DC coefficients of the blocks in the current macroblock;

determining whether one of at least two directionalities of the digital media data for the current macroblock is dominant based upon the calculated directional difference metrics;

when one of the at least two directionalities is determined to be dominant for the current macroblock, selecting predictors of the DCAC coefficients from the first, outer stage transform of the blocks of the current macroblock unidirectionally according to the determined dominant directionality;

encoding the DCAC coefficients from the first, outer stage transform of the blocks of the current macroblock relative to their predictors.

11. The method of claim 10 wherein said calculating directional difference metrics as a function of DCAC coefficients comprises:

calculating the directional difference metrics for the at least two directionalities as a function of the DCAC coefficients produced from the second, inner stage transform on DC coefficients of the blocks in the current macroblock from plural channels of the digital media data.

12. The method of claim 11 wherein the plural channels comprise a luminance channel and at least one chrominance channel.

13. At least one computer-readable recording medium carrying a compressed bitstream encoded according to the method of claim 1.

14. A digital media encoder (200) and/or decoder (300) comprising:

a data storage buffer for storing digital media data (210) to be encoded and/or decoded;

a processor (1510) programmed to:

determine (1200) whether one of at least two directionalities of the digital media data for blocks within a current macroblock is dominant based upon DCAC coefficients produced from an inner stage transform being applied to a block of DC coefficients produced from an outer stage transform being applied to the blocks within the current macroblock;

when one of the at least two directionalities is determined to be dominant for the blocks within the current macroblock, select predictors (1300, 1400) of the DCAC coefficients produced from the outer stage transform of the blocks within the current macroblock unidirectionally according to the determined dominant directionality; and

encode (280) and/or decode (320) the DCAC coefficients produced from the outer stage transform of the blocks within the current macroblock relative to their predictors.

15. The digital media encoder and/or decoder of claim 14, wherein the processor is further programmed to:

determine whether one of at least two directionalities of the digital media data in a locality of the current macroblock is dominant based upon DC coefficients of preceding neighboring macroblocks, where the DC coefficient of a respective preceding neighboring macroblock is produced from the inner stage transform being applied to a block of DC coefficients produced from the outer stage transform being applied to blocks within that respective preceding neighboring macroblock;

when one of the at least two directionalities is determined to be dominant in the locality of the current macroblock, select a predictor of the DC coefficient of the current macroblock according to the determined dominant directionality, where the DC coefficient of the current macroblock is produced from the inner stage transform being applied to a block of DC coefficients produced from the outer stage transform being applied to blocks within the current macroblock; and

encode and/or decode the DC coefficient of the current macroblock relative to its predictor.

16. The digital media encoder and/or decoder of claim 15, wherein the processor is further programmed to:

when one of the at least two directionalities is determined to be dominant in the locality of the current macroblock and except when the quantizers of the current macroblock differs from that of a neighboring macroblock in a direction according to the determined dominant directionality, select predictors of the DCAC coefficients of the current macroblock according to the determined dominant

directionality, where the DCAC coefficients of the current macroblock are produced from the inner stage transform being applied to a block of DC coefficients produced from the outer stage transform being applied to blocks within the current macroblock; and

encode and/or decode the DCAC coefficients of the current macroblock relative to their predictors.

17. The digital media encoder and/or decoder of claim 16, wherein said determination of whether one of at least two directionalities is dominant for blocks within a current macroblock and for the locality of the macroblock is based upon the respective coefficients in plural channels of the digital media data, the plural channels comprising a luminance channel and at least one chrominance channel.

18. At least one computer-readable recording medium carrying a computer-executable digital media processing program thereon for performing a method of processing digital media data, the method comprising:

calculating (1200) directional difference metrics for at least two directionalities of the digital media data for blocks within a current macroblock based upon DCAC coefficients produced from an inner stage transform being applied to a block of DC coefficients produced from an outer stage transform being applied to the blocks within the current macroblock;

determining (1200) based on the directional difference metrics whether one of the at least two directionalities of the digital media data for blocks within a current macroblock is dominant;

if one of the at least two directionalities is determined to be dominant for the blocks within the current macroblock, selecting predictors (1300, 1400) of the DCAC coefficients produced from the outer stage transform of the blocks within the current macroblock unidirectionally according to the determined dominant directionality; and

encoding (280)/decoding (320) the DCAC coefficients produced from the outer stage transform of the blocks within the current macroblock relative to their predictors.

19. The at least one computer-readable recording medium of claim 18, wherein the method of processing digital media data further comprises:

calculating directional difference metrics for at least two directionalities of the digital media data in a locality of the current macroblock based upon DC coefficients of preceding neighboring macroblocks, where the DC coefficient of a respective macroblock is produced from the inner stage transform being applied to a block of DC coefficients produced from the outer stage transform being applied to blocks within that respective macroblock;

determining based on the calculated directional difference metrics whether one of the at least two directionalities of the digital media data in the locality of the current macroblock is dominant;

if one of the at least two directionalities is determined to be dominant in the locality of the current macroblock, selecting a predictor of the DC coefficient of the current macroblock according to the determined dominant directionality; and

encoding/decoding the DC coefficient of the current macroblock relative to its predictor.

20. The at least one computer-readable recording medium of claim 19, wherein the method of processing digital media data further comprises:

if one of the at least two directionalities is determined to be dominant in the locality of the current macroblock and except when the quantizers of the current macroblock differs from that of a neighboring macroblock in a direction according to the determined dominant directionality, selecting predictors of the DCAC coefficients of the current macroblock according to the determined dominant directionality, where the DCAC coefficients of the current macroblock are produced from the inner stage transform being applied to a block of DC coefficients produced from the outer stage transform being applied to blocks within the current macroblock; and

encoding/decoding the DCAC coefficients of the current macroblock relative to their predictors.