




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification ⁵ : H04N 7/13, 7/137</p>	<p>A1</p>	<p>(11) International Publication Number: WO 92/12598 (43) International Publication Date: 23 July 1992 (23.07.92)</p>
<p>(21) International Application Number: PCT/GB92/00066 (22) International Filing Date: 13 January 1992 (13.01.92) (30) Priority data: 9100593.4 11 January 1991 (11.01.91) GB (71) Applicant: SONY BROADCAST & COMMUNICATIONS LIMITED [GB/GB]; Jays Close, Viables, Basingstoke, Hampshire RG22 4SB (GB). (72) Inventors: STONE, Jonathan, James ; 7 Forest Drive, Whitley Wood Lane, Reading, Berkshire RG2 8PQ (GB). HURLEY, Terence, Ralph ; 57 Bartlemy Road, Newbury, Berkshire RG14 6LA (GB).</p>		<p>(74) Agent: COTTER, Ivan, John; D Young & Co., 10 Staple Inn, London WC1V 7RD (GB). (81) Designated State: JP. Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>

(54) Title: COMPRESSION OF VIDEO SIGNALS



(57) Abstract

A digital video signal (component or composite) is compressed by spatial sub-bands filtering (12A or 12B) to form data sets constituting respective sub-bands (Fig.4) of the two-dimensional spatial frequency domain. The data sets for a field or frame of the signal are stored (28). A first sequencer (29A or 29B) controls writing, in accordance with a desired sequence, of the stored data to a quantiser (14-Fig.8) in which they are quantised in accordance with respective values (Fig.7), those values being such that the amount of quantisation of a data set constituting a sub-band to which dc luminance information of the signal is at least predominantly confined (and, in the case of a composite signal, also at least two of the data sets constituting sub-bands to which dc chrominance information of the signal is at least predominantly confined) is less than the average of the amounts of quantisation of the remaining data sets. The quantised data sets are then entropy encoded in an entropy encoder (16-Fig. 12) which has a first coding portion (64, 72) for coding quantised data representative of dc luminance information and a second coding portion (62) for coding quantised data representative of ac luminance information. A second sequencer (58), which may in fact be the same sequencer as the first sequencer (29A or 29B), controls the operation of the quantiser (14) so that each datum (sample) written thereto is appropriately quantised, and controls the operation of the entropy encoder (16) so that each quantised sample is directed to the appropriate one of the first (64, 72) and second (62) coding portions.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	ES	Spain	MG	Madagascar
AU	Australia	FI	Finland	ML	Mali
BB	Barbados	FR	France	MN	Mongolia
BE	Belgium	GA	Gabon	MR	Mauritania
BF	Burkina Faso	GB	United Kingdom	MW	Malawi
BG	Bulgaria	GN	Guinea	NL	Netherlands
BJ	Benin	GR	Greece	NO	Norway
BR	Brazil	HU	Hungary	PL	Poland
CA	Canada	IT	Italy	RO	Romania
CF	Central African Republic	JP	Japan	RU	Russian Federation
CG	Congo	KP	Democratic People's Republic of Korea	SD	Sudan
CH	Switzerland	KR	Republic of Korea	SE	Sweden
CI	Côte d'Ivoire	LI	Liechtenstein	SN	Senegal
CM	Cameroon	LK	Sri Lanka	SU	Soviet Union
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
DE	Germany	MC	Monaco	TC	Togo
DK	Denmark			US	United States of America

COMPRESSION OF VIDEO SIGNALS

This invention relates to the compression of video signals.

Compression of video signals on an intra-image basis (for
5 example, compression on an intra-field or intra-frame basis) makes use
of the redundancy present in pictures or images represented by the
signals to reduce the amount of information needed to represent the
pictures or images. The compression can be used to reduce bandwidth,
in the case of transmission of a video signal, or to reduce storage
10 capacity, in the case of storage of a video signal.

Intra-image compression can, as is known, be effected in the time
domain by the use of differential pulse code modulation, in which a
predictor is used to predict the values of samples representing pixels
based on previous pixel values. Since the image pixels are highly
15 correlated, the prediction is accurate and results in a small and
uncorrelated error (that is, a difference between the predicted and
actual values). The error samples are encoded and, since they can be
encoded using fewer bits than the samples representing the original
pixels, compression can be achieved.

20 Figure 1 of the accompanying drawings shows a known apparatus or
system for effecting intra-image compression of a video signal in the
two-dimensional spatial frequency domain. A video signal, which is in
digital form and comprises successive multi-bit (for example 8-bit)
samples or words each representing a respective pixel of an scanned
25 image or picture, is applied via an input 10 to a decorrelator 12. A
decorrelated version of the video signal is outputted by the
decorrelator 12 to a quantiser 14 and then to an entropy encoder 16,
which together compress the decorrelated version of the video signal
outputted by the decorrelator 12 to produce a compressed signal on an
30 output 18. The compressed signal can then be transmitted or stored.
(Note that, although the decorrelator 12, quantiser 14 and entropy
encoder 16 are shown for clarity as being separate items, they may in
practice be embodied in an at least partially combined form.) After
transmission or storage, the compressed signal can be restored
35 substantially to its original form by expansion by way of entropy
decoding, dequantising and correlation operations which employ
parameters converse to those used for decorrelation, quantisation and

entropy encoding, respectively, upon compression.

The operation of decorrelation performed in the decorrelator 12 relies upon the fact that neighbouring pixels of an image are highly correlated, whereby processing an image (for example, a field or frame of a video signal) to form decorrelated signal portions representing different components of the image in the two-dimensional spatial frequency domain enables a reduction in the amount of information needed to represent the image. Specifically, the decorrelated signal portions represent different spatial frequency components of the image to which the human psychovisual system has respective different sensitivities. The different decorrelated signal portions are subjected to different degrees of quantisation in the quantiser 14, the degree of quantisation for each signal portion depending upon the sensitivity of the human psychovisual system to the information in that portion. That is, each of the decorrelated signals is quantised in accordance with its relative importance to the human psychovisual system. This selective quantisation operation, which is a lossy operation in that it involves deliberate discarding of some frequency data considered to be redundant or of little importance to adequate perception of the image by the human psychovisual system, in itself enables some signal compression to be achieved. The quantiser 14 enables compression to be achieved in two ways: it reduces the number of levels to which the data inputted to it can be assigned, and it increases the probability of runs of zero value samples on the data it outputs. Note that, in video signal compression apparatus described in detail below, the ability to achieve signal compression provided by the operation of the quantiser 14 is not used to produce a bit (data) rate reduction in the quantiser itself. Instead, in that case, the ability to achieve signal compression provided by the operation of the quantiser is carried into effect in the entropy encoder 16 in that the reduction in information content achieved in the quantiser 14 enables a consequential bit (data) rate reduction to be achieved in the entropy encoder.

Further (non-lossy) compression, and bit (data) rate reduction, is provided in the entropy encoder 16 in which, in known manner, using for example variable length coding, the data produced by the quantiser 14 is encoded in such a manner that more probable (more frequently occurring) items of data produce shorter output bit sequences than less

probable (less frequently occurring) ones. In this regard, the decorrelation operation has the effect of changing the probability distribution of the occurrence of any particular signal level, which is substantially the same as between the different possible levels before decorrelation, into a form in which in which it is much more probable that certain levels will occur than others.

The compression/coding system or apparatus as shown in Figure 1 can be embodied in a variety of ways, using different forms of decorrelation. An increasingly popular form of implementation makes use of so-called transform coding, and in particular the form of transform known as the discrete cosine transform (DCT). (The use of DCT for decorrelation is in fact prescribed in a version of the compression system of Figure 1 described in a proposed standard prepared by JPEG (Joint Photographic Experts Group) and currently under review by the ISO (International Standards Organisation).) According to the transform technique of decorrelation, the signal is subjected to a linear transform (decorrelation) operation prior to quantisation and encoding. A disadvantage of the transform technique is that, although the whole image (for example, a whole field) should be transformed, this is impractical in view of the amount of data involved. The image (field) thus has to be divided into blocks (for example, of 8 x 8 samples representing respective pixels), each of which is transformed. That is, transform coding is complex and can be used on a block-by-block basis only.

A recently proposed approach to compression/coding in the frequency domain is that of sub-band coding. In this approach, the decorrelator 12 in the system of Figure 1 would comprise a spatial (two-dimensional) sub-band filtering arrangement (described in fuller detail below) which divides the input video signal into a plurality of uncorrelated sub-bands each containing the spatial frequency content of the image in a respective one of a plurality of areas of a two-dimensional frequency plane of the image, the sub-bands then being selectively quantised by the quantiser 14 in accordance with their positions in the sensitivity spectrum of the human psychovisual system. That is, decorrelation is achieved in this case by putting the energy of the overall image into different sub-bands of the two-dimensional spatial frequency domain. Sub-band filtering is believed to provide

better decorrelation than the transform approach. Also, unlike the transform technique, there is no restriction to operation on a block-by-block basis: the sub-band filtering can be applied directly to the video signal.

5 The proposed JPEG standard mentioned above requires that data decorrelated by means of a block-by-block DCT transformation operation be quantised and then entropy encoded. The data is required by the standard to be quantised and encoded in a particular sequence dictated by the order in which it is outputted by the decorrelator, which order
10 is in turn dictated by the way in which the input digital video signal is divided into blocks for transformation. The sequence is such that each successive one of a sequence of (for example) 8 x 8 arrays (blocks) of data each resulting from transformation of an 8 x 8 block of data of the input signal is, after quantisation, entropy encoded in
15 such a manner that one of the 64 data in the array is entropy encoded in a manner different than the other 63 data in the array. This requires switching of the entropy encoding at the block frequency, that is once every 64 data (samples). Moreover, the sequence is such that the 64 data in each block (in turn) are quantised and encoded in a
20 particular order specified in the standard. In the case of sub-band filtering, the data to be quantised and encoded is of a very different format than that obtained in the case of block-by block DCT transformation. On the face of it, the JPEG standard approach not only excludes the use of sub-band filtering, but, on the face of it, is
25 wholly incompatible with the use of sub-band filtering, which does not require a block-by-block approach. This is unfortunate because, as mentioned above, sub-band filtering is believed superior to block-by block DCT transformation.

According to one aspect of the invention there is provided a
30 method of compressing a video signal, the method comprising:

effecting spatial two-dimensional sub-band filtering of a digital video signal to form a plurality of data sets constituting respective sub-bands of the two-dimensional spatial frequency domain;

storing the data sets for a field or frame of the video signal;

35 quantising the stored data sets in accordance with respective values, said values being such that the amount of quantisation of one of the data sets constituting a sub-band to which dc luminance

information of the signal is at least predominantly confined is less than the average of the amounts of quantisation of the remaining data sets;

5 writing the stored data to a quantiser, for carrying out the quantising, in a desired sequence;

controlling the quantising in accordance with said desired sequence such that each datum written to the quantiser is appropriately quantised;

10 entropy encoding the quantised data sets such that quantised data representative of dc luminance information is coded by a first coding technique and quantised data representative of ac luminance information is coded by a second coding technique; and

15 controlling the entropy encoding in accordance with said desired sequence such that each quantised datum is subjected to the appropriate one of the first and second coding techniques.

The flexibility resulting from the fact that the data sets for a field or frame are stored and then written to the quantiser in a desired sequence, and a realisation that, while their respective formats are very different, the data obtained in the respective cases of sub-band filtering and block-by block DCT transformation have sufficient resemblance (at least in terms of information content) to one another that the former can be rearranged to resemble the latter, to a greater or lesser degree, means that the apparent incompatibility between sub-band filtering and the JPEG standard is not as great as at first appears. On the contrary, the same general approach as stipulated in the JPEG standard can, to a greater or lesser degree, be adopted.

25 The extent to which the JPEG standard is followed depends upon the application. If, for example, strict compliance with the standard is not required, for example if a proprietary piece of recording equipment is to be produced in which the designer can specify at his discretion the way in which the signal is to be compressed and subsequently expanded, there is no need to follow the above-described sequence of data quantisation and entropy encoding laid down in the standard. The designer can use a sequence that seems best in the circumstances.

35 Thus, according to one form of implementation of the invention

described hereinbelow:

said desired sequence in which the stored data is written to the quantiser is such that

(i) all the data of said one of the stored data sets
5 constituting the sub-band to which the dc luminance information is at least predominantly confined are quantised,

(ii) after which, upon each occurrence of an operation carried out for a number of times equal to the number of data in each stored data set, those data corresponding to a respective one of the spatial
10 positions in each of the remaining stored data sets are quantised in a predetermined order;

the quantising is controlled such that all the data written to the quantiser in step (i) are quantised in the same amount, namely in the amount appropriate to the sub-band to which the dc luminance
15 information is at least predominantly confined, and each of the data written to the quantiser in each of the operations of step (ii) is quantised in an amount appropriate to the data set of which it forms a part; and

the entropy encoding is controlled such that all the data
20 quantised upon being written to the quantiser in step (i) are subjected to the first coding technique and all the data quantised upon being written to the quantiser in all of the operations of step (ii) are subjected to the second coding technique.

The use of the two steps or stages set out at (i) and (ii) for
25 quantising the data means that the format of the data to be quantised is very different than in the case of the JPEG (DCT) standard; and that the entropy encoding has to be switched at the field or frame frequency rather than at the much higher frequency (block frequency) used in the case of the JPEG (DCT) standard. However, this form of sequencing is
30 believed superior to the JPEG sequence at least in some cases, in that it groups the dc and ac information together rather than intermingles it; and the invention enables this form of sequencing to be used if desired.

Even if the above form of sequencing is used, the ac data can be
35 treated in a generally similar manner to that specified in the JPEG standard. Thus, for example, said predetermined order may comprise successive groups of the remaining stored data sets in a sequence, as

between the groups, of sub-bands containing ac luminance information of increasing spatial frequency. In fact, to accomplish a form of zig-zag sequencing generally similar to that specified in the standard, said groups may comprise legs of a zig-zag pattern connecting the data of the same spatial position in the different stored data sets.

According to another form of implementation of the invention described hereinbelow:

said desired sequence in which the stored data is written to the quantiser is such that, upon each occurrence of an operation carried out for a number of times equal to the number of data in each stored data set, those data corresponding to a respective one of the spatial positions in each of the stored data sets are quantised in a predetermined order;

the quantising is controlled such that each of the data written to the quantiser in each of said operations is quantised in an amount appropriate to the data set of which it forms a part; and

the entropy encoding is controlled such that, for each of said operations, that one of the quantised data forming part of the data set constituting the sub-band to which the dc luminance information is at least predominantly confined is subjected to the first coding technique and all of the other data are subjected to the second coding technique.

In this case, the sequencing is very similar to that specified in the JPEG standard, in particular if said predetermined order comprises successive groups of the stored data sets in a sequence, as between the groups, of sub-bands containing ac luminance information of increasing spatial frequency; and if said groups comprise legs of a zig-zag pattern connecting the data of the same spatial position in the different stored data sets.

Thus, as explained in more detail below, the format of the data to be quantised is very similar to that in the case of the JPEG (DCT) standard, which has the advantage that quantisation can be sequenced in a very similar or even identical way to that used in the case of the standard. Thus, it may be possible to use an "off the shelf" quantiser chip or assembly intended for use in a JPEG compression apparatus (possibly with changes in the quantisation values). Also, in this case, the entropy encoding has to be switched at the frequency of carrying out the operations in which those data

corresponding to a respective one of the spatial positions in each of the stored data sets are quantised, that is at a frequency determined by the number of data sets (sub-bands), rather than at the field or frame frequency. If, as in the case of embodiments of the invention
5 described below, the number of data sets (sub-bands) is the same ($8 \times 8 = 64$) as the number of samples per block as specified in the JPEG standard, the frequency of switching the entropy encoding is the same as the frequency (block frequency) used in the case of the JPEG (DCT) standard. Thus, it may be possible to use an "off the shelf" entropy
10 encoder chip or assembly intended for use in a JPEG compression apparatus.

As is well known, a colour video signal can be in component or composite form. A component colour video signal comprises three separate signals which together represent the totality of the video
15 information. The three separate signals may, for example, be a luminance signal and two colour difference signals (Y, Cr, Cb) or three signals each representing a respective colour (R, G, B). A composite colour video signal, on the other hand, is a single signal comprising all the luminance and chrominance (colour) information.

20 Previously proposed colour video signal compression systems as described above all operate on component signals only. That is, taking the example of the system of Figure 1, three separate systems as shown in Figure 1 are needed, one for each of the three components. Also, if the signal is in composite form, there is a need for means to convert
25 it into component form prior to compression. Further, three expansion systems are needed to convert the transmitted or stored signals back to their original form, together with (if appropriate) means to convert the component signals back into composite form. The need to process the video signal in component form thus involves the expense and
30 inconvenience of considerable hardware replication.

While the invention is applicable in the case of component (or monochrome) video signals, a preferred feature of the invention is that it can be used also to compress composite colour video signals. This preferred feature takes advantage of a realisation by the inventors
35 that, due to the way in which luminance and chrominance information are combined in conventional broadcast standard (for example, NTSC and PAL) composite colour video signals, such a signal can be spatially sub-band

filtered such that the chrominance information can be (as is explained in detail below) concentrated in a certain area of the two-dimensional spatial frequency domain (that is, in certain of the sub-bands), whereby, if the data sets to which the dc chrominance information and dc luminance information are at least predominantly confined are quantised more lightly than the other data sets (which contain wholly or largely only the ac luminance information) are on average quantised, then since the dc information is more important to satisfactory appreciation of the image by the human psychovisual system than the ac luminance information it is in fact (surprisingly) possible satisfactorily to compress a composite colour video signal directly, that is without first converting it to component form and compressing each component individually.

Another advantage of the sub-band approach to signal decorrelation is that (as is also the case for the DCT approach) the sub-band approach is separable between two orthogonal spatial directions. Thus, the digital video signal is preferably separately spatially sub-band filtered in respective orthogonal spatial directions. The separable approach simplifies design.

A further advantage of the separable approach is that it enables the method to be performed such that a field or frame of the digital video signal is sub-band filtered in one of the orthogonal directions in a first one-dimensional sub-band filter arrangement, stored, and then transposed and sub-band filtered in the other of the orthogonal directions in a second one-dimensional sub-band filter arrangement which is of substantially the same construction as the first one-dimensional sub-band filter arrangement.

To further minimise hardware requirements, according to a preferred form of the method, in a first stage of the filtering in each of the orthogonal directions, the digital video signal is subjected to low pass filtering followed by decimation by two and also to high pass filtering followed by decimation by two, thereby to produce two intermediate outputs, and, in at least one subsequent stage of the filtering in each of the orthogonal directions, each of the intermediate outputs produced in the previous stage is subjected to low pass filtering followed by decimation by two and also to high pass

filtering followed by decimation by two. The use of such a "tree" or "hierarchical" structure of so-called quadrature mirror filters (QMFs), as opposed to the alternative possibility of a bank of filters operating in parallel, is preferred in that it reduces hardware requirements and also enables better reconstruction of the signal to be achieved on subsequent expansion. In this regard, the aliasing that will of necessity be introduced in the hierarchical QMF filtering or decomposition during the course of compression can in principle be removed completely during the course of a converse composition operation performed upon expansion (after transmission or storage of the compressed signal).

According to an alternative to the storage and transposition approach, the digital video signal may be sub-band filtered in a first one-dimensional sub-band filter arrangement configured to sub-band filter the signal in one of the orthogonal directions, and then sub-band filtered in a second one-dimensional sub-band filter arrangement configured to sub-band filter the signal in the other of the orthogonal directions. This approach may be preferable if the filter structure is constructed as a unit on silicon rather than by wiring together separate integrated circuits.

The location in the two-dimensional spatial frequency domain of the dc chrominance information is determined by the relationship between the frequency at which an analog composite colour video signal has been sampled to form the digital composite colour video signal, and the frequency of a colour sub-carrier frequency of the composite colour video signal. According to a preferred form of the method, the digital composite colour video signal has been formed by sampling an analog composite colour video signal at a frequency equal to four times the frequency of a colour sub-carrier frequency of the composite colour video signal.

According to the preferred form of the method disclosed below, the plurality of sub-bands make up a square array in the two-dimensional spatial frequency domain. The array may, for example, be a 4 x 4 array or an 8 x 8 array. However, it is equally feasible, and may in some cases be appropriate, to have different numbers of sub-bands in the two orthogonal directions, that is to employ a non-square, rectangular array. For instance, the sub-bands may make

up, in the two-dimensional spatial frequency domain, a rectangular array having a dimension of 8 in the direction of scanning of the video signal and a dimension of 4 in the direction orthogonal thereto.

According to another aspect of the invention there is provided apparatus for compressing a video signal, the apparatus comprising:

a spatial two-dimensional sub-band filtering arrangement operative to filter a digital video signal to form a plurality of data sets constituting respective sub-bands of the two-dimensional spatial frequency domain;

a store for storing the data sets for a field or frame of the video signal;

a quantiser operative to quantise the stored data sets in accordance with respective values, said values being such that the amount of quantisation of one of the data sets constituting a sub-band to which dc luminance information of the signal is at least predominantly confined is less than the average of the amounts of quantisation of the remaining data sets;

an entropy encoder operative to encode the quantised data sets, the entropy encoder comprising a first coding portion for coding quantised data representative of dc luminance information and a second coding portion for coding quantised data representative of ac luminance information; and

sequencing means operative to control writing of the stored data from the store to the quantiser in a desired sequence, to control the operation of the quantiser in accordance with said desired sequence such that each datum written thereto is appropriately quantised, and to control the operation of the entropy encoder in accordance with said desired sequence such that each quantised datum is directed to the appropriate one of the first and second coding portions.

The invention will now be further described, by way of illustrative and non-limiting example, with reference to the accompanying drawings, in which like references indicate like items throughout, and in which:

Figure 1 shows a video signal compression apparatus or system for achieving intra-image compression of a video signal in the frequency domain;

Figure 2 is a block diagram of one form of implementation of a decorrelator, in the form of a sub-band filtering arrangement, for use in the video signal compression apparatus;

Figure 3 is a detailed block diagram of a horizontal filter arrangement forming part of the sub-band filtering arrangement shown in Figure 2;

Figure 4 shows a sub-band filtered field of a video signal (luminance only) on a two-dimensional frequency plane;

Figure 5 is a block diagram of another form of implementation of a decorrelator, in the form of a sub-band filtering arrangement, for use in the video signal compression apparatus;

Figure 6 is a graph representing the response of the human psychovisual system to different spatial frequencies;

Figure 7 represents a quantisation matrix that would be used in a quantiser of the video signal compression apparatus if a sub-band filtered component (luminance) video signal were being processed in the quantiser, and shows also respective modifications to be made if, instead, a sub-band filtered composite video signal (NTSC or PAL) were being processed in the quantiser;

Figure 8 is a block diagram of the quantiser;

Figure 9 shows part of Figure 4 on an enlarged scale, and is used to explain the operation of the quantiser;

Figure 10 is a diagram showing how zig-zag scanning of the ac sub-bands is carried out in the quantiser;

Figure 11 shows the format of quantised data emerging from the quantiser for ac sub-bands;

Figure 12 is a block diagram of an entropy encoder forming part of the video signal compression apparatus;

Figure 13 is a representation of the contents of a fixed length code look-up table forming part of the entropy encoder;

Figure 14 shows a sub-band filtered field of an NTSC composite colour video signal, sampled at four times its colour sub-carrier frequency, on the two-dimensional frequency plane;

Figure 15 is a graph showing the two-dimensional frequency content of a field of an analog NTSC composite colour video signal;

Figure 16 shows a frame of an NTSC composite colour video signal, sampled at four times the colour sub-carrier frequency, on the

two-dimensional frequency plane;

Figure 17 is a view corresponding to Figure 4, but showing on the two-dimensional frequency plane both the sub-band filtered field of an NTSC composite colour video signal, and a sub-band filtered field of a PAL composite colour video signal, each sampled at four times its colour sub-carrier frequency;

Figure 18 shows how samples making up a field or frame of a digital video signal are divided into blocks to be processed by a linear transform, for example a DCT;

Figure 19 is a block diagram of a linear transform decorrelator;

Figure 20 is a block diagram of a linear transform decorrelator, like that of Figure 19 but modified to emulate decorrelation performed by sub-band filtering;

Figure 21 shows an 8 x 8 block of coefficients (samples) outputted by a linear transform circuit of the decorrelator of Figure 20; and

Figure 22 shows how 8 x 8 blocks of coefficients outputted by the linear transform circuit are written into a store of the decorrelator of Figure 20.

A method and apparatus for compressing a digital video signal will now be described with reference to the drawings. The basic construction of the apparatus is in accordance with Figure 1 (described above). The decorrelator 12 of the present apparatus is constituted by a sub-band filtering arrangement which, according to one form of implementation as shown in outline form at 12A in Figure 2, comprises a horizontal filter arrangement 20A, an intermediate field store 22, a transpose sequencer (address generator) 24, a vertical filter arrangement 26A, an output field store (FS) 28 and an output sequencer (address generator) 29A. As explained above, sub-band filtering can be effected on a separable basis. Thus, in Figure 2, filtering in the two orthogonal image directions, namely the horizontal direction (the direction of image scanning in the case of conventional video) and the vertical direction, is effected entirely independently and separately of one another by respective one-dimensional filtering operations performed in the horizontal and vertical filter arrangements 20A and 26A, respectively.

The horizontal filter arrangement 20A and vertical filter

arrangement 26A can be of substantially the same construction as one another. Thus, the construction of the horizontal filter arrangement 20A only will be described in detail.

It will be assumed that the filtering is to achieve 8 sub-bands
5 in each of the horizontal and vertical directions, that is to say that a square array of 64 (8 x 8) sub-bands is to be produced. It will further be assumed that the 64 sub-bands are (as is preferred) to be of equal extent to one another.

The horizontal filter arrangement 20A is preferably of a tree or
10 hierarchical structure as shown in Figure 3, comprising three successive filter stages 30, 32 and 34.

The first stage 30 comprises a low pass filter (LPF) 36 and a high pass filter (HPF) 38, each of which is followed by a respective decimator (DEC) 40. The LPF filter 36, HPF filter 38 and the
15 decimators 40 together make up a quadrature mirror filter (QMF). Each of the filters 36 and 38 can be a finite impulse response (FIR) filter of conventional form. In use, a line of a field of the input digital video signal is applied, sample-by-sample, to the first stage 30, to be low pass filtered and high pass filtered by the LPF 36 and HPF 38,
20 respectively. Thus, the LPF 36 and HPF 38 produce outputs comprising low pass filtered and high pass filtered versions of the input line, respectively, the outputs representing the spatial frequency content of the line in the upper and lower halves of the horizontal spatial frequency range. That is, the first stage 30 divides the input line
25 into two sub-bands in the horizontal direction. The decimators 40 decimate (sub-sample) the respective outputs by a factor of two, whereby the total number of samples outputted by the decimators 40 (together) is the same as the total number of samples in the line.

The second stage 32 is of similar construction to the first stage
30 30, except that there are two QMFs each as in the first stage and the output from each of the decimators 40 of the first stage is passed as an input to a respective one of the two QMFs. Thus, the second stage 32 produces four outputs representing the spatial frequency content of the line in four equal quarters of the horizontal spatial frequency
35 range. That is, the second stage 32 further divides the two sub-bands, into which the input line was divided in the first stage 30, into four sub-bands in the horizontal direction. The four decimators of the

second stage 32 decimate (sub-sample) the respective outputs by a factor of two, whereby the total number of samples outputted by the decimators of the second stage (together) is the same as the total number of samples in the line.

5 The third stage 34 is of similar construction to the first stage 30, except that there are four QMFs each as in the first stage and the output from each of the four decimators of the second stage 32 is passed as an input to a respective one of the four QMFs. Thus, the third stage 34 produces eight outputs representing the spatial
10 frequency content of the line in eight equal one-eighths of the horizontal spatial frequency range. That is, the third stage 34 divides the four sub-bands into which the input line was previously divided into the required eight sub-bands in the horizontal direction. The eight decimators of the third stage 34 decimate (sub-sample) the
15 respective outputs by a factor of two, whereby the total number of samples outputted by the decimators of the third stage (together) is the same as the total number of samples in the line.

The eight outputs of the third stage 34, that is of the horizontal filter arrangement 20A, are passed to the intermediate field
20 store 22 and stored at positions corresponding to respective one-eighths of a first line thereof. The above process of horizontal filtering is then repeated for all the other lines of the field of the input digital video signal. This results in the intermediate field store 22 containing a version of the field of the input digital video
25 signal that has been filtered into eight sub-bands in the horizontal direction (only). Each line of the field stored in the intermediate field store 22 is divided into eight portions each containing the horizontal spatial frequency information in a respective one of eight sub-bands of the horizontal spatial frequency range of the image that
30 the original field represented. Thus, the horizontally filtered field stored in the intermediate field store 22 can be considered to be divided into eight columns.

Referring back to Figure 2, the horizontally filtered field stored in the intermediate field store 22 is then fed (under the
35 control of the transpose sequencer 24) into the vertical filter arrangement 26A, in which it is filtered into eight sub-bands in the vertical direction in similar manner to that in which filtering into

eight sub-bands in the horizontal direction was achieved in the horizontal filter arrangement 20A. The horizontally and vertically filtered field is fed on a line-by-line basis into the output field store 28 to be passed from there to the quantiser 14. The store 28 can be considered to have been partitioned into an array of 64 (8 x 8) storage regions, in each of which a respective one of the 64 sub-bands is stored. Thus, successive fields of the input digital video signal are sub-band filtered and passed, duly filtered, to the quantiser 14 after a delay of two field intervals.

10 The transpose sequencer 24 produces read addresses for the intermediate field store 22, to control reading of the contents thereof into the vertical filter arrangement 26A, as follows. As will be recalled, the signal as stored in the intermediate field store 22 comprises the lines of the original field, each divided horizontally
15 into eight sub-bands. That is, the signal as stored in the intermediate field store 22 can, as mentioned above, be considered to comprise eight columns. To enable the signal stored in the intermediate field store 22 to be vertically filtered by hardware of the same construction (the vertical filter arrangement 26A) used to
20 horizontally filter it, it must be transposed, that is rotated through 90 degrees, as it is read to the vertical filter arrangement 26A, so that it comprises eight rows (as opposed to columns). The transpose sequencer 24 addresses the intermediate field store 22 in such a manner as to accomplish this.

25 The horizontally and vertically filtered field stored in the output field store 28, which has been sub-band filtered by a factor of eight in both directions, can thus be considered as having been divided into eight rows and eight columns, that is into an 8 x 8 sub-band array. The horizontally and vertically sub-band filtered
30 field, as stored in the output field store 28 of the sub-band filtering arrangement 12 ready for quantisation, can be represented (subject to the qualification mentioned below concerning sub-band scrambling) on a two-dimensional frequency plane as shown in Figure 4. In conventional manner for considering image (two-dimensional) signals, frequency is
35 represented in normalised form in Figure 4, the symbol π being equivalent to half the Nyquist limit sampling frequency. For the time being, it is assumed that the input digital video signal is a component

(luminance) signal, or even a monochrome signal, rather than a composite signal. Thus, the 64 sub-bands comprise a single sub-band, referred to hereinafter as the dc (zero spatial frequency) sub-band, which contains most or all of the dc information image intensity data, namely the sub-band (shown shaded) in the upper left hand corner of Figure 4, together with 63 ac sub-bands which contain edge data, that is components of the two-dimensional frequency spectrum of the image in respective sub-bands higher than dc (zero spatial frequency). In this regard, if the filtered signal in the output field store 28 were viewed on a monitor, it would be intelligible. Thus, a very heavily filtered version of the original signal would be seen in the upper left hand corner picture area (dc sub-band) and higher frequency components could be observed in the other 63 picture areas (ac sub-bands).

The sub-band filtering arrangement structure described above with reference to Figure 3 (unlike an alternative arrangement described below with reference to Figure 5), because of its hierarchical QMF structure, "scrambles" the order or sequence of the sub-bands. That is, due to a frequency inversion that takes place in each of the QMFs, if a field of the filtered signal in the output field store 28 were viewed on a monitor, there would not be a one-to-one correspondence between the field as viewed and the showing of Figure 4. Thus, while the dc sub-band would remain in the upper left-hand corner, the frequency plane locations of the 63 ac sub-bands would be different from (that is, scrambled with respect to) their locations in Figure 4. The locations would of course be the same for successive fields and can readily be determined from the structure of Figure 3. In other words, while each of the 64 storage regions into which the store 28 is partitioned stores a respective one of the 64 sub-bands, the relative positioning of the 63 storage regions containing the ac sub-bands is scrambled (in a known manner) with respect to the relative positioning of the ac sub-bands as shown in Figure 4.

In order that the scrambled locations of the 63 ac sub-bands are descrambled (that is, put into the pattern shown in Figure 4) before the sub-band filtered signal is passed to the quantiser 14, the output sequencer 29A (which can be located, as shown, in the sub-band filtering arrangement 12A, though it could be located elsewhere, for example in the quantiser 14), which is connected to the output field

store 28 to produce read addresses therefor to cause the data therein to be read out to the quantiser 14, is so designed that the data is read out in a descrambled manner, that is in such a manner that the sub-bands as supplied to the quantiser conform to Figure 4. (The operation of the sequencer 29A in this regard is described in more detail below with reference to Figures 9 and 10.)

Figure 5 shows at 12B a form of implementation of the sub-band filtering arrangement which can be used instead of that (12A) described above with reference to Figures 2 and 3. The sub-band filtering arrangement 12B comprises a horizontal filter arrangement 20B, a vertical filter arrangement 26B, an output field store 28, and an output sequencer 29B. As in the case of the sub-band filtering arrangement 12A of Figures 2 and 3, filtering in the horizontal and vertical directions is in this case also effected entirely separately of one another, namely by respective one-dimensional filtering operations performed in the horizontal and vertical filter arrangements 20B and 26B, respectively.

The horizontal filter arrangement 20B is of a conventional FIR structure, comprising a chain of an appropriate number of one-sample delay elements 40 tapped off to multipliers 42 (supplied with respective appropriate weighting coefficients WC) whose output signals are summed by adders 44 to produce a horizontally sub-band filtered output signal at the output 46 of the final adder. Similarly, the vertical filter arrangement 26B is of a conventional FIR structure, comprising a chain of an appropriate number of one-line delay elements 46 tapped off to multipliers 42 (supplied with respective appropriate weighting coefficients WC) whose output signals are summed by adders 44 to produce a horizontally and vertically sub-band filtered output signal at the output 48 of the final adder, which signal is stored on a field-by-field basis in the output field store 28. The output sequencer 29B (which can be located, as shown, in the sub-band filtering arrangement 12B, though it could be located elsewhere, for example in the quantiser 14), is connected to the output field store 28 to produce read addresses therefor to cause the data therein to be read out to the quantiser 14.

It should be noted that the intermediate field store 22 and the transpose sequencer 24 used in the sub-band filtering arrangement 12A

of Figures 2 and 3 are not necessary when the sub-band filtering arrangement 12B of Figure 5 is used. It should however be noted that the above-described sub-band frequency scrambling that occurs in the sub-band filtering arrangement 12A of Figures 2 and 3 also takes place
5 in the sub-band filtering arrangement 12B of Figure 5. Thus, the output sequencer 29B of the sub-band filtering arrangement 12B of Figure 5 has to perform descrambling.

Before the quantiser 14 is described in more detail, the principle on which it operates will be explained with reference to
10 Figures 6 and 7. Figure 6 is a graph representing an empirically determined equation approximately representing the response of the human psychovisual system to different spatial frequencies, the vertical axis representing the sensitivity of the human psychovisual system, the horizontal axis representing spatial frequency, and the
15 frequency value f_s representing the Nyquist limit sampling frequency. As can be seen from Figure 6, the human psychovisual system is most sensitive to lower frequencies, peaking at a value just above dc (zero spatial frequency), and the sensitivity rapidly drops as the frequency increases. It is therefore readily possible for the quantiser 14 to
20 achieve compression of the sub-band filtered video signal by selectively removing information, in conformity with the graph of Figure 6 (possibly also taking into account the amount of aliasing introduced into each sub-band by the sub-band filtering), to which the human psychovisual system is effectively insensitive. This is done by
25 quantising the 64 sub-bands of the sub-band filtered video signal by respective appropriate amounts. Specifically, it is assumed that circular symmetry extends the (one-dimensional) response curve of Figure 6 to two dimensions. (This assumption is believed justified in that the human psychovisual system is less sensitive to diagonal
30 frequencies than to horizontal and vertical frequencies.) The resultant generated surface is then integrated under each of the 64 sub-band regions to produce an array of 64 numbers (values) which act as thresholds for the purpose of quantisation of respective ones of the sub-bands in the quantiser 14. As will be evident, the numbers
35 determine the extent of quantisation for their respective sub-bands. If, as in the example described below, the numbers are used to achieve quantisation by virtue of their being used to divide data arriving from

the sub-band filtering arrangement 12, then the greater the number, the greater the quantisation threshold and the greater the probability of a sample in the relevant sub-band having a zero or near zero value after quantisation.

5 It should be appreciated that the above-described technique of establishing the 64 numbers to be used for quantising the different sub-bands represents one possible approach only and, even if this approach is used, the numbers derived by the somewhat theoretical method described above may be modified. In more detail, the quality or
10 viewer-acceptability of a picture represented by a video signal which has been compressed by the present (or any other) technique and thereafter expanded by a converse technique is, in the final analysis, a matter of subjective opinion. Thus, a final determination of the numbers to used for quantising the different sub-bands might well best
15 be achieved by selecting rough initial or starting point values by the theoretical method described above and then refining those initial values by viewer testing (trial and error) to produce values judged subjectively to be optimum.

The above-described 64 numbers can be stored in the form of a
20 quantisation matrix (naturally an 8 x 8 matrix in the case of an 8 x 8 sub-band filtered signal), for example in a look-up table in a programmable read only memory (PROM). Figure 7 shows an example of an 8 x 8 quantisation matrix produced for a particular design of sub-band filtering arrangement. The positioning of the numbers in the matrix of
25 Figure 7 corresponds to the positioning of the sub-bands in Figure 4. That is, for example, the number 68 applies to the dc sub-band and the number 8192 applies to the ac sub-band in the bottom right-hand corner in Figure 4. It will be seen that the dc sub-band is only lightly quantised (number = 68). Although the two ac sub-bands horizontally
30 and vertically adjacent to the dc sub-band are quantised marginally even more lightly than the dc sub-band (number = 64), the amount of quantisation (quantisation threshold) of the dc sub-band is, as can clearly be seen from Figure 7, considerably less than the average of the amounts of quantisation (quantisation thresholds) of the ac
35 sub-bands.

The following two factors must be borne in mind concerning the quantisation matrix.

(a) The relative values of the numbers, rather than their absolute values, are of importance. In this regard, as explained below, the numbers in the quantisation matrix may be scaled before they are used to effect quantisation of the sub-bands in the quantiser 14.

5 (b) Since, as mentioned above in the description of Figure 4, it is being assumed for the time being that the input digital video signal is a component (luminance) signal, rather than a composite signal, the numbers represented in Figure 7 apply to a component (luminance) signal. (The modifications made to the quantisation matrix
10 of Figure 7 in the case of processing a composite signal are explained below.)

In the light of the foregoing explanation of its principle of operation, the quantiser 14 will now be described with reference to Figures 8 to 11.

15 Figure 8 shows the quantiser 14 in block diagram form. The quantiser 14 comprises a divider 50 that receives data read thereto from the output field store 28 of the sub-band filtering arrangement 12A or 12B under the control of the output sequencer 29A or 29B, and outputs quantised data from the quantiser 14 to the entropy encoder 16
20 (Figure 1).

The above-mentioned quantisation matrix, referenced 52 in Figure 8, and stored for example in a look-up table in a PROM, is connected to one input of a multiplier 54. A scale factor generator 56 is connected to another input of the multiplier 54. A sequencer (address generator)
25 58 is connected to the quantisation matrix 52 to control it so that it outputs the appropriate one of the 64 numbers stored in the matrix at the correct time, that is so that each sample supplied to the quantiser is quantised in accordance with the sub-band in which it is located, and is connected to the entropy encoder 16 to supply thereto a timing
30 signal that indicates to the entropy encoder whether data being supplied by the quantiser 14 to the entropy encoder results from quantisation of the dc sub-band or quantisation of the ac sub-bands.

The scale factor generator 56 multiplies each of the 64 numbers outputted by the quantisation matrix 52 by a scale factor, whereby the
35 samples of the stored field supplied to the quantiser 14 are divided in the divider 50 by the product of the scale factor and the number currently outputted by the quantisation matrix 52. The scale factor is

usually kept constant throughout the period during which the same stored field is supplied to the quantiser 14 from the sub-band filtering arrangement 12A or 12B, whereby the values for the different sub-band samples as applied by the multiplier 54 to the divider 50 maintain the same relationship relative to one another over the field as do the numbers (shown in Figure 7) in the quantisation matrix 52. However, the absolute values applied by the multiplier 54 to the divider 50 are determined by the value of the scale factor. Variation of the scale factor therefore can vary the output data (bit) rate of the entropy encoder 16, that is of the entire compression apparatus, and can therefore be employed, for example, to keep the data rate (which can vary with image content) constant.

The quantiser 14 reads and processes a field of data stored in the output field store 28 of the sub-band filtering arrangement 12A or 12B, and passes it on after processing to the entropy encoder 16. The processing comprises, as explained above, and as described in more detail below, a selective quantisation operation used to achieve compression of the video signal. In addition, as explained below, the processing involves arrangement of the data outputted to the entropy encoder in a format that readies it for entropy encoding and bit rate reduction.

Since, in the quantiser 14 described above with reference to Figure 8, the quantisation is effected by dividing the input data (in the divider 50), the numbers (Figure 7) in the quantisation matrix 52 must be such that those for sub-bands that are to be quantised by a relatively large amount are greater than those for sub-bands that are to be quantised by a relatively small amount. Instead, the quantisation could be effected by multiplying the input data (in a multiplier taking the place of the divider 50), in which case the numbers in the quantisation matrix 52 would be such that those for sub-bands that are to be quantised by a relatively large amount are smaller than those for sub-bands that are to be quantised by a relatively small amount. (For example, in the latter case the numbers in the quantisation matrix 52 could be reciprocals of those shown in Figure 7.) It will be appreciated that, in both cases, the amount of quantisation of the dc sub-band is considerably less than the average of the amounts of quantisation of the ac sub-bands.

Figure 9 shows a part (the upper left-hand corner) of Figure 4 on an enlarged scale. More accurately, Figure 9 is a map of a sub-band filtered field as supplied to the quantiser 14 from the output field store 28 of the sub-band filtering arrangement 12A or 12B, each sub-band being stored (as mentioned above) in a respective one of an 8 x 8 array of regions into which the store 28 can be considered to be partitioned. In this regard, the stored field comprises an 8 x 8 array of sub-bands filtered from the corresponding field of the input video signal.

10 A field of, for example, an NTSC digital video signal has a horizontal extent of 910 samples and a vertical extent of 262 samples. The sub-band filtering described above is however carried out on the active part only of the field, which part comprises 768 samples in the horizontal direction and 248 samples in the vertical direction. (In
15 fact, there are 243 active samples, corresponding to the number of active lines, in the active part of an NTSC field. In order to produce numbers of active samples in both directions that are integrally divisible by 8, 5 blank lines are added to make the number of active samples in the vertical direction equal to 248.) Thus, each of the 64
20 sub-band areas in the active sub-band filtered field comprises $(768/8) \times (248/8) = 2976$ samples, that is an array of 96 x 31 samples (as shown in Figure 9). (The whole active field comprises, of course, 64 times that number of samples.) The output sequencer 29A or 29B of the sub-band filtering arrangement 12A or 12B is operative to output the
25 samples of the active field stored in the output field store 28 of the sub-band filtering arrangement 12A or 12B as follows.

The sequencer 29A or 29B first causes all of the 2976 samples forming the dc sub-band (the upper left-hand sub-band area in Figure 9), namely those in that one of the 64 regions of the output store 28
30 of the sub-band filtering arrangement 12A or 12B containing the data constituting that sub-band, to be fed in turn to the quantiser 14. This can be done by addressing the relevant regions of the output store 28 in an order akin to the raster scan employed to form the full active field, though in this case the area (and the number of samples) is
35 reduced by a factor of 64 as compared to a full field. The process is represented schematically by the arrowed lines drawn in the upper left-hand sub-band area in Figure 9. The resulting 2976 samples are

supplied in turn to the divider 50. While this process is taking place, the sequencer 58 (which, though shown as a separate item, could be combined with the output sequencer 29A or 29B of the sub-band filtering arrangement 12A or 12B) causes the quantisation matrix 52 to
5 output to the multiplier 54 the number (68) for the dc sub-band. Thus, all the 2976 samples of the dc sub-band are quantised (by the same amount) by being divided in the divider 50 by the product of the number (68) for the dc sub-band and the scale factor (from the scale factor generator 56), and passed on as a run or sequence of 2976 samples to
10 the entropy encoder 16. Also, while the above process is taking place, the sequencer 58 causes the timing signal that it supplies to the entropy encoder 16 to be such as to indicate to the entropy encoder that the quantised samples that it is receiving relate to the dc sub-band.

15 When the dc sub-band samples have been processed through the quantiser 14 as just described, the sequencer 58 causes the timing signal that it supplies to the entropy encoder 16 to be such as to indicate to the entropy encoder that the quantised samples that it is about to receive relate to the ac sub-bands. Thus, the timing signal
20 is changed once per field; that is, it has a frequency equal to the field frequency. The output sequencer 29A or 29B then causes writing to the quantiser 14 of the ac sub-band data, and the sequencer 58 causes a corresponding selection of the numbers to be outputted by the quantisation matrix 52, in a manner now to be described.

25 The ac sub-band data is processed through the quantiser 14 in a rather different manner than the dc sub-band data. An operation is carried out 2976 times, under the control of the output sequencer 29A or 29B, in each of which the respective 63 samples having a respective one of the 2976 spatial positions (pixel sites) in the 63 sub-bands are
30 passed to the divider and multiplied by their respective coefficients. This operation may be more readily understood by referring to Figure 9.

In the first of the above-mentioned 2976 operations, as a first step the first stored sample accessed is the top left-hand one (indicated by a dot) in the ac sub-band numbered 1 in Figure 9. That
35 sample is divided by the product of the scale factor and the number in the quantisation matrix 52 relating to that sub-band, that is the number 64: see Figure 7. Next, as a second step, the same process is

repeated for the top left-hand sample (again indicated by a dot) in the ac sub-band numbered 2 in Figure 9, the number outputted by the quantisation matrix 52 in this case being the number 64. As a third step, the process is repeated for the ac sub-band numbered 3 in Figure 5 9, the number outputted by the quantisation matrix 52 in this case being the number 84. The process is repeated until it has been carried out 63 times, that is for all of the 63 ac sub-bands. The order in which the sub-bands are accessed is in accordance with the sequence 1 to 63 in which the ac sub-bands are designated in Figure 10 (and, for 10 some only of the ac sub-bands, in Figure 9). It will be seen from Figure 10 that the order of processing or scanning of the ac sub-bands is a zig-zag order (shown partially by arrowed chain-dotted lines in Figure 9 for the top left-hand samples) in that it involves scanning the ac sub-bands in a diagonal direction and in opposite senses. 15 (Thus, the legs of the zig-zag comprise successive ones of a series of groups of the 63 ac sub-bands in a sequence as between the groups (legs of the zig-zag) of ac luminance information of increasing spatial frequency.) The above-explained zig-zag scanning technique is based upon, though considerably modified with respect to, a zig-zag scanning 20 technique (described below) that has been proposed as part of the above-mentioned JPEG (Joint Photographic Experts Group) standard, which (rather than sub-band filtering) requires the use of DCT coding with 8 x 8 sample blocks, to each of which an 8 x 8 DCT transform is applied, as mentioned at the beginning of this description.

25 The remaining ones of the above-mentioned 2976 (63-step) operations are carried out in the same manner as the first one, except that, in each case, a respective different one of the 2976 sample sites is used. Thus, for example, in the second operation the samples that are processed are those having the spatial positions indicated by 30 crosses in Figure 9, these being those immediately to the right of those, indicated by dots, that were processed in the first of the operations.

It will be understood from the foregoing explanation that the data inputted to and outputted by the quantiser 14 for the ac sub-bands 35 (only) has a format as represented in Figure 11. That is, 2976 successive series (hereinafter referred to as "scans") - represented in Figure 11 by horizontal strips - of 63 quantised samples are sent to

the entropy encoder 16, each such scan relating to a respective one of the 2976 sub-band pixel sites and each such scan having employed the zig-zag technique of scanning the 63 ac sub-bands as described above. The total number of samples sent to the entropy encoder 16 per field (including the dc sub-band and the ac sub-bands) is the same as the number of samples in the stored sub-band filtered field written to the quantiser. However, as will be evident from the foregoing explanation, the data sent to the entropy encoder no longer has any resemblance to a video field.

During the writing of the dc and ac data from the field store 28 to the quantiser 14 under the control of the sequencer 29A or 29B, the sequencer 58 is operative to control the quantisation matrix 52 such that each sample supplied to the quantiser is appropriately quantised. Specifically, the matrix 52 first continuously outputs the number (68) for the dc sub-band for a period having a duration of 2976 samples, and then outputs the 63 numbers for the ac sub-bands in a 63-stage sample-by-sample zig-zag manner corresponding to the manner in which the samples are written from the field store 28 to the quantiser 14.

The aim of reducing information in the video field by the quantising operation performed in the quantiser 14, and therefore enabling compression to be achieved by virtue of the quantising operation, is achieved by the division operation performed in the divider 50. Thus, particularly for the higher frequency sub-bands, and particularly for image positions that contain little ac spatial frequency information, the sample outputted by the divider 50 will have a zero or very low value, being constituted wholly or mostly by bits of the value zero. It should, however, be noted that, at least in the apparatus presently being described, no reduction in bit (data) rate is carried out in the quantiser 14. That is, the bit length of each sample outputted by the divider 50 is the same as that of the sample inputted to it. However, the presence of long runs of zero value samples in the data outputted by the quantiser 14, and the reduction in the number of levels to which the data inputted thereto can be assigned, enables a consequential bit rate reduction to be effected in the entropy encoder, as described below.

The entropy encoder 16 of the video signal compression apparatus may be embodied in the form shown in Figure 12. The entropy encoder 16

shown in Figure 12 complies with a so-called "baseline" version of the above-mentioned JPEG standard, which version sets out minimal requirements for complying with the standard, whereby it is in many respects of known form or based on known technology and will therefore not be described in great detail.

The entropy encoder 16 shown in Figure 12 comprises a switch 60 controlled by the above-mentioned timing signal provided to the entropy encoder 16 by the sequencer 58 (Figure 8) of the quantiser 14. When the timing signal indicates that the data emerging from the quantiser 14 relates to the ac sub-bands, that is when such data is one of the 2976 successive scans (each having a length of 63 samples) represented in Figure 11, the switch 60 directs the data to a run length detector/data modeller 62. When, on the other hand, the timing signal indicates that the data emerging from the quantiser 14 relates to the dc sub-band, that is when such data is the run or sequence of 2976 samples of the dc sub-band preceding the 2976 successive scans represented in Figure 11, the switch 60 directs the data to a differential pulse code modulator (DPCM) 64. The switch 60 is thus changed over once per field.

The detector/modeller 62 is connected to a PROM 66 containing a variable length code (VLC) look-up table and to a PROM 68 containing a fixed length code (FLC) look-up table. An output of the detector/modeller 62 is connected via a multiplexer 70 to the output 18 of the apparatus.

An output of the DPCM 64 is connected to a data modeller 72, an output of which is in turn connected via the multiplexer 70 to the output 18 of the apparatus. In similar manner to the detector/modeller 62, the modeller 72 is connected to a PROM 74 containing a VLC look-up table and to a PROM 76 containing an FLC look-up table. The VLC PROMs shown at 66 and 74 may in fact be the same PROM: they are shown as being separate in Figure 12 largely for the sake of clarity. Similarly the FLC PROMs shown at 68 and 76 may in fact be the same PROM. Further, rather than being (as shown) a separate item, the modeller 72 can be a part (sub-set) of the detector/modeller 62.

The operation of the entropy encoder 16 shown in Figure 12 will now be described, considering first the case in which the data arriving from the quantiser 14 relates to the ac sub-bands and is therefore

directed by the switch 60 to the detector/modeller 62.

The detector/modeller 62 examines each of the 2976 63-sample scans (Figure 11) arriving from the quantiser 14 and looks for runs of consecutive zero value samples each preceded and followed by a sample of non-zero value. The detector/modeller 62 models the incoming data by converting each such run of zero consecutive value samples to a word pair of the following form:

[RUNLENGTH,SIZE][AMPLITUDE].

The two components or "nibbles" (RUNLENGTH and SIZE) of the first word of the pair each have a length of 4 bits. The bit pattern of the first nibble (RUNLENGTH) represents in binary form the number of consecutive zero value samples in the run and is generated by a counter (not shown) that counts the number of consecutive zero value samples following a previous non-zero value. (Run lengths from 0 to 15 are allowed and a runlength continuation is indicated by a code [F,0].) The bit pattern of the second nibble (SIZE) represents the number of bits to be used to indicate the amplitude of the sample of non-zero (value) amplitude that follows the consecutive run of zero value samples and is looked up from the table - represented in Figure 13 - contained in the FLC PROM 68, the left hand part of Figure 13 representing ranges of actual values (in decimal form) and the right hand part representing values of SIZE for the different ranges. The second word (AMPLITUDE) of the pair represents the amplitude of the sample of non-zero value in the form of a number of bits determined by the value of SIZE. For a positive non-zero value, AMPLITUDE is the result of truncating the non-zero value (in binary form) to have only the number of bits specified by SIZE. For a negative non-zero value, the non-zero value is decremented by one and the same truncation procedure is followed. To illustrate the nature of the word pair by way of an example, suppose that the detector/modeller 62 detects a run of 4 samples of zero value followed by a sample having a value (amplitude) of +7. In this case, the word pair will be as follows:

[4,3][111].

The number 4 (or, more accurately, its binary equivalent, namely 0100) for RUNLENGTH indicates that the length of the run of zero value samples is 4. The number 3 (or, more accurately, its binary equivalent, namely 0011) for SIZE indicates (as can be seen from Figure

13) that 3 bits are used to represent the number +7, namely the amplitude (in decimal form) of the sample of non-zero value (amplitude). The number 111 is in fact the amplitude (+7) of the sample of non-zero value expressed in binary form and truncated to 3
5 bits.

It will be appreciated that the above operation will be carried out for the whole of each scan and that a sequence of word pairs will be generated for each scan. The number of word pairs (that is, the length of the sequence of word pairs) generated for each scan will
10 depend upon the picture content. In general, the greater the number and length of runs of zero value samples, the lesser the number of word pairs.

The operation of the detector/modeller 62 as so far described represents only the first of two stages of data (bit) rate reduction
15 carried out in the detector/modeller. This first stage represents a reduction in bit rate resulting from the above-described reduction of information effected in the quantiser 14 that results (without perceptible degradation in picture content) in a large number of samples of zero value (and, more especially, runs thereof) emerging
20 from the quantiser, especially in the data relating to the ac sub-bands.

The second stage of data rate reduction effected in the detector/modeller 62 is achieved as follows. The first of each of the above-mentioned word pairs is replaced in the data outputted from the
25 detector/modeller 62 with a code therefor looked up in the VLC PROM 66. The VLC PROM 66 stores a respective such code for each possible value of the first word. The codes are of different lengths, and their lengths are selected such that the length of each code is, at least approximately, inversely proportional to the probability of the
30 associated word value occurring. In this way, a further reduction in the data (bit) rate, resulting from entirely loss-free compression, is achieved.

The operation of the entropy encoder 16 shown in Figure 12 will now be described for the case in which the data arriving from the
35 quantiser 14 relates to the dc sub-band and is therefore directed by the switch 60 to the DPCM 64. The dc sub-band (unlike the ac sub-bands) is subjected to DPCM treatment. Since the dc sub-band

contains the intensity information of the original image (field), it has similar statistics to the original image. The ac sub-bands, on the other hand, contain sparse image edge information separated by zero value data and thus have completely different statistics to the dc sub-band. Consequently, it is believed desirable to entropy encode the ac and dc sub-band data separately and in respective different manners to minimise the overall data rate.

Specifically, the dc sub-band data is treated, firstly, in the DPCM 64, prior to entropy encoding proper. The DPCM 64 uses a previous sample predictor with no quantisation of the error data, because the fact that the dc sub-band data represents only a small proportion of the overall data means that high complexity DPCM treatment is difficult to justify. The DPCM 64 decorrelates (adjusts the probability distribution of) the dc sub-band samples so that a greater degree of compression can be achieved in the modeller 72.

Next, entropy encoding proper, resulting in a reduction in the data rate, is carried out in the data modeller 72. The modeller 72 operates similarly to the detector/modeller 62, except that there is no detection of runs of zero value samples, such runs being much less likely in the dc sub-band.

The modeller 72 models the incoming data by converting the incoming data to a sequence of word pairs of the following form:

[SIZE][AMPLITUDE].

As in the case of the ac sub-band data, SIZE is looked up from the FLC table of Figure 13 (in the FLC PROM 76) and indicates the number of bits used to represent AMPLITUDE. The bits used to represent AMPLITUDE are determined in the same way (truncation) as in the case of ac sub-band data. The word SIZE is then encoded in that it is replaced in the data outputted from the modeller 72 with a code therefor looked up in the VLC PROM 74. The VLC PROM 74 stores a respective such code for each possible value of the word. The codes are of different lengths, and their lengths are selected such that the length of each code is, at least approximately, inversely proportional to the probability of the associated word value occurring. In this way, a further reduction in the data (bit) rate, resulting from entirely loss-free compression, is achieved.

Figure 14 is a graph, corresponding to Figure 4, showing,

on the two-dimensional frequency plane, what the inventors have discovered happens when a field of a digital NTSC composite video signal, sampled at a frequency equal to four times the colour sub-carrier frequency f_{sc} (f_{sc} is approximately equal to 3.58 MHz), is sub-band filtered in a video signal compression apparatus as described above. The dc and ac luminance data is distributed among the 64 sub-bands in substantially the same way as described above for a component (luminance) signal. Surprisingly, however, it was found that the chrominance data, or at least the chrominance data that is needed, is largely (substantially) restricted to two only of the sub-bands (shown shaded in Figure 14), namely to those two adjacent sub-bands (hereinafter referred to as "dc chrominance sub-bands") at the bottom centre in Figure 14. Attempts have been made on an ex post facto basis to explain this phenomenon.

As regards the horizontal positioning of the dc chrominance information, this seems on consideration to be appropriate since it should be centred around the position $\pi/2$ along the horizontal axis of Figure 14 by virtue of the use of a sampling frequency equal to $4.f_{sc}$. Thus, if a sampling frequency of other than $4.f_{sc}$ were used, the dc chrominance information would be displaced horizontally from the position shown in Figure 14. If this were the case, the horizontal positioning of the sub-bands to be treated as the dc chrominance sub-bands would differ from that described above.

As regards the vertical positioning of the dc chrominance information in Figure 14, this can be explained as follows. Figure 15 is a graph showing the two-dimensional frequency content of a field of an analog NTSC composite colour video signal, the horizontal axis being in units of MHz and the vertical axis being in units of cycles per picture height (cph). It is of course known that analog NTSC is characterised by a luminance bandwidth of 5.5 MHz and a chrominance bandwidth of 1.3 MHz modulated about the colour sub-carrier frequency of 3.58 MHz. It is also known that the number of sub-carrier cycles per line is 227.5, as a result of which the phase of the sub-carrier is shifted by 180 degrees for each line. This is responsible for a modulation of the chrominance signal vertically, which, as shown in Figure 15, leads to the chrominance being centred at a spectral position of 131.25 cph. This appears to explain the vertical

positioning of the chrominance information in Figure 14. Thus, the process of modulation generates lower and upper sidebands. Since the vertical carrier frequency is at the Nyquist limit frequency, the upper sidebands are on the other side of the Nyquist limit and thus do not form part of the frequency plane of Figure 14. Therefore, for NTSC, the dc chrominance data will appear at the bottom of Figure 14.

As regards the horizontal extent of the dc chrominance information, the fairly harsh filtering (horizontal bandwidth restriction) to which the colour (chrominance) information is subjected before it is modulated onto the luminance information appears to explain why the horizontal extent of the chrominance is restricted as shown in Figure 14, namely so that it falls largely within two horizontally adjacent ones of the 64 sub-bands employed in this case, that is so that the horizontal extent is equal to about $\pi/4$. (In fact, as explained below, the dc chrominance data in fact "spills over" somewhat into the two sub-bands in the bottom row of Figure 14 that are horizontally adjacent to those shown shaded.)

It seems on reflection that the vertical extent of the needed colour information in Figure 14 is restricted to about the height of one of the sub-bands, namely about $\pi/8$, for the following reason. It is probable that the dc chrominance information is wholly or largely restricted to the two sub-bands shown shaded at the bottom of Figure 14. It is likewise probable that ac chrominance appears in at least some of those sub-bands above the two shown shaded at the bottom of Figure 14. However, since the human psychovisual system has a low sensitivity to high frequency (ac) chrominance information, it appears to produce subjectively acceptable results if any such sub-bands that are co-occupied by ac luminance and ac chrominance information are treated as if they are occupied only by ac luminance information.

However, whatever the explanation, the restricted bandwidth (in both directions) of the needed colour information has proven very fortunate because, as is explained below, it leads to the advantageous effect that, with very minor modification, the apparatus as described above can handle an NTSC composite colour video signal. Thus, conversion of the signal to component form, and tripling of the hardware to handle the three components separately, is not necessary, leading to a large saving in expense.

The only modification that has to be made to the apparatus as described above to enable it to handle an NTSC colour composite signal is to change the numbers in the quantisation matrix 52 that determine the amount of quantisation of the sub-bands that contain the dc chrominance data, namely the two dc chrominance sub-bands as shown shaded in Figure 14. Specifically, instead of being heavily quantised as high frequency ac luminance sub-bands of relatively little importance, the two sub-bands should be relatively lightly quantised so as to preserve the dc chrominance information. The amount of quantisation is in fact desirably reduced to about the same level as applied to the dc luminance sub-band. The necessary effect can therefore be achieved by changing the two bottom centre numbers in the quantisation matrix as represented in Figure 7 from their values of 1856 and 2491, for a component (luminance) signal, to 68 (or thereabouts) for an NTSC composite signal. This is shown schematically in Figure 7.

In principle, no changes other than the above-described change to two numbers in the quantisation matrix 52 are necessary to enable the apparatus to handle a digital NTSC composite colour video signal. In particular, it is to be noted that the (now lightly quantised) dc chrominance sub-bands can be handled in the quantiser 14 and entropy encoder 16 together with, and in the same manner as, the ac luminance sub-bands.

Although, in principle, only the above-described change in the quantisation is necessary to enable the apparatus to handle a digital NTSC colour composite signal, another change that can advantageously be made is as follows. The zig-zag sequence or order in which, for a component (luminance) signal, the 63 sub-bands other than the dc luminance sub-band are quantised and then entropy encoded is, as explained above, shown in Figure 10. It will be seen that, in the case of a digital NTSC colour composite signal, the dc chrominance sub-bands have the positions 49 and 57 in the sequence. This could result in a decrease in the efficiency of compression in that the dc chrominance sub-bands are much more likely than the adjacent sub-bands in the sequence to contain non-zero value samples: that is, they could break up runs of zero value samples. (This is even more likely in the case of PAL than NTSC because, as explained below, in the case of PAL there

are four dc chrominance sub-bands positioned in the centre of the frequency plane as shown in Figure 14.) Thus, preferably, the apparatus is further modified in that the sequencer 29A (or 29B) is modified to change the zig-zag sequence so that the dc chrominance sub-bands occupy (in any specified order) the first positions in the sequence and the remaining sub-bands occupy the remaining positions in the sequence in the same order as before. That is, in the case of an NTSC signal, and using the same numbering system for the sub-bands as shown in Figure 10, the sequence will comprise, in the following order, sub-band 49 (or 57), sub-band 57 (or 49), sub-bands 1 to 48, sub-bands 50 to 56, and sub-bands 58 to 63. (The changed sequence that would be adopted in the case of a PAL signal, as will be clear from the description given below with reference to Figure 17, will be sub-bands 24, 31, 32 and 39 (in any order), sub-bands 1 to 23, sub-bands 25 to 30, sub-bands 33 to 38, and sub-bands 40 to 63.) The sequencer 58 in the quantiser 14 (if separate from the sequencer 29A or 29B) is modified in correspondence with the way in which the sequencer 29A or 29B is modified in order to ensure that each sub-band is appropriately quantised. That is, instead of outputting the 63 numbers for the sub-bands other than the dc luminance sub-band as shown in Figure 7 in the same zig-zag order as that in which the sub-bands other than the dc luminance sub-band are numbered 1 to 63 in Figure 10, the sequencer 58 is modified so that it outputs those numbers in an order which is modified in the same way in which the zig-zag sequence of quantising the sub-band filtered samples is (as was just explained) modified.

Further consideration was given to the phenomenon of spectral concentration of the colour information by examining the two-dimensional frequency plane for a frame (as opposed to a field) of a digital NTSC composite colour video signal sampled at 4.fsc, as shown in Figure 16. It will be seen that the composite data in the centre of the frequency plane is composed of four distinct regions due to modulation of the negative frequencies. These four regions are identical except for frequency inversion and a phase shift. Ideally, as explained below, the chrominance data should be restricted to a small number of the sub-bands. Figure 16 indicates that the use of 64 (8 x 8) sub-bands is a good choice in this respect.

Ideally, the horizontal extent or span of the sub-bands should

equal the baseband chrominance bandwidth for efficient compression. This is because, in this case, the chrominance information falls exactly within the relevant sub-bands, that is it occupies the whole of those sub-bands and does not occupy parts of adjacent sub-bands, so
5 that all of the dc chrominance information is lightly quantised and no substantial amount of adjacent ac luminance information is lightly quantised. In other words, a smaller span would lead to the chrominance data falling into a greater number of sub-bands (which is in conflict with the above-mentioned requirement of keeping the number
10 of chrominance sub-bands as small as possible) and a greater span would lead to the adjacent luminance data not being appropriately quantised.

It will be seen from Figure 16 that there is in fact a small overlap or "spill over" of chrominance data into adjacent sub-bands which are treated as ac luminance sub-bands, whereby the overlapping
15 parts of the chrominance will be (heavily) quantised in accordance with the quantisation thresholds set for those adjacent sub-bands. In practice, it is believed that the results will nonetheless be subjectively acceptable. The overlap occurs in the horizontal direction because, as can be seen from Figure 16, the horizontal extent
20 of each sub-band is approximately equal to 0.9 MHz, whereas the chrominance data has a bandwidth (two sidebands) of 1.3 MHz, which is slightly larger. Provided, of course, that the overlap is not so large that a significant amount of low-frequency chrominance information spills over into adjacent sub-bands which are treated in the
25 quantisation process as ac luminance sub-bands, the overlap will generally be tolerable because, as explained above, it will comprise higher frequency chrominance information to which the human psychovisual system is not very sensitive. However, the overlap could be avoided, in theory, by slightly increasing the size of the sub-bands
30 in either or both directions, that is by slightly decreasing the total number of sub-bands. Thus, an inspection of Figure 16 indicates that the overlap would be reduced if a 7 x 7 or a 6 x 6 array were used. While such an array is realisable in theory, it could not be realised in the case of the "tree" or "hierarchical" QMF structure described
35 with reference to Figures 2 and 3 because this can only produce, in each direction, a number of sub-bands which is an integral power of two. Thus, if the tree structure is to be used, the overlap described

above could be avoided only by going down to a 4 x 4 array. While a 4 x 4 array is usable and produces acceptable results, it would result in the extent of the sub-bands that would have to be used as chrominance sub-bands (which, similarly to Figure 14, would be the two at the bottom centre of the 4 x 4 array) being substantially greater than the extent of the dc chrominance data. Also, it would reduce the efficiency of compression by virtue of the fact that the number of sub-bands would be greatly reduced. The reason for this is as follows.

The amount of compression achievable by virtue of the quantisation step decreases, up to a certain extent, as the number of sub-bands decreases. This is because the ratio between the number of ac luminance sub-bands and the number of dc (luminance and chrominance) sub-bands will increase with the total number of sub-bands and the ac sub-bands are on average more heavily quantised than the dc sub-bands. Thus, for example, in above-described case in which there are 64 sub-bands, of which one is a dc luminance sub-band and two (for NTSC) - or four (for PAL, see below) - are dc chrominance sub-bands, either 61 (for NTSC) - or 59 (for PAL) - of the 64 sub-bands are ac luminance sub-bands. That is, either 61/64 or 59/64 of a field can be relatively heavily quantised on average, thereby enabling a higher degree of compression to be achieved than would be the case if the number of sub-bands were less than 64. (Thus, for example, if 16 (4 x 4) sub-bands were used, only 13/16 of a field (for NTSC) would be ac luminance sub-bands.) Therefore, it is in general desirable to use as large a number of sub-bands as is practical, bearing in mind, however, that hardware realisation will become impractical if too many sub-bands are used. Also, if a large increase (over an 8 x 8 array) is made in the number of sub-bands, there will be no net benefit (or at least not a greatly increased benefit) because more than two of the sub-bands (for NTSC) or more than four of the sub-bands (for PAL) may have to be treated (due to extensive overspill of chrominance information) as dc chrominance sub-bands. At present, the use of an 8 x 8 square array (or a non-square array of similar size) is believed to provide a good compromise between the above constraints, though, as mentioned above, a 4 x 4 array is usable. Also arrays having horizontal and vertical extents of 4 and 8, and 8 and 4, respectively, are usable, the latter being considered promising. At the very least, it is highly preferable

for the number of ac luminance sub-bands to exceed the number of dc luminance and chrominance sub-bands.

As an alternative to ignoring limited overspill or increasing the size of the sub-bands to reduce or remove overspill, it is possible to
5 take account of the fact that some chrominance information appears in bands adjacent to these treated (in the quantisation operation) as dc chrominance sub-bands by quantising the adjacent sub-bands to an extent intermediate that to which they would be quantised if considered as ac
10 luminance sub-bands only, and that to which the sub-bands treated as dc chrominance sub-bands are quantised. The actual extent of quantisation of the adjacent sub-bands might well have to be established empirically.

As mentioned above, the use of a sampling frequency equal to four times the colour sub-carrier frequency is preferred since it has the
15 effect of centring the dc chrominance sub-bands about $\pi/2$ in the horizontal direction, that is locating them in the horizontal sense where shown in Figure 14. However, other sampling frequencies can be used.

The foregoing description with reference to Figures 14 to 17 has
20 concentrated on NTSC composite colour video signals. It is to be noted, however, that the technique outlined above can be applied to other broadcast standard composite colour video signals. The application of the technique to PAL composite colour video signals will now be described.

25 Figure 17 is a view corresponding to Figure 4, but showing on the two-dimensional frequency plane both the sub-band filtered field of an NTSC composite colour video signal, and a sub-band filtered field of a PAL composite colour video signal, each sampled at four times its colour sub-carrier frequency. It will be seen that, in the case of
30 PAL, the chrominance information occupies (in the case of an 8 x 8 array of sub-bands) the four sub-bands (shown shaded) clustered at the centre, rather than, as in the case of NTSC, the two at the bottom centre, namely those numbered 24, 31, 32 and 39 in Figure 10.

The only modification that has to be made to the apparatus as
35 described above to enable it to handle a PAL colour composite signal is to change the numbers in the quantisation matrix 52 that determine the amount of quantisation of the sub-bands that contain the chrominance

data in the case of PAL, namely the four PAL dc chrominance sub-bands as shown shaded in the centre of Figure 17. Specifically, instead of being heavily quantised as high frequency ac luminance sub-bands of relatively little importance, the four sub-bands should be relatively
5 lightly quantised so as to preserve the dc chrominance information. As in the case of NTSC, for PAL also the amount of quantisation is in fact desirably reduced to about the same level as applied to the dc luminance sub-band. The necessary effect can therefore be achieved by changing the four numbers clustered in the centre of the quantisation
10 matrix as represented in Figure 7 from their values of 260,396,396 and 581, for a component (luminance) signal, to 68 for a PAL composite signal. This is shown schematically in Figure 7.

Further, in the case of PAL also, the apparatus is desirably further modified (as already indicated above) to change the zig-zag
15 sequence of treatment of the 63 sub-bands other than the dc luminance sub-band so that the four dc chrominance sub-bands come first.

Since, in the case of PAL, the chrominance data occupies 4 of the 64 sub-bands, whereas in the case of NTSC the chrominance data occupies only 2 of the 64 sub-bands, there is a slightly lower potential for
20 compression (as compared to NTSC) for PAL. Specifically, as indicated above, only 59/64 of a field in the case of PAL, as opposed to 61/64 of a field in the case of NTSC, is occupied by ac luminance sub-bands and therefore can be relatively heavily quantised on average.

The invention can, of course, be embodied in other ways than that
25 described above by way of example. For instance, although the above-described apparatus operates on a field-by-field basis, which will generally be more convenient, it could instead operate on a frame-by-frame basis. In this case the sub-bands would have twice the number of samples in the vertical direction and the various field
30 stores would be replaced by frame stores.

Further, although the above-described apparatus operates only on an intra-field basis, whereby sub-band filtering is effected in two dimensions or directions only, namely the horizontal and vertical
35 spatial directions, it could in principle be extended to operate also on an inter-field or inter-frame basis, whereby sub-band filtering would in this case be effected in three dimensions or directions, namely the horizontal and vertical spatial directions and the temporal

dimension or direction.

As alluded to in the introduction to this specification, the format (described above with reference to Figures 9 to 11) of the data written from the field store 28 to the quantiser 14, under the control of the sequencer 29A or 29B, for quantisation and entropy encoding, and the consequential outputting of the numbers of the quantisation matrix 52 (Figure 7), under the control of the sequencer 58, are very different than in the case of the JPEG (DCT) standard. The same applies to the timing signal supplied to the switch 60 of the entropy encoder 16 by the sequencer 58, in that the entropy encoder has in the above case to be switched at the field or frame frequency rather than at the much higher frequency (block frequency) used in the case of the JPEG (DCT) standard. While this form of sequencing is believed superior to the JPEG sequence at least in some cases, in that it groups the dc and ac information together rather than intermingles it, it is subject to the disadvantage that difficulty might arise if the compression apparatus must operate in close conformity with the JPEG standard and/or if there is a desire to construct the apparatus using a quantiser 14 and/or an entropy encoder 16 designed specifically for use in accordance with the JPEG standard. As will now be described, this difficulty can be overcome by modifying the operation as described above to operate in close accordance with the JPEG standard. The modification is possible by virtue of a realisation of an extent of commonality between decorrelation effected by the technique of sub-band filtering and decorrelation effected by way of the very different technique of linear block transformation, for example by way of a DCT technique. The commonality will now be explained, commencing with a brief review of the linear block transformation method.

Figure 18 shows how the samples making up a field or frame of a digital video signal are divided into blocks or arrays of samples which are each to be processed by a linear transform, for example a DCT. It is assumed, by way of example, that the item depicted in Figure 18 is a frame of a 4:2:2 component (luminance or colour difference) signal according to CCIR Recommendation 601. If DCT were employed in a compression apparatus like those embodying the invention as described above, it might instead be the case that a field (or frame) of an NTSC or PAL composite signal would instead in fact be processed. The only

difference in that case is that there would be a different number of blocks since a field (or frame) of an NTSC or PAL composite signal has a different extent (number of samples) in both the horizontal and vertical directions than a CCIR 601 4:2:2 frame.

5 The frame shown in Figure 18 has a horizontal extent of 720 samples and a vertical extent of 576 samples. Prior to being processed by a linear transform, for example a DCT, the frame is divided by suitable hardware into blocks BL of (for example) 8 x 8 samples. Since 720 and 576 are each integrally divisible by 8, the frame is divided
10 into an array of $(720/8 \times 576/8 =)$ 6480 blocks, the array having a horizontal extent of $(720/8 =)$ 90 blocks and a vertical extent of $(576/8 =)$ 72 blocks.

Figure 19 shows a linear transform decorrelator 12C for carrying out the above-outlined operation. A digital input video signal is
15 applied via an input 10 to a blocking circuit 80 that divides each field or frame of the signal into 8 x 8 sample blocks. In a manner analogous to a raster scan, the blocking circuit 80 sequentially outputs the blocks to a linear transform circuit 82 which transforms each block. For convenience, it will be assumed that the linear
20 transform circuit 82 performs a DCT transform; and the circuit will thus hereinafter be referred to as a DCT circuit. However, as indicated above, other suitable linear block transforms known in the art can be used.

The transformation performed by the DCT circuit 82 on each 8 x 8
25 block BL of samples results in the circuit outputting an 8 x 8 block BL(T) of transformed samples which are (somewhat confusingly) referred to in the art as "coefficients". Each coefficient is a sample or measure of the frequency content of the video signal at a respective one of an 8 x 8 array of positions in the two-dimensional frequency
30 domain or plane corresponding to a respective one of the samples inputted to the DCT circuit 82. The coefficient blocks BL(T) are supplied from the DCT circuit 82 to a quantiser 14 and entropy encoder 16 for compression of the signal, as described above with reference to Figure 1. As explained below, the decorrelator 12C causes the
35 coefficients to be supplied to the quantiser 14 in a rather different

manner than that in which the sequencers 29A and 29B of Figures 2 and 5, respectively, cause the sub-band filtered samples in the stores 28 of Figures 2 and 5 to be written to the quantiser.

Figure 20 shows a DCT decorrelator 12D that employs DCT
5 decorrelation followed by a coefficient reordering process that emulates sub-band filtering. The decorrelator 12D of Figure 20 is essentially the same as the decorrelator 12C of Figure 19, except that:
(i) the DCT circuit 82 is followed by a reorder circuit (address generator) 83 followed by a store 84, the reorder circuit 83 being
10 operative, as described below, to cause the coefficients making up the blocks BL(T) emerging from the DCT circuit 82 to be written into the store 84 in a very different manner to that in which they are outputted from the DCT circuit 82 in the case of Figure 19; and (ii) the writing of data from the store 84 to the quantiser 14 is controlled by an
15 output sequencer (address generator) 29D, which can operate similarly to the output sequencer 29A (29B) of Figure 2 (Figure 5).

In the case of Figure 20, the store 84 can be considered to be partitioned into a number of regions equal to the number of coefficients per block (64, that is 8 x 8, in the above example), each
20 such region having a capacity equal to the number of blocks (6480, that is 90 x 72, in the above example). In this case, the 64 coefficients making up each coefficient block BL(T) are spread out over the whole of the store 84 rather than being outputted as a unit to the quantiser 14 as in the case of Figure 19. More specifically, each of the 64
25 coefficients making up each coefficient block BL(T) is written into a respective one of the 64 regions into which the store is partitioned. The exact positioning of each coefficient within its respective region of the store 84 will now be explained with reference to Figures 21 and 22.

30 Figure 21 shows one of the 8 x 8 coefficient blocks BL(T). Figure 22 shows the store 84 of Figure 20 partitioned, as mentioned above, into 64 (8 x 8) regions R each having a capacity equal to 6480 (90 x 72) coefficients. The way in which each coefficient is positioned by the reorder circuit 83 in the store 84 is as follows. Assume that
35 the coefficient block BL(T) is that corresponding to the first input sample block BL, namely that shown in the upper left-hand corner of Figure 18. Employing a convention in which the coefficients in the

block BL(T) are identified as $c(m,n)$, where m varies from 0 to 7 and represents the horizontal position of the coefficient within the block and n varies from 0 to 7 and represents the vertical position of the coefficient within the block, the origin being the coefficient $c(0,0)$ in the upper left-hand corner in Figure 21, and employing an identical convention to identify the regions $R(m,n)$ of the store 84 as shown in Figure 22, each coefficient $c(m,n)$ of the coefficient block BL(T) corresponding to the input sample block BL shown in the upper left-hand corner of Figure 18 is stored in the upper left-hand one of the 90×72 array of storage positions in that one of the regions R identified by the same values of m and n as the coefficient. These coefficients are thus stored in positions represented (for some only of the regions R) by dots in Figure 22.

A similar process is then carried out for the coefficients $c(m,n)$ of the second coefficient block BL(T), namely that corresponding to the input sample block BL which is horizontally adjacent to and on the right of the input sample block shown in the upper left-hand corner of Figure 18. The coefficients of the second block BL(T) are stored in the next set of storage positions in the regions R of the store 84, namely those horizontally adjacent to and on the right of the positions in which the coefficients of the first block BL(T) were stored. The coefficients of the second block BL(T) are thus stored in positions represented (for some only of the regions R) by crosses in Figure 22.

The foregoing process is then repeated, in a manner analogous to a raster scan, for the coefficients $c(m,n)$ of each of the remaining coefficient blocks BL(T), until the coefficients of the final block BL(T), namely that corresponding to the input sample block BL which is shown in the bottom right-hand corner of Figure 18, are stored in the final set of storage positions in the regions R of the store 84, namely those in the bottom right-hand corners of the regions R. That is, the coefficients of the final block BL(T) are stored in positions represented (for some only of the regions R) by circles in Figure 22.

In more general terms, employing a convention in which the coefficient blocks BL(T) are identified as $BL(T)(p,q)$, where p varies from 0 to 90 and represents the horizontal position of the block and q varies from 0 to 72 and represents the vertical position of the block, the origin being the coefficient block corresponding to the input

sample block BL in the upper left-hand corner in Figure 18, and employing an identical convention to identify the storage positions $s(p,q)$ of each of the regions $R(m,n)$ of the store 84, each coefficient $c(m,n)$ of each coefficient block $BL(T)(p,q)$ is stored in that one of the regions R identified by the same values of m and n as the coefficient and, within that region, in that one of the storage position s of that region having the same values of p and q as the coefficient block.

If the structure of the data content of the store 84 as represented in Figure 22 is analysed, it will be seen that, starting from the region $R(0,0)$ in the upper left-hand corner, the content of each region increases in horizontal spatial frequency as one goes right (horizontally) and increases in vertical spatial frequency as one goes down (vertically). That is, for example, the region $R(0,0)$ will contain dc spatial frequency information (that is, the coefficients $c(0,0)$ of all of the blocks $BL(T)$), the region $R(0,7)$ will contain the highest vertical frequency information and dc horizontal frequency information, the region $R(7,0)$ will contain the highest horizontal frequency information and dc vertical frequency information, and the region $R(7,7)$ will contain the highest diagonal frequency information. Thus, the reordering process effected by the reorder circuit 83 results in the content of the store 84 being such that the contents of the different regions R thereof are data sets which are in substance the same as would have been obtained if, instead of being decorrelated in the decorrelator 12D of Figure 20, the video signal had been decorrelated in a decorrelator in the form of a sub-band filtering arrangement, for example either the arrangement 12A described above with reference to Figures 2 and 3 or the arrangement 12B described above with reference to Figure 5. That is, the content of the store 84 of the decorrelator 12D of Figure 20 as read out to the quantiser 14 is substantially the same as the contents of the stores 28 of the decorrelators (sub-band filtering arrangements) 12A and 12B of Figures 2 and 3, and Figure 5, respectively, as read out to the quantiser 14.

Thus, the decorrelator 12D of Figure 20 emulates the sub-band filtering carried out in the decorrelators (sub-band filtering arrangements) 12A and 12B, as a consequence of which the decorrelator 12D could (though this is not done in the present invention, since the

emulated sub-band filtering technique does not preserve the advantages of actual sub-band filtering over the DCT approach) be used in direct substitution for the decorrelator 12A or 12B in a video signal compression signal apparatus which can handle a digital composite colour video signal (or a component video signal). In this regard, assuming that the decorrelator 12D is configured to process, for example, an NTSC signal on a field-by-field basis, the content of the store 84 will correspond to Figure 14. Thus, the region R(0,0) of the store 84 will be quantised in the quantiser 14 (as described above with reference to Figures 7 and 8) on the basis that it contains dc luminance information, the contents of the regions R(3,7) and R(4,7) will be quantised on the basis that they contain dc chrominance information, and the contents of the other 61 regions will be quantised on the basis that they contain ac luminance information. Likewise, if the decorrelator 12D is configured to process a PAL signal on a field-by-field basis, the content of the store 82 will correspond to the relevant parts of Figure 17. Thus, the region R(0,0) of the store 84 will be quantised on the basis that it contains dc luminance information, the contents of the regions R(3,3), R(3,4), R(4,3) and R(4,4) will be quantised on the basis that they contain dc chrominance information, and the contents of the other 59 regions will be quantised on the basis that they contain ac luminance information.

The foregoing description with reference to Figures 18 to 22 substantiates the above suggestion of a commonality or duality between sub-band filtering and linear transformation in that the two-dimensional spatial frequency information obtained in the case of the former is present also in the coefficients obtained in the case of the latter and can be recovered by data reordering. Thus, it is possible to compress a digital composite colour video signal (without splitting the composite signal into its components), not only by using sub-band filtering for decorrelation, but also by emulating sub-band filtering by using linear transform (for example, DCT) decorrelation followed by data reordering. However, as will now be described, realisation of the commonality leads to the further development that merely by altering the operation of the sequencer 29A (or 29B) of the decorrelator 12A (or 12B) of Figures 2 and 3 (or Figure 5), and correspondently modifying the operation of the sequencer 58 of the

quantiser 14, a video signal compression apparatus embodying the invention, in which the data format outputted by the decorrelator 12A (or 12B) is substantially the same as that of the JPEG standard, can be achieved.

5 In the foregoing regard, in the apparatus described above with reference to Figures 2 to 17, the sub-bands stored in respective ones of the 64 regions into which the output store 28 of the decorrelator 12A or 12B is partitioned comprise respective data sets representing dc
10 luminance information, ac luminance information and (if a composite colour video signal is being compressed) dc chrominance information of the video signal in the two-dimensional frequency domain. It was demonstrated above that there is a commonality or duality between sub-band filtering and transform decorrelation in that, in the case of transform decorrelation as described with reference to Figures 18 and
15 19, the data sets obtained in the case of sub-band filtering are still present (in that each coefficient block BL(T) contains a respective member of each of the data sets) and the data sets can be put into storage (in the store 84 of Figure 20) in the same manner as in the case of sub-band filtering so as thereby to emulate sub-band filtering
20 whereby the stored data sets can be treated after outputting from the store 84 in substantially exactly the same way as if they had been obtained by sub-band filtering. Pursuing the commonality further, it is in fact the case that, since the data sets obtained in the case of sub-band filtering are still present in the case of transform decorrelation, it is possible to reorder the data outputted by the
25 decorrelator 12A or 12B to conform to the format obtained by the use of transform coding per se, that is as would be obtained if the DCT decorrelator 12C of Figure 19 were used.

This is accomplished as follows. Instead of the sequencer 29A or
30 29B being operative (as described above with reference to Figures 9 to 11) to first scan or output all of the 2976 samples (for NTSC) of the storage region of the store 28 holding the dc sub-band and then to zig-zag scan the storage regions holding the remaining 63 sub-bands (each made up of 2976 samples) 2976 times, each time scanning 63 of the
35 samples having a common one of the 2976 possible spatial positions, the sequencer 29A or 29B zig-zag scans all 64 of the storage regions (each made up of 2976 samples) 2976 times, each time scanning all 64 of the

samples having a common one of the 2976 possible spatial positions. That is, the operation of the sequencer 29A or 29B is modified with respect to that described with reference to Figures 9 to 11 in that the samples in the field store 28 are outputted by zig-zag scanning them in an order which is the same as that shown in Figure 10, save that the order in the present case is a 64-stage one (rather than a 63-stage one) starting with the area in Figure 10 that is not numbered and then carrying on in the order of the areas numbered 1 to 63. The 64-stage zig-zag would be as shown by arrowed chain-dotted lines in Figure 21 if the dots in Figure 21 were considered to represent (instead of the coefficients of a transformed coefficient block BL(T) in the DCT decorrelator 12D of Figure 20) those pixels represented by dots in Figure 9.

The data that in this case is inputted to and outputted from the quantiser 14 has a very different form than that described with reference to Figure 11. Instead of there being a run of 2976 samples relating to the dc sub-band followed (as shown in Figure 11) by 2976 scans or sequences of 63 samples (one for each ac sub-band) each relating to a respective one of the 2976 sub-band spatial positions, there are 2976 successive scans or sequences of 64 samples (one for each of the 64 sub-bands) each relating to a respective one of the 2976 sub-band spatial positions. It is therefore necessary for the sequencer 58 of the quantiser 14 to operate in a correspondingly different manner. That is, instead of first outputting the same number from the quantisation matrix 52 as shown in Figure 7 (that for the dc sub-band) continuously for a period having a duration of 2976 samples, and then cyclically outputting the 63 numbers for the other sub-bands, 2976 times, in a 63-stage, sample-by-sample zig-zag manner, as described above, so as to conform with the manner of reading the output field store 28 as described above, in this case the sequencer 58 always cyclically outputs all 64 numbers of the quantisation matrix 52 in a zig-zag manner which is the same as the 64-stage zig-zag scanning of the 64 samples of the same spatial position in the 64 regions of the store 28 containing the respective sub-bands.

Also, the timing signal supplied to the entropy encoder 16 of Figure 13 by the sequencer 58 must be altered to reflect the fact that there is a difference in the timing of receipt by the entropy encoder

of data relating to dc frequency information (which is switched by the switch 60 to the DPCM 64) and ac frequency information (which is switched by the switch 60 to the run length detector/data modeller 62). As described above with reference to Figures 9 to 11, in that case the switch 60 switches the data to the DPCM 64 for the initial run of samples (duration of 2976 samples, in the case of NTSC) of the dc sub-band, and then switches the data to the detector/modeller 62 for the other 63 sub-bands (duration of 63×2976 samples). That is, the switch 60 is changed over once per field (or frame). In the present case, the sequencer 58 must be operative to change over the switch 60 rather more frequently. Specifically, the switch 60 is changed over every 64 samples, that is once per 64-stage zig-zag scan, to supply the first sample of each scan (namely that in the region of the store 28 containing the dc sub-band), after quantisation, to the DPCM 64, and to supply the remaining 63 samples of each zig-zag scan (namely those in the other 63 storage regions of the store 28), after quantisation, to the detector/modeller 62. In short, while the switch 60 has to be changed over (by the sequencer 58) once per field (or frame) in the case of Figures 9 to 11, in the present case the switch 60 has to be changed over every 64 samples, that is once every 64-stage zig-zag scan.

Thus, in summary of the above-described modification, the format of the data to be quantised is very similar to that in the case of the JPEG (DCT) standard, which has the advantage that quantisation can be sequenced in a very similar or even identical way to that used in the case of the standard. Thus, it may be possible to use an "off the shelf" chip or assembly intended for use in a JPEG compression apparatus (possibly with changes in the quantisation values, that is the numbers in the quantisation matrix 52 as shown in Figure 7) for the quantiser 14. Also, the entropy encoder 16 is switched at the frequency of carrying out the 64-stage zig-zag scanning operations in which those samples corresponding to a respective one of the spatial positions in each of the stored data sets are quantised, that is at a frequency determined by the number of data sets (sub-bands), rather than at the field or frame frequency. If, as described, the number of data sets (sub-bands) is the same ($8 \times 8 = 64$) as the number of samples per block as specified in the JPEG standard, the frequency of switching

the entropy encoding is the same as the frequency (block frequency) used in the case of the JPEG (DCT) standard. Thus, it may be possible to use an "off the shelf" chip or assembly intended for use in a JPEG compression apparatus for the entropy encoder 16.

- 5 Reference is directed to our UK Patent Application No 9100591.8, (Publication No GB-A-) which was filed on 11 January 1991 and includes claims directed to the foregoing disclosure.

CLAIMS

1. A method of compressing a video signal, the method comprising:
5 effecting spatial two-dimensional sub-band filtering of a digital video signal to form a plurality of data sets constituting respective sub-bands of the two-dimensional spatial frequency domain;
storing the data sets for a field or frame of the video signal;
quantising the stored data sets in accordance with respective
10 values, said values being such that the amount of quantisation of one of the data sets constituting a sub-band to which dc luminance information of the signal is at least predominantly confined is less than the average of the amounts of quantisation of the remaining data sets;
15 writing the stored data to a quantiser, for carrying out the quantising, in a desired sequence;
controlling the quantising in accordance with said desired sequence such that each datum written to the quantiser is appropriately quantised;
20 entropy encoding the quantised data sets such that quantised data representative of dc luminance information is coded by a first coding technique and quantised data representative of ac luminance information is coded by a second coding technique; and
controlling the entropy encoding in accordance with said desired
25 sequence such that each quantised datum is subjected to the appropriate one of the first and second coding techniques.
2. A method according to claim 1, wherein:
said desired sequence in which the stored data is written to the
30 quantiser is such that
(i) all the data of said one of the stored data sets constituting the sub-band to which the dc luminance information is at least predominantly confined are quantised,
(ii) after which, upon each occurrence of an operation carried
35 out for a number of times equal to the number of data in each stored data set, those data corresponding to a respective one of the spatial positions in each of the remaining stored data sets are quantised in a

predetermined order;

the quantising is controlled such that all the data written to the quantiser in step (i) are quantised in the same amount, namely in the amount appropriate to the sub-band to which the dc luminance information is at least predominantly confined, and each of the data written to the quantiser in each of the operations of step (ii) is quantised in an amount appropriate to the data set of which it forms a part; and

the entropy encoding is controlled such that all the data quantised upon being written to the quantiser in step (i) are subjected to the first coding technique and all the data quantised upon being written to the quantiser in all of the operations of step (ii) are subjected to the second coding technique.

3. A method according to claim 2, wherein said predetermined order comprises successive groups of the remaining stored data sets in a sequence, as between the groups, of sub-bands containing ac luminance information of increasing spatial frequency.

4. A method according to claim 2, wherein:

the digital video signal is a digital composite colour video signal, the stored data sets are quantised in accordance with respective values which are such that the amounts of quantisation of each of said one of the data sets constituting the sub-band to which dc luminance information of the signal is at least predominantly confined and of at least two of the data sets constituting sub-bands to which dc chrominance information of the signal is at least predominantly confined are less than the average of the amounts of quantisation of the other data sets; and

said predetermined order comprises, first, said at least two of the data sets constituting sub-bands to which the dc chrominance information of the signal is at least predominantly confined and, thereafter, successive groups of the other stored data sets in a sequence, as between the groups, of sub-bands containing ac luminance information of increasing spatial frequency.

5. A method according to claim 1, wherein:

said desired sequence in which the stored data is written to the quantiser is such that, upon each occurrence of an operation carried out for a number of times equal to the number of data in each stored data set, those data corresponding to a respective one of the spatial positions in each of the stored data sets are quantised in a predetermined order;

the quantising is controlled such that each of the data written to the quantiser in each of said operations is quantised in an amount appropriate to the data set of which it forms a part; and

the entropy encoding is controlled such that, for each of said operations, that one of the quantised data forming part of the data set constituting the sub-band to which the dc luminance information is at least predominantly confined is subjected to the first coding technique and all of the other data are subjected to the second coding technique.

15

6. A method according to claim 5, wherein said predetermined order comprises successive groups of the stored data sets in a sequence, as between the groups, of sub-bands containing ac luminance information of increasing spatial frequency.

20

7. A method according to claim 5, wherein:

the digital video signal is a digital composite colour video signal, the stored data sets are quantised in accordance with respective values which are such that the amounts of quantisation of each of said one of the data sets constituting the sub-band to which dc luminance information of the signal is at least predominantly confined and of at least two of the data sets constituting sub-bands to which dc chrominance information of the signal is at least predominantly confined are less than the average of the amounts of quantisation of the other data sets; and

30

said predetermined order comprises, first, said one of the data sets constituting the sub-band to which the dc luminance information of the signal is at least predominantly confined, second, said at least two of the data sets constituting sub-bands to which the dc chrominance information of the signal is at least predominantly confined and, thereafter, successive groups of the other stored data sets in a sequence, as between the groups, of sub-bands containing ac luminance

35

information of increasing spatial frequency.

8. A method according to any one of claims 3, 4, 6 and 7, wherein
said groups comprise legs of a zig-zag pattern connecting the data of
5 the same spatial position in the different stored data sets.

9. A method according to claim 4 or claim 7, wherein the amounts of
quantisation of the sub-bands constituted by the data sets to which the
dc luminance information and the dc chrominance information is at least
10 predominantly confined are at least approximately the same as one
another.

10. A method according to any one of claims 4, 7 and 9, wherein the
number of said other data sets exceeds the number of the data sets
15 constituting the sub-bands to which the dc luminance information and
the dc chrominance information is at least predominantly confined.

11. A method according to any one of claims 4, 7, 9 and 10, wherein
the composite colour video signal is an NTSC signal.

20

12. A method according to any one of claims 4, 7, 9 and 10, wherein
the composite colour video signal is a PAL signal.

13. A method according to any one of claims 1 to 12, wherein the
25 sub-bands make up a square array in the two-dimensional spatial
frequency domain.

14. A method according to claim 13, wherein the square array is a 4
x 4 array.

30

15. A method according to claim 13, wherein the square array is an 8
x 8 square array.

16. A method according to any one of claims 1 to 12, wherein the sub-
35 bands make up, in the two-dimensional spatial frequency domain, a
rectangular array having a dimension of 8 in the direction of scanning
of the video signal and a dimension of 4 in the direction orthogonal

thereto.

17. A method according to any one of the preceding claims, wherein the digital video signal is separately spatially sub-band filtered in
5 respective orthogonal spatial directions.

18. A method according to claim 17, wherein a field or frame of the digital video signal is sub-band filtered in one of the orthogonal directions in a first one-dimensional sub-band filter arrangement,
10 stored, and then transposed and sub-band filtered in the other of the orthogonal directions in a second one-dimensional sub-band filter arrangement which is of substantially the same construction as the first one-dimensional sub-band filter arrangement.

19. A method according to claim 17 or claim 18, wherein, in a first stage of the filtering in each of the orthogonal directions, the digital video signal is subjected to low pass filtering followed by decimation by two and also to high pass filtering followed by decimation by two, thereby to produce two intermediate outputs, and, in
15 at least one subsequent stage of the filtering in each of the orthogonal directions, each of the intermediate outputs produced in the previous stage is subjected to low pass filtering followed by decimation by two and also to high pass filtering followed by decimation by two.
20

25

20. A method according to claim 17, wherein the digital video signal is sub-band filtered in a first one-dimensional sub-band filter arrangement configured to sub-band filter the signal in one of the orthogonal directions, and then sub-band filtered in a second
30 one-dimensional sub-band filter arrangement configured to sub-band filter the signal in the other of the orthogonal directions.

21. Apparatus for compressing a video signal, the apparatus comprising:

35 a spatial two-dimensional sub-band filtering arrangement operative to filter a digital video signal to form a plurality of data sets constituting respective sub-bands of the two-dimensional spatial

frequency domain;

a store for storing the data sets for a field or frame of the video signal;

5 a quantiser operative to quantise the stored data sets in accordance with respective values, said values being such that the amount of quantisation of one of the data sets constituting a sub-band to which dc luminance information of the signal is at least predominantly confined is less than the average of the amounts of quantisation of the remaining data sets;

10 an entropy encoder operative to encode the quantised data sets, the entropy encoder comprising a first coding portion for coding quantised data representative of dc luminance information and a second coding portion for coding quantised data representative of ac luminance information; and

15 sequencing means operative to control writing of the stored data from the store to the quantiser in a desired sequence, to control the operation of the quantiser in accordance with said desired sequence such that each datum written thereto is appropriately quantised, and to control the operation of the entropy encoder in accordance with said
20 desired sequence such that each quantised datum is directed to the appropriate one of the first and second coding portions.

22. Apparatus according to claim 21, wherein the entropy encoder comprises a switch for directing quantised data to the appropriate one
25 of the first and second coding portions and the sequencing means is operative to control the operation of the entropy encoder by supplying thereto a timing signal for controlling the switch.

23. Apparatus according to claim 21 or claim 22, wherein the
30 sequencing means comprises a first sequencer operative to control writing of the stored data from the store to the quantiser, and a second sequencer operative to control the operation of the quantiser and to control the operation of the entropy encoder.

35 24. Apparatus according to claim 21, claim 22 or claim 23, wherein: said desired sequence in which the sequencing means is operative to control writing of the stored data from the store to the quantiser

is such that

(i) all the data of said one of the stored data sets constituting the sub-band to which the dc luminance information is at least predominantly confined are quantised,

5 (ii) after which, upon each occurrence of an operation carried out for a number of times equal to the number of data in each stored data set, those data corresponding to a respective one of the spatial positions in each of the remaining stored data sets are quantised in a predetermined order;

10 the sequencing means is operative to control the operation of the quantiser such that all the data written to the quantiser in step (i) are quantised in the same amount, namely in the amount appropriate to the sub-band to which the dc luminance information is at least predominantly confined, and each of the data written to the quantiser
15 in each of the operations of step (ii) is quantised in an amount appropriate to the data set of which it forms a part; and

the sequencing means is operative to control the operation of the entropy encoder such that all the data quantised upon being written to the quantiser in step (i) are directed to the first coding portion and
20 all the data quantised upon being written to the quantiser in all of the operations of step (ii) are directed to the second coding portion.

25 25. Apparatus according to claim 24, wherein the sequencing means is so operative that said predetermined order comprises successive groups of the remaining stored data sets in a sequence, as between the groups, of sub-bands containing ac luminance information of increasing spatial frequency.

30 26. Apparatus according to claim 24, wherein:

in order to enable compression of a digital composite colour video signal, the quantiser is operative to quantise the stored data sets in accordance with respective values which are such that the amounts of quantisation of each of said one of the data sets
35 constituting the sub-band to which dc luminance information of the signal is at least predominantly confined and of at least two of the data sets constituting sub-bands to which dc chrominance information of

the signal is at least predominantly confined are less than the average of the amounts of quantisation of the other data sets; and

the sequencing means is so operative that said predetermined order comprises, first, said at least two of the data sets
5 constituting sub-bands to which the dc chrominance information of the signal is at least predominantly confined and, thereafter, successive groups of the other stored data sets in a sequence, as between the groups, of sub-bands containing ac luminance information of increasing spatial frequency.

10

27. Apparatus according to claim 21, claim 22 or claim 23, wherein:

said desired sequence in which the sequencing means is operative to control writing of the stored data from the store to the quantiser is such that, upon each occurrence of an operation carried out for a
15 number of times equal to the number of data in each stored data set, those data corresponding to a respective one of the spatial positions in each of the stored data sets are quantised in a predetermined order;

the sequencing means is operative to control the operation of the quantiser such that each of the data written to the quantiser in each
20 of said operations is quantised in an amount appropriate to the data set of which it forms a part; and

the sequencing means is operative to control the operation of the entropy encoder such that, for each of said operations, that one of the quantised data forming part of the data set constituting the sub-band
25 to which the dc luminance information is at least predominantly confined is directed to the first coding portion and all of the other data are directed to the second coding portion.

28. Apparatus according to claim 27, wherein the sequencing means is
30 so operative that said predetermined order comprises successive groups of the stored data sets in a sequence, as between the groups, of sub-bands containing ac luminance information of increasing spatial frequency.

35 29. Apparatus according to claim 27, wherein:

in order to enable compression of a digital composite colour video signal, the quantiser is operative to quantise the stored data

sets in accordance with respective values which are such that the amounts of quantisation of each of said one of the data sets constituting the sub-band to which dc luminance information of the signal is at least predominantly confined and of at least two of the data sets constituting sub-bands to which dc chrominance information of the signal is at least predominantly confined are less than the average of the amounts of quantisation of the other data sets; and

the sequencing means is so operative that said predetermined order comprises, first, said one of the data sets constituting the sub-band to which the dc luminance information of the signal is at least predominantly confined, second, said at least two of the data sets constituting sub-bands to which the dc chrominance information of the signal is at least predominantly confined and, thereafter, successive groups of the other stored data sets in a sequence, as between the groups, of sub-bands containing ac luminance information of increasing spatial frequency.

30. Apparatus according to any one of claims 25, 26, 27 and 29, wherein said groups comprise legs of a zig-zag extending diagonally of the array pattern connecting the data of the same spatial position in the different stored data sets.

31. Apparatus according to claim 26 or claim 29, wherein the quantiser is so operative that the amounts of quantisation of the sub-bands constituted by the data sets to which the dc luminance information and the dc chrominance information is at least predominantly confined are at least approximately the same as one another.

32. Apparatus according to any one of claims 26, 29 and 31, wherein the number of said other data sets exceeds the number of data sets constituting the sub-bands to which the dc luminance information and the dc chrominance information is at least predominantly confined.

33. Apparatus according to any one of claims 26, 29, 31 and 32, which is capable of compressing an NTSC composite colour video signal.

34. Apparatus according to any one of claims 26, 29, 31 and 32, which is capable of compressing a PAL composite colour video signal.

35. Apparatus according to any one of claims 21 to 34, wherein the
5 sub-bands make up a square array in the two-dimensional spatial frequency domain.

36. Apparatus according to claim 35, wherein the sub-bands make up a
4 x 4 square array.

10

37. Apparatus according to claim 35, wherein the sub-bands make up an
8 x 8 square array.




38. Apparatus according to any one of claims 21 to 34, wherein the
15 sub-bands make up, in the two-dimensional spatial frequency domain, a rectangular array having a dimension of 8 in the direction of scanning of the video signal and a dimension of 4 in the direction orthogonal thereto.

20 39. Apparatus according to any one of claims 21 to 38, wherein the spatial two-dimensional sub-band filtering arrangement is operative to separately spatially sub-band filter the digital video signal in respective orthogonal spatial directions.

25 40. Apparatus according to claim 39, wherein the spatial sub-band filtering arrangement comprises first and second one-dimensional sub-band filter arrangements of substantially the same construction as one another, the first filter arrangement being connected to receive the digital video signal, storage means for storing a field or frame of the
30 digital video signal after it has been sub-band filtered in one of the orthogonal directions in the first one-dimensional sub-band filter arrangement, and a transpose sequencer operative to transpose the stored field or frame and write the transposed field or frame to the second one-dimensional sub-band filter arrangement to be sub-band
35 filtered in the other of the orthogonal directions.

41. Apparatus according to claim 39 or claim 40, wherein, for filtering in each of the orthogonal directions, the spatial sub-band filtering arrangement comprises a first stage having a low pass filter followed by a decimator for subjecting the digital video signal to low pass filtering followed by decimation by two, and also a high pass filter followed by a decimator for subjecting the digital video signal to high pass filtering followed by decimation by two, thereby to produce two intermediate outputs, and at least one subsequent stage which duplicates the construction of the previous stage whereby each of the intermediate outputs produced in the previous stage is subjected to low pass filtering followed by decimation by two and also to high pass filtering followed by decimation by two.

42. Apparatus according to claim 39, wherein the spatial sub-band filtering arrangement comprises a first one-dimensional sub-band filter arrangement connected to receive the digital video signal and configured to sub-band filter the signal in one of the orthogonal directions, and a second one-dimensional sub-band filter arrangement connected to receive the signal filtered by the first filter arrangement and configured to sub-band filter it in the other of the orthogonal directions.






FIG. 3

3/12

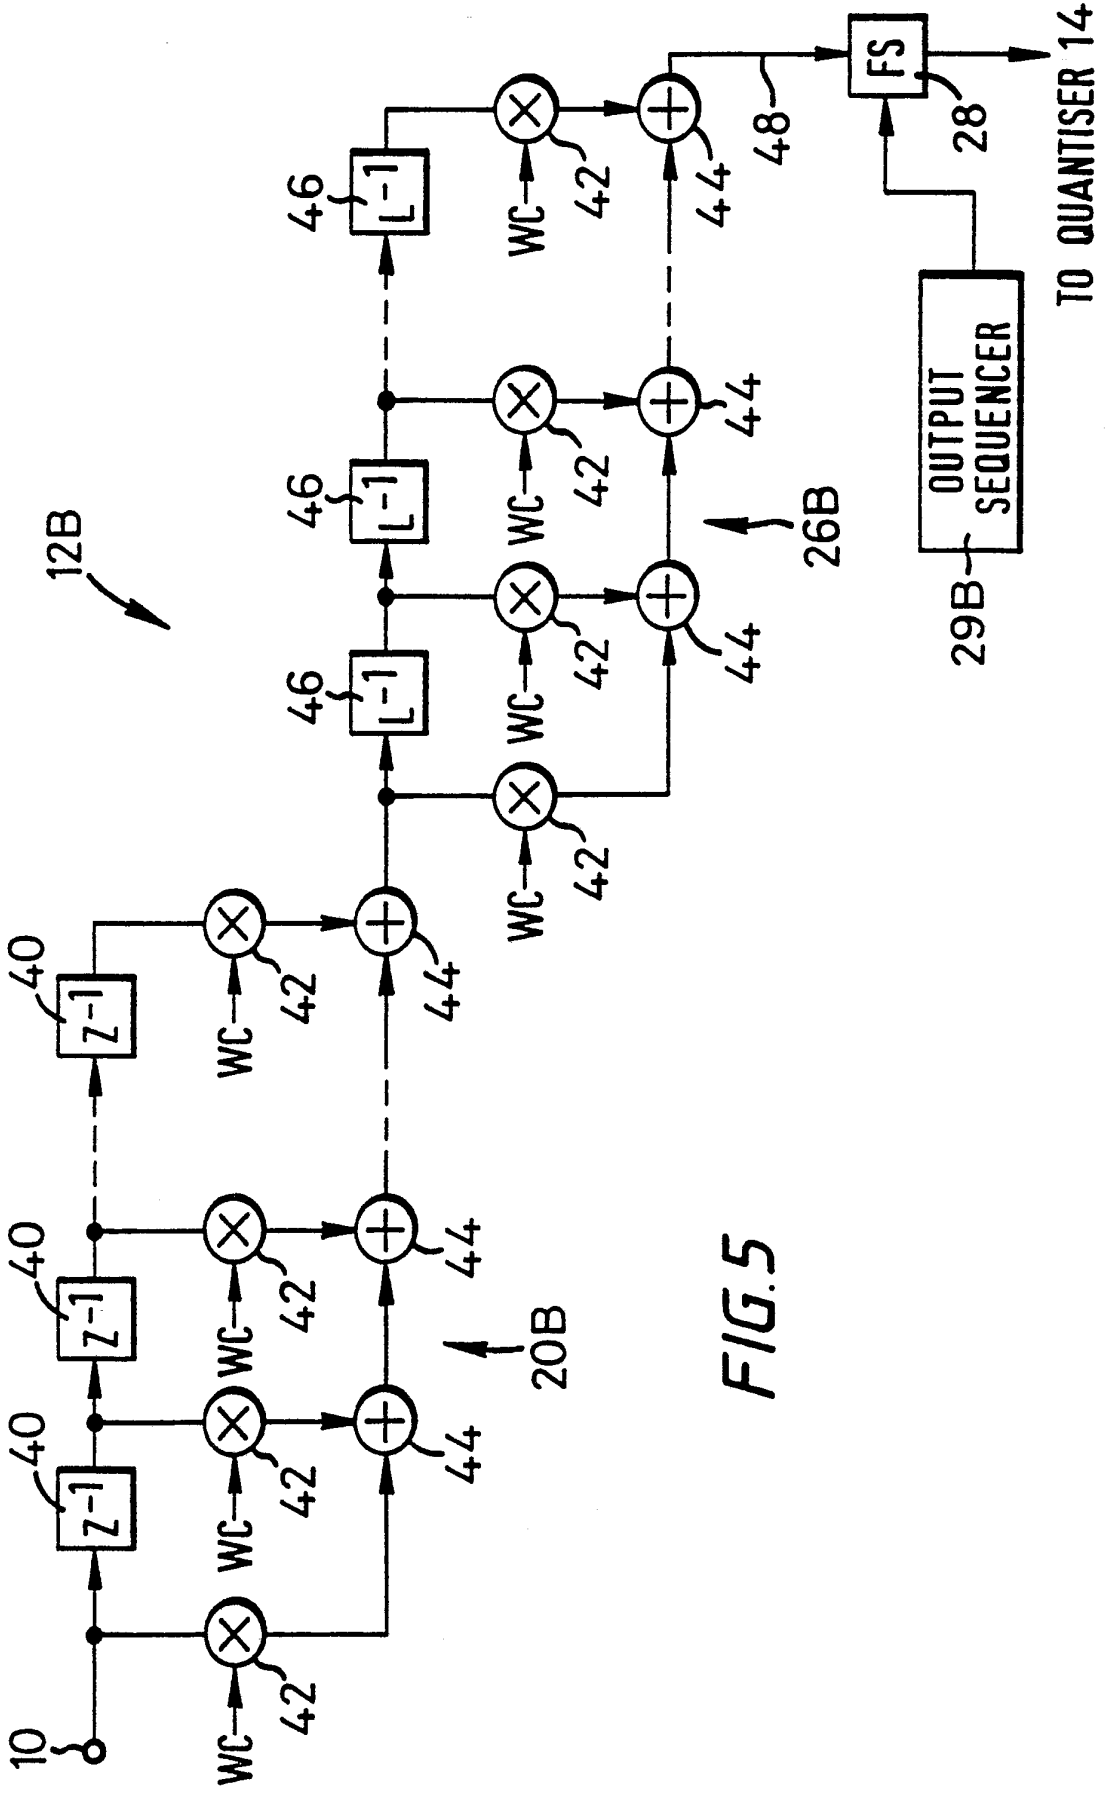


FIG. 5

4/12





FIG. 7



5/12

DATA FROM OUTPUT FIELD STORE 28
OF SUB-BAND FILTERING ARRANGEMENT 12A OR 12B




FIG. 8




FIG. 9

6/12

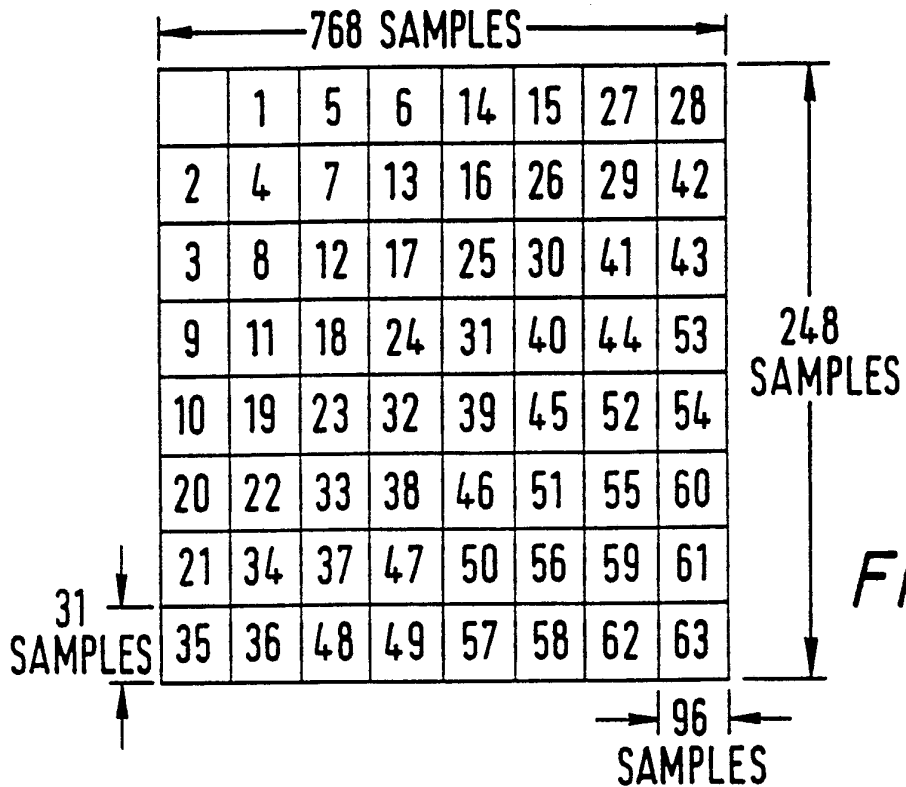


FIG.10




FIG.11

7/12

FIG. 12




FIG. 13

FLC	SIZE
-1, 1	1
-3, -2, 2, 3	2
-7, ... -4, 4, ... 7	3
-15, ... -8, 8, ... 15	4
-31, ... -16, 16, ... 31	5
-63, ... -32, 32, ... 63	6
-127, ... -64, 64, ... 127	7
-255, ... -128, 128, ... 255	8
-511, ... -256, 256, ... 511	9
-1023, ... -512, 512, ... 1023	10
-2047, ... -1024, 1024, ... 2047	11
-4095, ... -2048, 2048, ... 4095	12

8/12




FIG.14

dc chrominance sub-bands





FIG.15


9/12

FIG.16



10/12

FIG. 17



17/12




FIG. 18




FIG. 19




FIG. 20

12/12




FIG. 21





FIG. 22

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 92/00066

I. CLASSIFICATION OF SUBJECT MATTER (if several classification symbols apply, indicate all) ⁶		
According to International Patent Classification (IPC) or to both National Classification and IPC Int.Cl. 5 H04N7/13; H04N7/137		
II. FIELDS SEARCHED		
Minimum Documentation Searched ⁷		
Classification System	Classification Symbols	
Int.Cl. 5	H04N	
Documentation Searched other than Minimum Documentation to the Extent that such Documents are Included in the Fields Searched ⁸		
III. DOCUMENTS CONSIDERED TO BE RELEVANT⁹		
Category ¹⁰	Citation of Document, ¹¹ with indication, where appropriate, of the relevant passages ¹²	Relevant to Claim No. ¹³
Y	EP,A,0 396 368 (AT & T) 7 November 1990	1, 13, 14, 17-20, 21, 35, 36, 39-42
A	see column 1, line 1 - column 8, line 26; figures 1-6	2, 5, 22-24, 27
Y	---	
Y	OPTICAL ENGINEERING. vol. 27, no. 7, July 1988, BELLINGHAM US pages 574 - 586; G. KARLSSON ET AL.: 'Subband coding of video for packet networks'	1, 13, 14, 17-20, 21, 35, 36, 39-42
A	see paragraph 3 -paragraph 4.2; figures 4-8	2, 5, 22, 24, 27

	-/--	
<p>¹⁰ Special categories of cited documents :</p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier document but published on or after the international filing date</p> <p>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p> <p>"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</p> <p>"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step</p> <p>"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.</p> <p>"&" document member of the same patent family</p>		
IV. CERTIFICATION		
Date of the Actual Completion of the International Search	Date of Mailing of this International Search Report	
14 MAY 1992	- 9. 06. 92	
International Searching Authority	Signature of Authorized Officer	
EUROPEAN PATENT OFFICE	DOCKHORN H. S. 	

III. DOCUMENTS CONSIDERED TO BE RELEVANT (CONTINUED FROM THE SECOND SHEET)		
Category °	Citation of Document, with indication, where appropriate, of the relevant passages	Relevant to Claim No.
A	<p>SIGNAL PROCESSING IMAGE COMMUNICATION. vol. 2, no. 3, October 1990, AMSTERDAM, NL pages 333 - 341; K. IRIE ET AL.: 'Adaptive sub-band DCT coding for HDTV signal transmission' see paragraph 2; figures 1-4; tables 1-4</p> <p>---</p>	1,2,5, 13, 17-24, 39-42
A	<p>Proceedings ICASSP'89, vol. 3, 23-26 May 1989, IEEE, NEW YORK, US, pages 1945-1948, R. SAFRANEK ET AL.: 'A perceptually tuned sub- band coder with image dependent quantization and post-quantization data compression' see the whole document</p> <p>---</p>	1,21
A	<p>IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS. vol. 35, no. 2, February 1988, NEW YORK US pages 258 - 261; H. GHARAVI: 'Low bit-rate video transmission for ISDN application' see paragraph 2.5; figure 2</p> <p>---</p>	1,2,5, 21,24,27
A	<p>IEEE TRANSACTIONS ON COMMUNICATIONS. vol. COM34, no. 3, March 1986, NEW YORK US pages 310 - 319; E. DUBOIS ET AL.: 'Encoding and progressive transmission of still pictures in NTSC composite format using transform domain methods' see the whole document</p> <p>---</p>	1,2,4,21

**ANNEX TO THE INTERNATIONAL SEARCH REPORT
ON INTERNATIONAL PATENT APPLICATION NO. GB 9200066
SA 55276**

This annex lists the patent family members relating to the patent documents cited in the above-mentioned international search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information. 14/05/92

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP-A-0396368	07-11-90	CA-A- 2014935 JP-A- 2305191	04-11-90 18-12-90
