

(12) PATENT
(19) AUSTRALIAN PATENT OFFICE

(11) Application No. AU 199873714 B2
(10) Patent No. 727601

(54) Title
Method for data compression

(51)⁶ International Patent Classification(s)
G06K 009/36 G06K 009/46

(21) Application No: **199873714** (22) Application Date: **1998.05.06**

(87) WIPO No: **W098/50886**

(30) Priority Data

(31) Number **60/045915** (32) Date **1997.05.07** (33) Country **US**

(43) Publication Date : **1998.11.27**
(43) Publication Journal Date : **1999.01.21**
(44) Accepted Journal Date : **2000.12.14**


(71) Applicant(s)
Landmark Graphics Corporation

(72) Inventor(s)
Ira D. Hale

(74) Agent/Attorney
DAVIES COLLISON CAVE, 1 Little Collins Street, MELBOURNE VIC 3000

(56) Related Art
US 5297236
US 4754492

4 OPI DATE 27/11/98 APPLN. ID 73714/98
 AOJP DATE 21/01/99 PCT NUMBER PCT/US98/09290 
 AU9873714

(51) International Patent Classification ⁶ : G06T		(11) International Publication Number: WO 98/50886
		(43) International Publication Date: 12 November 1998 (12.11.98)
(21) International Application Number: PCT/US98/09290 (22) International Filing Date: 6 May 1998 (06.05.98) (30) Priority Data: 60/045,915 7 May 1997 (07.05.97) US		(81) Designated States: AL, AU, BA, BB, BG, BR, CA, CN, CU, CZ, EE, GE, HU, IL, IS, JP, KP, KR, LC, LK, LR, LT, LV, MG, MK, MN, MX, NO, NZ, PL, RO, SG, SI, SK, TR, TT, UA, UZ, VN, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG). Published <i>Without international search report and to be republished upon receipt of that report.</i>
(71) Applicant: LANDMARK GRAPHICS CORPORATION [US/US]; 15150 Memorial Drive, Houston, TX 77079-4304 (US). (72) Inventor: HALE, Ira, D.; 8003 South Oneida Court, Englewood, CO 80112 (US). (74) Agent: BUSH, Gary, L.; Bush, Riddle & Jackson, L.L.P., Suite 180, 950 Echo Lane, Houston, TX 77024 (US).		
(54) Title: METHOD FOR DATA COMPRESSION		
		
(57) Abstract		
<p>The widely used JPEG standard algorithm for two-dimensional image compression may be adapted for compression of arrays of any dimension and data type, specifically for arrays of seismic data. Because the JPEG algorithm processes, more or less independently, small subsets (8 x 8 blocks) of larger images or arrays of data, such adaptations are particularly useful in applications that cannot maintain a large, uncompressed, multi-dimensional array in computer memory. JPEG-like methods enable compression and decompression of large arrays by iteration over sub-arrays that are small enough to reside in memory. These algorithms lead to the concept of a <i>compressed virtual memory</i>. Special care must be taken in JPEG-like algorithms to avoid blocking artifacts, which are discontinuities between blocks of data that are compressed and decompressed independently. Fortunately, computationally efficient methods for suppressing these artifacts are well known. Of these methods, one has been adopted that enables much of the JPEG method to be reused. The JPEG-like method of the invention uses the JPEG methods for discrete cosine transform (although the forward and inverse transforms are reversed), and for Huffman encoding of the quantized transform coefficients. The method differs from JPEGs primarily in additional steps taken to avoid blocking artifacts, and in the quantization of transform coefficients.</p>		

- 1 -

METHOD FOR DATA COMPRESSION

BACKGROUND OF THE INVENTION

5

Field of the Invention

The invention relates generally to data compression methods and in particular to an improvement of the JPEG method for image compression as it is applied to seismic data.

10

Description of the prior art

The JPEG standard algorithm for image compression (Pennebaker and Mitchell, 1993) consists of the following three steps, performed for each 8 x 8 block of pixels in a two-dimensional array:

15

1. Transform the 8 x 8 block of pixels, using a discrete cosine transform.
2. Quantize (scale and round) the transform coefficients into small



integers.

3. Encode the bits, using few bits to represent the most frequent integers.

The decompression algorithm inverts each of these steps, in reverse order. Both
5 algorithms may be easily extended to compress and decompress arrays of any dimension.

Figure 1 displays a 2-D array of seismic data that has not been compressed. The 2-D array of 32-bit floating-point numbers of Figure 1 is a constant-time slice extracted from a 3-D seismic survey. Figure 2 displays a
10 zoomed subset of the same array.

Figure 3 shows the same zoomed subset, after compression and decompression of the entire 2-D array using a JPEG-like algorithm. The compression ratio for the entire array is approximately 103:1, meaning that the original array of 32-bit floating-point numbers contained about 103 times as many
15 bits as the compressed array.

For such large compression ratios, this JPEG-like algorithm produces the blocking artifacts visible in Figure 3. At lower compression ratios, these discontinuities between blocks become less visible, but they may still be significant, particularly when further processing or interpretation is performed
20 after decompression.

The artifacts in Figure 3 are the result of each block of 8 x 8 samples being compressed and decompressed, independently, with no attempt to maintain

- 3 -

continuity between the blocks.

In spite of these artifacts, the ability to compress and decompress such small subsets of data independently is a desirable feature. In particular, it enables access to some small part of a large compressed array, without decompressing 5 the entire array. It also enables the compression algorithm to adapt to spatial variations in data amplitude and spectrum. These features are lacking in compression methods based on wavelet transforms (eg., Bradley et al., 1993; Wickerhauser, 1994). The problem addressed in this application is to obtain these features, without the blocking artifacts.

10 One solution to this problem is to compress data using overlapping blocks, so that decompressed sample values can be computed without reference to values adjacent to block boundaries. This solution was used by Yeo and Liu (1995), in their adaptation of the JPEG algorithm to volume rendering of 3-D medical images. Unfortunately, the use of overlapping blocks increases the 15 number of blocks to be compressed, which increases computation times and decreases compression ratios.

References to prior work cited in this specification

20 Bradley, J. N., C. M., and Hopper, T., 1993, The FBI wavelet/scalar quantization standard for gray-scale fingerprint image compression: Visual Information Processing II, SPIE Proceedings, 293-304. (<ftp://ftp.c3.lanl.gov/pub/WSQ.>)

25 Jawerth, B., and Sweldens, W., 1995, Biorthogonal smooth local trigonometric bases: J. Fourier Anal. Appl., 2. (<http://cm.bell-labs.com/who/wim/papers/-papers.html>).

Jawerth, B., Liu, Y., and Sweldens, W., 1996, Signal compression with smooth local trigonometric bases: <http://cm.bell-labs.com/who/wim/papers/-papers.html>.



- 4 -

Malvar, H. S., and Staelin, D. H., 1989, The LOT – transform coding without blocking effects: IEEE Transactions on Acoustic, Speech, and Signal Processing, 37, no. 4, 553-559.

5 Malvar, H. S., 1990, Lapped transforms for efficient transform/subband coding: IEEE Transactions on Acoustic, Speech, and Signal Processing, 38, no. 6, 969-978.

Pennebaker, W. B., and Mitchell, J. L., 1993, JPEG still image data compression standard: Van Nostrand Reinhold.

10 Princen, J. P., and Bradley, A. B., 1986, Analysis/synthesis filter bank design based on time domain aliasing cancellation: IEEE Transaction on Acoustics, Speech, and Signal Processing, 34, no. 5, 1153-1161.

Wickerhauser, M. V., 1994, Adapted wavelet analysis from theory to software: A. K. Peters.

15 Yeo, B., and Liu, B., 1995, Volume rendering of DCT-based compressed 3D scalar data: IEEE Transactions on Visualization and Computer Graphics, 1, no. 1, 29 – 43.

The reference to any prior art in this specification is not, and should not be taken as, an acknowledgment or any form of suggestion that that prior art forms part of the common general knowledge in Australia.

20



SUMMARY OF THE INVENTION

The present invention provides a method for compressing a one-dimensional array x of N samples representative of characteristics of the earth, comprising the steps of:

5 (a) partitioning said array x into blocks of M samples, where $M < N$,
 (b) folding the samples $x(j) \equiv x(M+j)$ across each block boundary l ,
 according to

$$y(-j) = f(j) x(j) + f(-j) x(-j),$$

$$y(-j) = f(j) x(-j) - f(-j) x(j),$$

$$l = 1, 2, \dots, N/M - 1,$$

$$j = 1, 2, \dots, M/2 - 1,$$

$$y(j) = x(j), \text{ otherwise,}$$

where

$$f(j) = \sin \left[\frac{\pi}{4} \left(1 + \frac{2j}{M} \right) \right],$$

15 (c) transforming the folded samples in each block of array y according to

$$z(k) = \sum_{j=0}^{M-1} y(j) b(j) \sqrt{\frac{2}{M}} \cos \left[\frac{\pi(2k+1)j}{2M} \right],$$

$$k=0,1,\dots,M-1,$$

20

where

$$b(j) = \begin{cases} \frac{1}{\sqrt{2}}, & j = 0 \\ 1, & \text{otherwise} \end{cases}$$

(d) quantizing the transformed samples in each block of array z to produce integers, and



- 5A -

(e) encoding said integers into a stream of bits representing the compressed array.

A preferred embodiment of the invention relates to a method of data compression which uses JPEG methods for discrete cosine transform and for 5 Huffman encoding of the quantized transform coefficients. The method is an improvement over standard JPEG methods primarily in the additional steps to avoid blocking artifacts, and in the quantization of transform coefficients.

Embodiments of the invention are described hereinafter in detail, by way of example only, with reference to the accompanying drawings.

10

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 represents a constant-time slice extracted from a 3-D seismic survey, where the display is a 2-D array of 32-bit floating-point numbers which has not been compressed;

15

Figure 2 is a zoomed subset of the 2-D array of Figure 1;

Figure 3 shows data of Figure 2 after compression and decompression via a straightforward adaptation of the JPEG algorithm, showing the blocking artifacts, the discontinuities between the 8 x 8 sample-blocks that were compressed independently;

20

Figure 4 shows the data of Figure 2 after compression and decompression via a straightforward adaptation of the JPEG algorithm that has been designed to suppress blocking artifacts;

25

Figure 5 shows a matrix C_{III}^{-1} corresponding to an inverse DCT-III of length 32 where black pixels correspond to positive values; white pixels correspond to negative values;

Figure 6 shows a matrix corresponding to an inverse blocked DCT-III with block length 8;



Figure 7 shows columns 11 and 19 from the inverse blocked DCT-III matrix of Figure 6, where the blocking artifacts in JPEG compression are caused by discontinuities such as that shown between samples 15 and 16, and the first sample in each cosine lies off the curve plotted and in this figure because of the discontinuity in the function $b(j)$ defined by equation (1c) below;

Figure 8 shows smooth tapered cosines obtained by unfolding the truncated cosines displayed in Figure 7 which are columns 11 and 19 of the inverse folded-DCT-III matrix of Figure 9;

Figure 9 shows a matrix corresponding to an inverse folded-DCT-III where the columns of this matrix overlap from block to block, with values that smoothly taper to zero such that weighted sum of these columns will not exhibit blocking artifacts;

Figure 10 shows a matrix corresponding to the operation used to unfold the inverse blocked DCT-III, where the matrix in Figure 9 equals the product of this matrix and that displayed in Figure 6;

Figure 11 shows an unfolding operation which is a pair-wise combination and replacement of sample values across DCT block boundaries and where folding is the inverse of this operation;

Figure 12 illustrates the steps to decompress the highlighted sample in block A, which requires that one first decode, dequantize, and inverse DCT-III blocks A, B, C, and D, and then unfold the four highlighted samples where the folding and unfolding operations are a pair-wise mixing of samples across DCT

block boundaries.

DESCRIPTION OF THE PREFERRED EMBODIMENT OF THE INVENTION

Embodiments of the invention relate to an improved method of data compression based on blending data in adjacent blocks during compression with a JPEG-like algorithm. JPEG-like data compression has been described in various forms by many authors (e.g., Princen and Bradley, 1986; Malvar and Staelin, 1989; Malvar, 1990; Wickerhauser, 1994, Jawerth et al., 1995). The advantage of JPEG-like data compression is that it tends to increase compression ratios, and only slightly increases computation times.

Figure 4 displays the same subset from the 2-D array of Figure 1, after compression and decompression using a second JPEG-like algorithm based on the method of an embodiment of the invention described below. The compression ratio for the entire array is approximately 112:1. Blocking artifacts in these decompressed data are absent.

The difference between the prior JPEG-like algorithms and the new JPEG-like algorithm according to the invention lies in a modification of the blocked discrete cosine transform specified in the JPEG standard. This modification in the algorithm of this invention attenuates blocking artifacts, while retaining the desirable features of the prior JPEG-like algorithm

DISCRETE COSINE TRANSFORM



The transform used prior JPEG compression is a discrete cosine transform, called DCT-II. According to the invention, a different discrete cosine transform is used, called DCT-III. (For an enumeration of discrete cosine transforms, see Wickerhauser, 1994, p. 84.)

5 DCT-III

The forward DCT-III is defined by

$$z(k) = \sum_{j=0}^{M-1} y(j)b(j)\sqrt{\frac{2}{M}}\cos\left[\frac{\pi(2k+1)j}{2M}\right],$$

$$k = 0, 1, \dots, M-1, \quad (1a)$$

and the corresponding inverse transform is

$$y(j) = \sum_{k=0}^{M-1} z(k)b(j)\sqrt{\frac{2}{M}}\cos\left[\frac{\pi(2k+1)j}{2M}\right],$$

$$j = 0, 1, \dots, M-1, \quad (1b)$$

where

- 8 -

(1c)

SUBSTITUTE SHEET (RULE 26)

$$b(j) = \begin{cases} \frac{1}{\sqrt{2}}, & j=0, \\ 1, & \text{otherwise.} \end{cases}$$

The forward and inverse transforms can be represented as matrix multiplications, as in $z = C_{II}y$ and $y = C_{II}^{-1}z$, where the matrix C_{II} has elements

$$C_{II}(k,j) = b(j) \sqrt{\frac{2}{M}} \cos \left[\frac{\pi(2k+1)j}{2M} \right]. \quad (2)$$

The inverse transform matrix C_{II}^{-1} is illustrated in Figure 5, $M = 32$. Any vector of 32 real numbers can be represented as a weighted sum of the columns of this matrix.

In a one-dimensional compression algorithm, a vector y of sample values might be approximated with only a few columns from the matrix C_{II}^{-1} . The weights for each column would be given by transform coefficients in the vector z . High compression ratios require that many of the coefficients in z are negligible, nearly zero. Such small coefficients will be quantized to zeros, which may be efficiently encoded using few bits.

The JPEG standard describes compression of two-dimensional images, not one-dimensional vectors. The prior DCT-II used in JPEG compression is a 2-D transform. However, because the discrete cosine transform of multi-dimensional data may be performed as a cascade of 1-D transforms along each data dimension, only the 1-D transform is described for simplicity.

Blocked DCT-III

For long data vectors, such as seismic traces, a single discrete cosine transform of the entire vector is unlikely to yield many transform coefficients that are negligible. Therefore, like JPEG, a blocked DCT is used, with transform length $M = 8$, and pad data vectors are set to zeros, as necessary, to a length N that is a multiple of 8. The matrix corresponding to an inverse blocked DCT-III is illustrated in Figure 6, for $N = 32$ and $M = 8$. Besides being more appropriate for compression, the blocked transform is also more efficient, with a computational cost that grows only linearly with the length N of data vectors.

Like the prior DCT-II used in JPEG compression, the DCT-III has several useful properties. First, a DCT-III transforms real numbers to real numbers, so that no complex arithmetic is required. Also, like the prior DCT-II, DCT-III is a unitary transform: $C_{III}^{-1} = C_{III}^T$. [This property is implied by equations (1) above.] The forward blocked DCT-III matrix corresponding to Figure 6 is simply the transpose of the inverse blocked DCT-III displayed there.

Another useful property is that an inverse DCT-III is equivalent to a forward DCT-II: $C_{III}^{-1} = C_{II}$. This relationship between the DCT-III and the DCT-II, and a choice of transform length $M = 8$, enables use of the highly optimized DCT-II algorithms used in JPEG compression, by simply swapping the forward and inverse algorithms.

All of the properties cited above are useful, but why not simply use the prior JPEG's DCT-II? The answer lies in a modification according to the invention of the DCT-III to avoid the blocking artifacts caused by JPEG compression.

Blocking artifacts appear when only a subset of the columns of the matrix shown in Figure 6 is used in an approximation of a vector. Suppose, for example, that such a vector is highly compressed by using only two columns with indices $k = 11$ and $k = 19$. As illustrated in Figure 7, any non-trivial combination of these two columns yields a discontinuity between sample indices $j = 15$ and $j = 16$ in the approximation. Such discontinuities create the blocking artifacts visible in images that have been compressed using the JPEG algorithm.

Folded DCT-III

To avoid the artifacts caused by using a blocked DCT-III in compression, the blocked cosines illustrated in Figure 7 is replaced with the smooth tapered cosines illustrated in Figure 8. Compare the complete inverse transform matrix displayed in Figure 9 with the inverse blocked DCT-III matrix of Figure 6. The number of samples in each cosine has increased to 16 (except at the ends) so that cosines in adjacent blocks now overlap, and that each cosine smoothly tapers to zero. A weighted sum of these cosines will not produce the discontinuity shown in Figure 7.

The transforms corresponding to these tapered cosines are often called

local cosine transforms or lapped orthogonal transforms (Wickerhauser, 1994, p. 370; there is a similarity between Figure 8 and Wickerhauser's Figure 11.5). Here, these transforms are referred to as *folded cosine transforms*, to reflect the manner in which they are computed. Specifically, the transform used in the compression algorithm according to the invention is a folded-DCT-III. The inverse transform corresponding to the matrix shown in Figure 9 is called an inverse folded-DCT-III.

5 Folding is a way to achieve smooth, tapered, 16-sample cosines using highly optimized, blocked, $M = 8$ DCT algorithms. Wickerhauser (1994, p. 103),
10 describes this method as "a remarkable observation made independently by several individuals", and goes on to discuss its application to compression. The folding operation used in compression in this invention is one of many described by Wickerhauser, but was inspired by the work of Jawerth and Sweldens (1995) and Jawerth et al. (1996). The latter authors discuss aspects of folding that are
15 particularly relevant to compression.

6 The cleverness and efficiency of folding lies in the fact that the inverse folded-DCT-III matrix displayed in Figure 9 is the product of an unfolding matrix, shown in Figure 10, and the inverse blocked DCT-III matrix displayed in Figure
20 6. Because each row and column of the unfolding matrix contains no more than two non-zero elements, the computational cost of unfolding is an almost insignificant addition to the cost of the inverse blocked DCT-III.

In practice, the matrices displayed in Figures 6, 9, or 10 are never actually

constructed. Rather, operations are performed on sample values that have the same effect as multiplication by these matrices. For the inverse blocked DCT-III matrix of Figure 6, this operation is the highly efficient forward blocked DCT-II algorithm used in JPEG compression. For the unfolding matrix of Figure 10, this 5 operation is a pair-wise mixing of sample values across the block boundaries at indices $j = 8, 16, \dots$, as illustrated in Figure 11.

The folding operation is simply the inverse of the unfolding operation, a different pair-wise mixing of the same sample values. Because folding and unfolding are centered about block boundaries, these operations are most 10 concisely specified in terms of shifted sample values, which are defined by $y_l(j) \equiv y(lM + j)$. The symbol l is a block index and j is a sample-within-a-block index.

Then, folding is accomplished via

$$\begin{aligned}
 y_l(j) &= f(j)x_l(j) + f(-j)x_l(-j), \\
 y_l(-j) &= f(j)x_l(-j) - f(-j)x_l(j), \\
 l &= 1, 2, \dots, N/M - 1, \\
 j &= 1, 2, \dots, M/2 - 1, \\
 y_l(j) &= x_l(j), \quad \text{otherwise,}
 \end{aligned} \tag{3a}$$

and unfold is accomplished via

$$x_l(j) = f(j)y_l(j) - f(-j)y_l(-j),$$

$$x_l(-j) = f(j)y_l(-j) + f(-j)y_l(j),$$

$$l = 1, 2, \dots, N/M - 1,$$

$$j = 1, 2, \dots, M/2 - 1,$$

$$x_l(j) = y_l(j), \quad \text{otherwise,} \quad (3b)$$

where the folding function $f(j)$ is defined by

$$f(j) = \sin \left[\frac{\pi}{4} \left(1 + \frac{2j}{M} \right) \right], \quad (3c)$$

The smooth tapered cosines in Figure 8 and the matrix of such cosines in Figure 9 were computed by applying equations (3b) to the columns of the matrix displayed in Figure 6.

The forward folded-DCT-III of sampled values $x(j)$ is determined by first computing $y(j)$ via equations (3a), and then, for each block of 8 samples, computing $z(k)$ via equation (1a). Likewise, the inverse folded-DCT-III is determined by first computing $y(j)$ via equation (1b) for each block, and then computing $x(j)$ via equations (3b).

10

As discussed by Wickerhauser (1994, p. 105), Jawerth and Sweldens

(1995) and Jawerth et al. (1996), many alternative but similar folding operations are possible. The folding operation defined by equations (3) is chosen for two reasons.

First, the folding operation is unitary. The unfolding matrix shown in Figure 5 10 is the transpose of the corresponding folding matrix (not shown). To verify this property, observe that the folding function of equation (3c) satisfies $f^2(j) + f^2(-j) = 1$, and then express the folding and unfolding operations of equations (3a) and (3b) as multiplications by 2×2 matrices. Analytically invert the 2×2 folding matrix to see that its inverse equals its transpose, which equals the 2×2 10 unfolding matrix. The entire folding operation is therefore unitary, because it consists of these 2×2 mixings of sample values across block boundaries.

Second, the folding operation according to the invention ensures that a constant function, such as $x(j) = 1$, yields transform coefficients $z_{11}(k)$ in each block that are non-zero for only $k=0$, thereby enhancing compression of constant (or 15 slowly varying) data. In the terminology of Jawerth et al. (1996), the forward folded-DCT-III achieves "*resolution of the constants*".

To verify this second property of the folded-DCT-III, analytically apply the folding operation of equations (3a) to constant samples values $x_{11}(j) = 1$, and verify that the result is $y_{11}(j) = \sqrt{M}C_{111}(0,j)$, where $C_{111}(k,j)$ is defined in equation (2). 20 In other words, the folding function is chosen so that constant sample values are folded to precisely match (to within scale factor \sqrt{M}) the first ($k = 0$) cosine of the DCT-III. Because this cosine is orthogonal to all other cosines of that transform,

only the $k = 0$ transform coefficient in each block will be non-zero after the forward folded-DCT-III. The value of each non-zero coefficient will be $z_i(0) = \sqrt{M}$.

More precisely, this property holds for all but the first and last blocks.

Although resolution of the constants for the first and last blocks can be obtained 5 by modifying the folding and unfolding matrices so that they are not unitary in their corners, less compression is accepted in these blocks and a strictly unitary folding and unfolding operation is maintained.

Comparison with DCT-II algorithm used in JPEG compression

The forward DCT-II used in JPEG compression achieves resolution of the 10 constants without folding, because the first cosine of DCT-II (e.g., the first row of the matrix in Figure 5) is constant. Indeed, this property is the reason why DCT-II, not DCT-III, is specified in the JPEG compression standard. Using the folding trick with DCT-II, might be considered, because folding before JPEG compression and unfolding after JPEG decompression might be performed, 15 without modifying the standard JPEG algorithm.

Unfortunately, the folding and unfolding operations of equations (3) are not appropriate for DCT-II. Specifically, equations (3b) would not yield smooth tapered cosines, like those displayed in Figure 8.

Jawerth et al. (1995) describe alternative folding and unfolding operations 20 that are appropriate for DCT-II, and that maintain resolution of the constants. However, these operations are not unitary. They are in fact ill-conditioned,

tending to amplify discontinuities in sample values that occur near block boundaries, and thereby reduce the efficiency of compression.

In contrast, the folding and unfolding operations of equations (3) are unitary, as are the forward and inverse DCT-III of equations (1). If F denotes the 5 folding matrix, then the forward folded-DCT-III can be expressed as $z = C_{III} F x$, and observe that

$$\begin{aligned}
 z^T z &= x^T F^T C_{III}^T C_{III} F x \\
 &= x^T F^{-1} C_{III}^{-1} C_{III} F x \\
 &= x^T x \\
 \sum_{k=0}^{N-1} z^2(k) &= \sum_{j=0}^{N-1} x^2(j). \tag{4}
 \end{aligned}$$

In other words, the sum of squared sample values after folded-DCT-III equals that sum before folded-DCT-III. This property may be used to estimate the distortion in decompressed data caused by quantizing the transform coefficients.

QUANTIZATION AND ENCODING

As described above, the transform used in the compression algorithm of the invention is just the inverse of that used in JPEG compression, but with folding included to reduce blocking artifacts. The quantization and encoding methods used are also adapted from those used in JPEG compression. The

differences between the JPEG methods and the invention are described below.

Quantization

Because JPEG compression is intended for images, data before compression are either 8-bit or 12-bit values. (Color images are represented by red-green-blue triplets of such samples.) After a 2-D transform via DCT-II, 8-bit data may require up to 11 bits per sample, because the largest sample value after 2-D DCT-II is up to 8 times greater than that before the transform. In general, the largest sample value after DCT-II is up to $M^{D/2}$ times the largest value before transform, where D is the number of dimensions transformed. In other words,

$$|z|_{\max} \leq M^{D/2} |x|_{\max}. \quad (5)$$

This upper bound is achieved when the data exactly match one of the cosines used in the DCT-II, such as the constant first cosine $C_{00}(k=0,j)$.

The same factor $M^{D/2}$ also applies to the folded-DCT-III used in the compression algorithm of this invention. Therefore, after transforming 8-bit integer data with a folded-DCT-III, the method of the invention quantizes and encodes 11-bit integers in 2-D compression, and 13-bit integers in 3-D compression.

The quantization step in both the JPEG compression algorithm and the method of the invention is a scaling operation designed to reduce the number of

bits to be encoded, and it is this reduction in the number of bits that is responsible for the distortion in data that are compressed and then decompressed. In such lossy compression methods this distortion is accepted in return for high compression ratios.

5 In JPEG compression, the transform coefficients within each block are quantized differently, with high wavenumbers (spatial frequencies) being represented with fewer bits than low wavenumbers. This wavenumber-dependent scaling in JPEG compression is usually optimized for human visual perception.

10 In the compression method of the invention, wavenumbers are all quantized equally. Errors are not introduced that are wavenumber-dependent. One reason for that is that for seismic data, there is often high interest in high wavenumbers. For example, subsurface faults imaged with seismic data correspond to high wavenumbers. Another reason is that seismic data are often analyzed by computer algorithms independent of the human visual system.

15 Local quantization

When compressing 8-bit image data, sample values lie between -128 and +127; and it may be assumed that low-amplitude blocks of data are insignificant and may be safely quantized to zero during compression. The JPEG standard, in particular, makes this assumption, for it permits only one set of quantization scale factors for an entire image. As discussed above, these scale factors may vary for different coefficients (different wavenumbers) within a block, but the

same set is used for every block. In this sense, JPEG quantization is *global*.

When compressing 32-bit floating-point data, it is preferred to perform a *local* quantization, with scale factors that vary from block to block. Before compressing such data, the maximum sample amplitude $|x|_{\max}$ might not be known, and reading every sample before compression to determine this value may be costly. Furthermore, low-amplitudes may not be properly assumed to be insignificant; seismic data, in particular, often require considerable processing before this assumption is valid. Therefore, although a single scale factor is used to quantize all transform coefficients within a block, a scale factor may be permitted to vary from block to block. Specifically, in equation (5) above, $|z|_{\max}$ denotes the maximum coefficient within each transformed block, and a different scale factor s is computed for each block.

Local quantization yields lower compression ratios (produces more bits-per-sample) than global quantization. An obvious reason is that additional bits are required to store the quantization scale factors for each block compressed. A less obvious reason is that local quantization may quantize fewer samples to zero than global quantization. Therefore, a choice of either local or global quantization is permitted in the compression algorithm of the invention.

The local quantization does not handle large dynamic range within a single block. For example, low-amplitude reflections in unprocessed seismic data may be hidden beneath high-amplitude surface waves. Within blocks contaminated with high-amplitude noise, the compression algorithm of the invention, even with

local quantization, may quantize low-amplitude signal to zero, so that the signal could not be recovered by processing after decompression. Therefore, high-amplitude noise should be attenuated before compression.

Compressed virtual memory

5 Virtual memory presents the illusion of a memory larger than that physically available. This illusion is most effective for applications that tend to access data located near other data that were recently accessed. Such applications have good *locality of reference*.

10 Applications working with 2-D or 3-D arrays often exhibit good locality of reference. For example, a seismic interpretation application may display consecutive 2-D slices of data from a 3-D array. With a sufficiently fast, local decompression algorithm, blocks may be decompressed containing the samples for such slices on demand, without decompressing the entire 3-D array. The compression algorithm of the invention is especially useful in such applications.

15 Key to such applications is the ability to compress or decompress a subset of a large array, without compressing or decompressing the entire array. For compression based on a blocked DCT-II, like that used in JPEG compression, this feature comes easily. Specifically, to decompress a single sample, it is only necessary to decode, dequantize, and inverse DCT-II the block containing that 20 sample. Once decompression for one sample is accomplished, decompression of all of the samples of its block is accomplished. Assuming locality of reference,

the computational cost of decompressing the other samples in that block will not be wasted.

For the compression algorithm of the invention, based on a folded-DCT-III, some additional work is required. Consider 2-D compression and the four blocks 5 of 8 x 8 samples illustrated in Figure (12). In order to decompress the sample corresponding to the filled circle in block A, it is required to (1) decode, dequantize, and inverse DCT-III all four blocks (A, B, C, and D), and (2) unfold the four samples corresponding to the filled circles. The unfolding across block boundaries is just as illustrated in Figure 11 and described by equation (3b), 10 performed first for one dimension and then for the other. Although four blocks are required, once decompression of the sample in block A is accomplished most of the work necessary to decompress neighboring samples has been performed. Again, assuming locality of reference, this additional computation will not be wasted.

15 Another difference between the quantization step of the invention and JPEG's stems from our need to quantize 32-bit floating-point data, which have a much larger dynamic range than 8- or 12-bit image data. To quantize a floating-point value z into an integer i with $B+1$ bits (including a sign bit), the following algorithm is used

$$i = \begin{cases} \lfloor z \times s + 1/2 \rfloor, & z \geq 0, \\ \lceil z \times s - 1/2 \rceil, & z < 0, \end{cases} \quad (6)$$

where s is the quantization scale factor. To avoid overflow, $|i| < 2^B$ is required. This restriction and equation (6) lead to the following equation for the scale factor:

$$s = \frac{(2^B - 1/2)(1 - \epsilon)}{|z|_{\max}} \quad (7)$$

where ϵ is float-epsilon, the smallest positive number that can be subtracted from 5 1, using floating-point arithmetic, to obtain a number not equal to 1.

Huffman encoding

The JPEG compression standard permits two methods for encoding the integers produced by quantization --- Huffman coding and arithmetic coding --- and many other encoding methods are possible. The JPEG's Huffman encoding 10 is used in the compression algorithm of the invention, because it is computationally fast, is simple to implement, and is freely available.

With one exception, the Huffman encoding and decoding algorithms follow the JPEG specification. The JPEG standard specifies special encoding of the DC ($k=0$) transform coefficients. JPEG compression exploits the fact that these DC 15 coefficients are often highly correlated among neighboring blocks. This special treatment is not used in the method of the invention, for two reasons: (1) it

introduces block-to-block dependencies that complicate compression and decompression of single blocks, and (2) seismic data typically have relatively small DC coefficients. (Although small, the DC coefficients are seldom negligible, because of the short $M=8$ transform used.) Therefore, the DC coefficient is
5 encoded just as the other (AC) coefficients are encoded.

LOCAL COMPRESSION

The primary advantage of a JPEG-like algorithm is that part of a multi-dimensional array can be compressed or used without processing all of it. In contrast, compression algorithms based on wavelet transforms (e.g., Bradley
10 et al., 1993; Wickerhauser, 1994) lack this feature. While the JPEG standard does not explicitly support this feature, the blocked DCT-II used in JPEG compression makes it feasible. This capability is exploited in the compression algorithm of the invention based on a folded-DCT-III.

Each block of data in an array may be thought of as analogous to a page
15 of virtual memory. For 2-D compression, each page would contain $64 = 8 \times 8$ samples; for 3-D compression, each page would contain $512 = 8 \times 8 \times 8$ samples. Samples are decompressed as they are paged in, and wavenumber compressed as they are paged out wavenumber. A working set of uncompressed pages are kept in memory, while most of the pages remain
20 compressed and stored either in memory or on disk. (If stored on disk, pages may be combined to enhance I/O efficiency.) If the working set is large enough,

- 25 -

and if our application has good locality of reference, the computation cost of compressing and/or decompressing each page will be amortized over the access of most of the samples within those pages.

5

SUMMING UP

The adaptation of the JPEG compression algorithm according to embodiments of the invention enables much of that algorithm to be reused. The JPEG methods are reused for length-8 DCTs, by simply swapping the forward and inverse transforms. The JPEG quantization method is modified to avoid preferential treatment of low wavenumbers, and to handle block-to-block variations in data amplitudes. Sample values are also folded across block boundaries, before a forward DCT, and such values are unfolded after an inverse DCT. This folding and unfolding suppresses the blocking artifacts visible in images that have been compressed with the JPEG algorithm.

10 For comparing the performance of the method of the invention with other algorithms, two useful measures of performance are computation time and distortion, for a specified compression ratio. Preliminary benchmarks with an early implementation of the algorithm of the invention are encouraging. The computation times are roughly half and distortions almost identical to those for a wavelet-based algorithm, for a wide range of compression ratios (Diller 1997, personal communication). The computation times are believed to be lower for the method of the invention than for wavelet-based methods (due to fewer multiplies and adds, and more local use of memory).

15

20

Throughout this specification and the claims which follow, unless the context requires otherwise, the word "comprise", and variations such as "comprises" and "comprising", will be understood to imply the inclusion of a stated integer or step or group of integers or steps but not the exclusion of any other integer or step or group of integers or steps.

[The next page is page 27]



WHAT IS CLAIMED IS:

1 1. A method for compressing a one-dimensional array x of N samples
2 representative of characteristics of the earth, comprising the steps of:
3 (a) partitioning said array x into blocks of M samples, where $M < N$,
4 (b) folding the samples $x_i(j) = x(iM + j)$ across each block boundary i ,
5 according to

6 $y_i(-j) = f(j) x_i(j) + f(-j) x_i(-j),$

7 $y_i(j) = f(j) x_i(-j) - f(-j) x_i(j),$

8 $i = 1, 2, \dots, N/M - 1,$

9 $j = 1, 2, \dots, M/2 - 1,$

10 $y_i(j) = x_i(j), \text{ otherwise,}$

11 where

$$f(j) = \sin \left[\frac{\pi}{4} \left(1 + \frac{2j}{M} \right) \right],$$

12 (c) transforming the folded samples in each block of array y according to

$$z(k) = \sum_{j=0}^{M-1} y(j) b(j) \sqrt{\frac{2}{M}} \cos \left[\frac{\pi(2k + 1)j}{2M} \right],$$

13 $k = 0, 1, \dots, M - 1,$

14

15 where

16

$$b(j) = \begin{cases} \frac{1}{\sqrt{2}}, & j = 0 \\ 1, & \text{otherwise} \end{cases}$$

17

18 (d) quantizing the transformed samples in each block of array z to produce integers, and

19 (e) encoding said integers into a stream of bits representing the 20 compressed array.

1 2. The method of Claim 1, applied to a multi-dimensional array,
2 wherein the (a) block, (b) fold, and (c) transform steps are applied as a cascade
3 of operations along each array dimension.

1 3. The method of Claim 2, wherein a subset of said multi-dimensional
2 array is compressed.

1 4. The method of Claim 3, wherein said multi-dimensional array
2 represents seismic signals.

3 5. The method of Claim 1, further comprising decompressing said
4 compressed array by inverting steps (a) through (e) in reverse order.

- 29 -

6. The method of claim 5, applied to a multi-dimensional array, wherein the inverse of the (a) block, (b) fold, and (c) transform steps are applied as a cascade of operations along each array dimension.

5 7. The method of claim 6, wherein a subset of said multi-dimensional array is decompressed.

8. The method of claim 7, wherein said multi-dimensional array represents seismic signals.

10

9. A method of compressing a one-dimensional array substantially as hereinbefore described with reference to Figures 4 to 12.

15 DATED this 13th day of October 2000

Landmark Graphics Corporation

By its Patent Attorneys

DAVIES COLLISON CAVE



FIG. 1





FIG. 2



SUBSTITUTE SHEET (RULE 26)




FIG. 3.




FIG. 4.




FIG. 5.




FIG. 6.

SUBSTITUTE SHEET (RULE 26)




FIG. 7.




FIG. 8.




FIG. 9.




FIG. 10.





FIG. 11.

FIG. 12.
SUBSTITUTE SHEET (RULE 26)