Defence Research Establishment Pacific
Centre de recherches pour la défense pacifique

[redacted]

DEFENCE RESEARCH ESTABLISHMENT PACIFIC
Research and Development Branch
Department of National Defence

Canada

Defence Research Establishment Pacific
Centre de recherches pour la défense pacifique

DEFENCE RESEARCH ESTABLISHMENT PACIFIC

CFB Esquimalt, FMO Victoria, B.C. V0S 1B0

DREP Technical Memorandum 94-146

Block truncation compression of sonar images

Garry J. Heard

December 12, 1994

Approved By:

Chief DREP

Research and Development Branch
Department of National Defence

Canadä

Block truncation compression of sonar images

Garry J. Heard

Defence Research Establishment Pacific, FMO Esquimalt, B.C. Canada, V0S 1B0

604/363-2905, fax 604/363-2856, heard@orca.drep.dnd.ca

(December 12, 1994)

Abstract

This paper describes the Block Truncation Compression algorithm and applies it to the problem of storing sonar lofargram and frequency-azimuth (FRAZ) images. Four simulation test cases are presented that demonstrate the algorithm's capabilities, and two real data examples are shown that illustrate the use of the method with actual lofargram and FRAZ images. Only qualitative comparisons of the original and restored images are given, as the technique greatly changes the picture details, but maintains the original image's appearance. This compression algorithm has desirable features such as data independence, minimal computation, and fixed compress/decompress processing times that suit it ideally to real-time applications. Compression ratios of 4:1 are achieved with the basic implementation of the algorithm.

KEYWORDS: Compression, Coding, Sonar, Lofargram, Displays, Image Processing, BTC

Typeset using REVTEX

INTRODUCTION

In this article, block truncation compression (BTC) is discussed and applied to the storage of time-frequency spectral images commonly called lofargrams by the underwater acoustics community. Lofargrams appear in many research related fields including underwater acoustics, sonar, seismology, and signal processing. In addition to lofargrams, BTC is also applied to the closely-related intensity-vs-frequency and bearing display commonly called a FRAZ (frequency-azimuth).

Many sonar systems make use of lofargrams as the primary display of output information. Lofargrams come in all sizes, but a typical display encountered in a research application might consist of 1024 by 512 pixels each requiring 1 byte of storage (*i.e.*, 256 colours or grey levels). This 512 kByte image usually represents just one directional beam of the sonar system, and often may not represent the entire time-history or frequency bandwidth that the sonar is capable of. To store a sonar's entire output can require a vast amount of storage. As an example, consider the Canadian Ocean Acoustic Measurement System (COAMS)[1]. This 2-km-long towed array sonar system continuously forms 64 beam-lofargrams, each in several frequency resolutions, and keeps 6 hours of lofargram information on-line. The total disk storage required for this system's lofargram images is almost 4 Gbytes. Currently the images are stored in an uncompressed format, but if the BTC scheme described in this paper were to be adopted, then that same amount of disk storage would be sufficient for 24 hours of information. Alternatively, BTC would allow 1 Gbyte to store the 6 hours of data currently on-line and would allow for such benefits as multiple redundancy or reduced size and power consumption. The need for effective data compression schemes is also apparent when one considers the archival storage of lofargrams and the transmission of lofargram data from place to place in one machine, or from machine to machine.

In today's technological world, memory and disk sizes are continuously increasing, while the costs of these devices are continuously decreasing. Despite this trend to ever larger and cheaper mass storage systems, applications and data storage requirements are pressing the limits of current technology. In the future, a means of storing and moving large amounts of data on devices with limited storage and bandwidth capabilities will become ever more desirable for an increasing number of applications. Compression technology is one answer to these needs. Today, compression technology can be found in many everyday appliances such as the video game, phone, computer, fax, and television. Despite these common applications, compression is seldom used in the day-to-day research workplace. This is probably due to a lack of familiarity with compression methods by researchers in other fields.

In the next section the BTC method and some of its features relevant to the current application are described, then in the following section the BTC algorithm is applied to a number of simulated and real data sets, and finally, in the last section the results of the investigation into the use of BTC with sonar images are briefly summarized.

BLOCK TRUNCATION COMPRESSION

In this section, the block truncation compression algorithm[2,3] is explained. It is assumed that the technique is to be applied to an image consisting of 8-bit numbers representing the colours (or grey-levels) in the original image. The BTC algorithm divides the original image into *super-pixels* or blocks. In most implementations, each block consists of 16 pixels arranged in a 4x4 array. Each pixel in a block is inspected and if the value of the pixel is greater than or equal to the mean of pixel values in the block, then a 1 is written to the corresponding bit of a 1-bit quantization representation (*i.e.*, 16 bits representing the original 16 bytes) of the block. If the value of the pixel is less than the mean, then a 0 is written into the corresponding bit of the 1-bit quantization. The original pixels are replaced by only 2 values (0 and 1) that in turn represent two colour values, a and b. The original 16 bytes of the 4x4 block are reduced to 4 bytes. Two bytes are used for the 1-bit quantization map of the original block, and one byte is used for each of a and b.

The values, a and b, are chosen to maintain the original image's first and second moments in each block. The original image's first moment, or mean, is defined as

$$m_1 = \frac{1}{n} \sum_{i=0}^{n-1} p_i. \tag{1}$$

where p_i is the value of the i^{th} pixel in a block, and n is the number of pixels in each block. The original image's second moment, or variance, is defined as

$$m_2 = \frac{1}{n} \sum_{i=0}^{n-1} p_i^2. \tag{2}$$

The mean of the compressed image is given by

$$m_{1c} = \frac{1}{n} \left[(n-q)\, a + qb \right], \tag{3}$$

where q is the number of pixels in the block with values greater than or equal to the mean. The variance of the compressed image is given by

$$m_{2c} = \frac{1}{n} \left[(n-q)\, a^2 + qb^2 \right]. \tag{4}$$

By equating Equation 1 with Equation 3, and Equation 2 with Equation 4, and solving for a and b we obtain

$$a = m_1 - \sqrt{\frac{q}{n-q}(m_2 - m_1^2)}, \tag{5}$$

and

$$b = m_1 + \sqrt{\frac{n-q}{q}(m_2 - m_1^2)}. \tag{6}$$

Two special cases must be considered. First, when all pixels in the block are the same, set $a = b = m_1$. This prevents a divide-by-zero error when calculating a and b. Second, it is possible to arrive at values of a and b that cannot be represented by an unsigned 8-bit integer. When this situation occurs, simply set the value to 0 or 255, whichever is the closest.

In order to decompress an image it is simply necessary to read the 1-bit quantization map and the a and b values for each block in the original picture; wherever a 0 appears in the 1-bit quantization write the value a in the corresponding element of the 4x4 block and wherever a 1 appears in the 1-bit quantization, write the value b in the corresponding block element. In this way, the image is reconstructed 4 rows at a time using only 2 of the possible 256 colours in any given block. From this last statement it is obvious that BTC is a lossy compression algorithm. Despite this loss of picture information, the human eye often has trouble telling the difference between the original and reconstructed images.

The current application is to use BTC with lofargram and FRAZ images that are commonly generated by sonar systems. Figure 1 is an example of a lofargram generated for a simulated 100 Hz tone embedded in white Gaussian noise. In many lofargrams, spectral normalization[4] is applied to remove broadband information. Normalized images generally consist of a flat background of random noise represented by a speckled image with an overall light-grey colour, while the signal is represented by darker pixels. Constant frequency signals appear as straight dark lines parallel to the time axis of the figure, which is usually oriented vertically. Experience has shown that approximately 75 pixels/inch is an optimum resolution for lofargram images. At this plotting resolution individual pixels are just resolved from a comfortable viewing distance by a person with normal eyesight.

Perhaps the most common complaint against BTC is that it can introduce a block shaped texture in reconstructed images where there is a smooth gradient in the original image. One approach to solving this problem has been to use smoothing filters that take advantage of the correlation between a pixel and its nearest neighbours[5]. In lofargrams, blockiness should not a be problem, since broadband information that might give rise to a gradient in the image colour is usually removed by the spectral normalizer. In addition, experience has shown that a characteristic *blockiness* is desirable in lofargrams. In lofargrams, blockiness often occurs when a number of signals clustered near a common frequency are under-resolved by the spectral analysis. In many cases, attempting to further resolve such features can lead to increased confusion of the sonar operator.

BTC has excellent line and edge preservation[2], a feature of considerable importance when dealing with lofargram images, where the features of primary interest are generally lines from continuous frequency signals. Any compression technique that is unable to preserve narrow, possibly single-pixel-wide, lines is unsuitable for use with lofargram images.

BTC is also very resistant to errors[2]. This is of considerable interest when images must be transmitted to remote receivers. It is fairly easy to see why BTC is error resistant. First, the original image is broken into blocks, so individual errors will affect only one block of the image. A single bit error in the 1-bit quantization will affect only one pixel, and a bit error in a or b will only affect one block. Errors are therefore restricted to a block of the reconstructed image, which represents only a tiny fraction of the entire image. As the error rate is increased, the degradation will affect more and more blocks, but will do so at a nearly linear rate.

BTC requires very little computation. Modestly sized computers can compress and decompress large images quickly. In fact, real-time applications are possible with specialized integrated circuits[2,6]. In its most basic implementation, the compression ratio, compression time, and decompression time of BTC are independent of the image. This feature allows for very straight forward estimates of timing and storage requirements, issues that can be of extreme importance when real-time processing is desired.

Finally, many enhancements and modifications of the basic BTC algorithm are possible. One of the simplest modifications involves varying the threshold level that controls the assignment of a or b to a given pixel. Another modification is to preserve the third moment in addition to the first two moments. Some authors modify the 1-bit quantization image: for example Siedband *et al.*[5] employ a two-dimensional smoothing operator that relies on the correlation between pixels to improve the reconstructed image quality. BTC can also be used as one stage of a multi-step process. For example, Neagoe[7] employs BTC to encode the error image resulting from a previous application of feedback transform coding (FTC), and Wu and Coll[8] use BTC as the first stage of a three level process consisting of BTC followed by vector quantization (VQ) and the adaptive discrete cosine transform (ADCT). In this paper, only the basic algorithm has been investigated. The main reason for this decision is that a real-time application is the driving force, and the basic algorithm has the desirable features of simplicity (minimal computation) and well defined fixed processing times regardless of the input data. Once experience with the basic algorithm has been obtained, enhancements may be tried in the future. The following section presents examples of simulated and real images processed with the basic BTC algorithm.

TEST CASES AND REAL DATA

In this section, four test cases and two real data examples are presented that show the effectiveness of BTC for compression of the data contained in a sonar display. Each test case consists of a simulated lofargram composed of one or two tonal components (100 and 101 Hz) and normally distributed Gaussian noise. The first test case is that of a single tonal at 100 Hz with a signal-to-noise ratio (SNR) of −23 dB. In this article SNR is defined by

$$SNR = 10 \log \left(\frac{\sum_{i=1}^{L} b_i^2}{\sum_{i=1}^{L} n_i^2} \right), \tag{7}$$

where b_i and n_i are the tonal and noise sample amplitudes and L is the length of the test time-series ($L = 307,200$ samples).

In each case, the standard algorithm has been used without modification. Images were compressed using a 4x4 block size with a resulting compression ratio of 4:1. The original images were composed of 222 frequency bins and 300 time bins with 8-bit quantization. Each image in its original state requires 66,600 Bytes of storage. The compressed data files required only 16,654 Bytes. The four extra bytes are included at the beginning of each compressed file to store the original number of frequency and time bins.

Figure 1 shows the original lofargram image on the left-hand side and the image after compression and restoration with the BTC algorithm on the right-hand side. Close inspection of the original and restored images reveals that a slight loss of information has occurred, but that in general the restored image retains all important details of the original image. Even the random patterns in the noise background have been preserved. Figure 2 shows the histogram of the original image as a dashed-line and the histogram of the restored image as a solid-line. Note that the BTC algorithm induces a bi-modal distribution of the image pixels. Basically, this figure indicates that the restored image is composed of just two levels, a high-mean value and a low-mean value. The bi-modal distribution preserves the first and second moments of the original image. Empirical results indicate that the mean and variance are preserved to within 2% accuracy for the four test cases shown in this paper. An interesting consequence of this bi-modal distribution of image pixels is that images that have been compressed and restored with BTC tend to photocopy and fax better than the original images. Note also that isolated high-value outliers in the original image are removed from the restored image, but that their loss is not directly detectable by the human eye. Instead, this loss of outliers is responsible for an overall shift in the brightness of the restored images, the restored images generally appearing darker (when printed with PostScript®). In order to compare the original and restored images fairly, a simple re-normalization of the original image is made so that the maximum value in both images is the same.

The results from the second test case are shown in Fig. 3. In this case, two equal intensity lines ($SNR = -23$ dB for each line) are separated by 1 Hz. It is important to note that at the spectral resolution employed for these test cases (0.684 Hz) the two signals are separated by just one spectral bin. Both signals will fall within the same block when the BTC algorithm is employed and, therefore, level variations in one signal will also effect the final appearance of the other signal. This situation has been chosen on purpose to provide a more demanding test of the BTC algorithm's capabilities. As for the previous test, the restored image compares very favourably with the original, although a slight degradation is seen as the lines in the restored image appear more broken than in the original.

The results from the third test case are shown in Fig. 4. In this case, the line at 100 Hz is weak ($SNR = -23$ dB), while the line at 101 Hz is appreciably stronger ($SNR = -17$ dB). Once again the original and restored images compare favourably. Note that the stronger line appears marginally stronger and the weaker line marginally weaker in the restored image than in the original image. This change in relative strength appears to be the result of the interaction of the signals in the same block. In addition to the change in relative levels, note also that some speckle has been introduced, particularly near the stronger line.

The results from the fourth test case are shown in Fig. 5. In this case, both lines are strong (−3 dB). Again, the comparison is favourable, but some artifacts appear as the increased speckle in the pixels between the two signals and the relatively uniform banding on each side of the strong signals.

The four test cases represent a broad range of input conditions, but do so in isolation. Real data, on the other hand, is composed of numerous signals, both narrow- and broadband in nature, with a multitude of weak and strong signals all present simultaneously. In order to examine the BTC algorithm's performance with real data, a section of a recording made during the Heard Island Feasibility Test[9,10] has been analyzed, passed through a split-window three pass normalizer[4] and displayed in Fig. 6. Several different types of signals are found in this lofargram. The Heard Island signal appears as the relatively weak line at 57 Hz. Another weak tonal appears at about 90 Hz, and some signals with broader frequency content appear in the frequency range 115–145 Hz. The original and restored images compare favourably, but for this real data case there are some notable differences that were not so apparent with the simulated data. First, notice that although the first and second moments of the original image have been preserved, the restored image is perceived to be darker. Secondly, notice that speckle has been increased in the restored image. Even isolated outliers in the original image generate a few dark pixels in the restored image. Despite these negative comments the restored image still contains the same information as the original image. A small change in the quantization levels would probably be quite effective at reducing the apparent changes in the restored image.

So far, we have considered only the compression of lofargram images. Sonar systems also commonly create other images known as FRAZ diagrams. These diagrams are closely related to the lofargram, but convey information about spectral content and source bearing rather than the frequency-time information of lofargrams. Figure 7 shows the effect of applying the BTC algorithm to another section of the Heard Island Feasibility Test data. In this example, the Heard Island signal occurs at a frequency of 57 Hz and a bearing of 132° Rel. The image is generally smoother than a lofargram image because of a smaller number of bearing pixels than time pixels, and the time averaging that is implied in a FRAZ diagram. The comparison between the original and restored FRAZ indicates that the image degradation is minimal. Some increased blockiness is noted in the restored image, but this could be minimized by a smoothing operation[5].

CONCLUSIONS

This paper has considered the application of BTC to the problem of storing sonar data in the form of lofargram and FRAZ images. The basic BTC algorithm is computationally simple and independent of the data presented to it. This feature together with the fixed processing times makes real-time applications relatively simple. Perhaps the biggest complaint about BTC is that the restored images tend to have a blocky appearance. This effect has been seen in the sonar images also, but here it is of less importance because of the low-resolution nature of the original images. BTC has several properties such as line preservation, error resistance, and single pixel resolution that also match this algorithm to the current application. In this work, only the basic algorithm has been investigated; however, many modifications to the algorithm have been tried and some of these together with new ideas, such as changing the block shape from square to rectangular to take advantage of the nature of the input images may result in better performance of the algorithm in similar future applications.

The four test cases were designed to study the performance of the BTC algorithm when varying combinations of weak and strong signals were present with noise in a single block simultaneously. It is desirable that a quantitative measure of the restored images usefulness be available; however, it is well known that various measures of error do not correlate well with human perception of the image quality. Because of this fact, only qualitative statements concerning the original and restored images have been made. In general, there is some apparent degradation of the images, but it is quite minimal in most cases. Perhaps the most severe losses occur when a weak and strong line are present in the same block. Under these conditions the weak line is perceived to be slightly weaker and the strong line slightly stronger than they really are. This should not be a restrictive problem as it occurs only when the lines are extremely close together. Under such circumstances the user would probably employ greater spectral resolution to increase the separation of the lines (*i.e.*, place them in different blocks).

When the BTC algorithm is applied to the real data examples, the results are found to be very encouraging. It may be possible to improve the appearances of both the lofargram and the FRAZ by application of a smoothing operation similar to that used by Siedband *et al.*[5]. This extra processing may turn out to be unnecessary; further experience with the technique will determine if this is the case.

The BTC algorithm is straight-forward and results in moderate compression ratios with data that would generally not compress well with exact methods. BTC should be given serious consideration for the archival and real-time storage of sonar data.

REFERENCES

[1] D.W. Craig, and L. McKee, "High Performance Shipboard Acoustics Data Acquisition, Imaging" Sea Technology **31**(1), 51–53, January (1990).

[2] E.J. Delp and O.R. Mitchell, "Image Compression Using Block Truncation Coding" IEEE Trans. on Communications **COM-27**(9), 1335–1342, September (1979).

[3] A. Kruger, "Block Truncation Compression" Dr. Dobb's Journal, 48–55, April (1992).

[4] W.A. Struzinski and E.D. Lowe, "A performance comparison of four noise background normalization schemes proposed for signal detection systems" J. Acoust. Soc. Am. **76**(6), 1738–1742, December (1984).

[5] M.P. Siedband, H. Zhang, J. Gray, "Block truncation image compression with bit map compensation" Computers in Physics, 52–56, Nov/Dec (1989).

[6] L.G. Chen, Y.C. Liu, T.D. Chiuch, Y.P. Lee, "A real-time video signal processing chip" IEEE Trans. on Consumer Electronics **39**(2), 82–92, May (1993).

[7] V.I. Neagoe, "Predictive Ordering Technique and Feedback Transform Coding for Data Compression of Still Pictures" IEEE Trans. on Communications **40**(2), 385–396, February (1992).

[8] Y. Wu and D.C. Coll, "BTC-VQ-DCT Hybrid Coding of Digital Images" IEEE Trans. on Communications **39**(9), 1283–1287, September (1991).

[9] W.H. Munk and A.M.G. Forbes, "Global Ocean Warming: An Acoustic Measure?" J. Phys. Oceanogr. **19**, 1765–1778 (1989).

[10] A.B. Baggeroer and W.H. Munk, "The Heard Island Feasibility Test," *Physics Today*, **45**(9), 22–30, 1992.

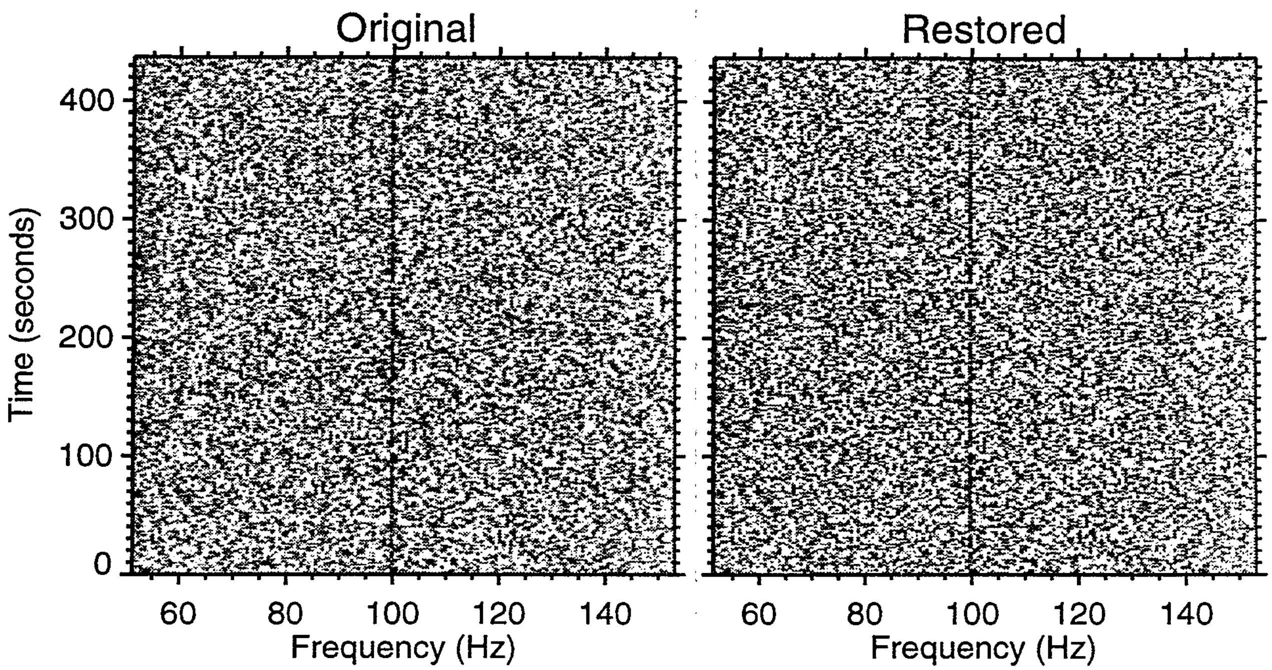

FIG. 1. Example lofargram for a simulated weak ($SNR = -23$ dB) 100 Hz tone in Gaussian white noise.

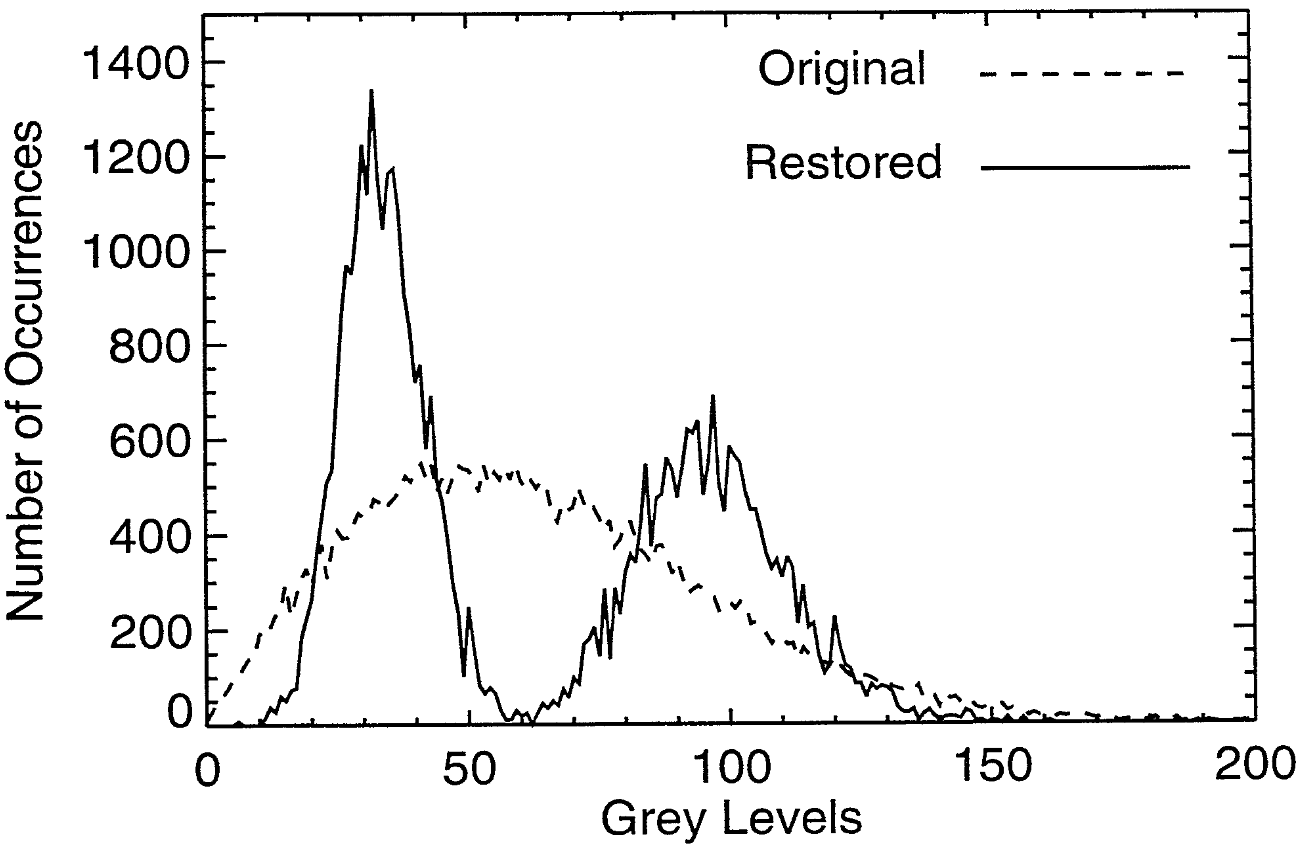

FIG. 2. Histogram of original and restored images for test case 1.

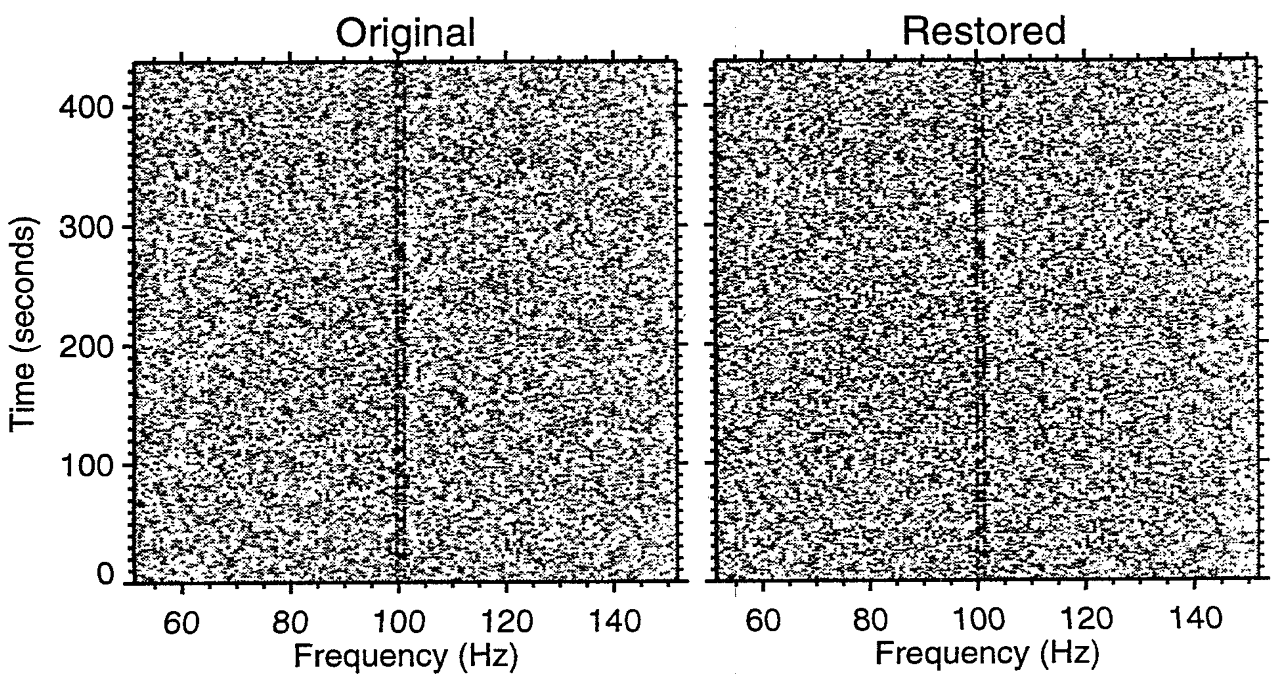

FIG. 3. Example lofargram for simulated equal strength ($SNR = -23$ dB) 100 Hz and 101 Hz tones in Gaussian white noise. Tonals are separated by only 1 pixel.

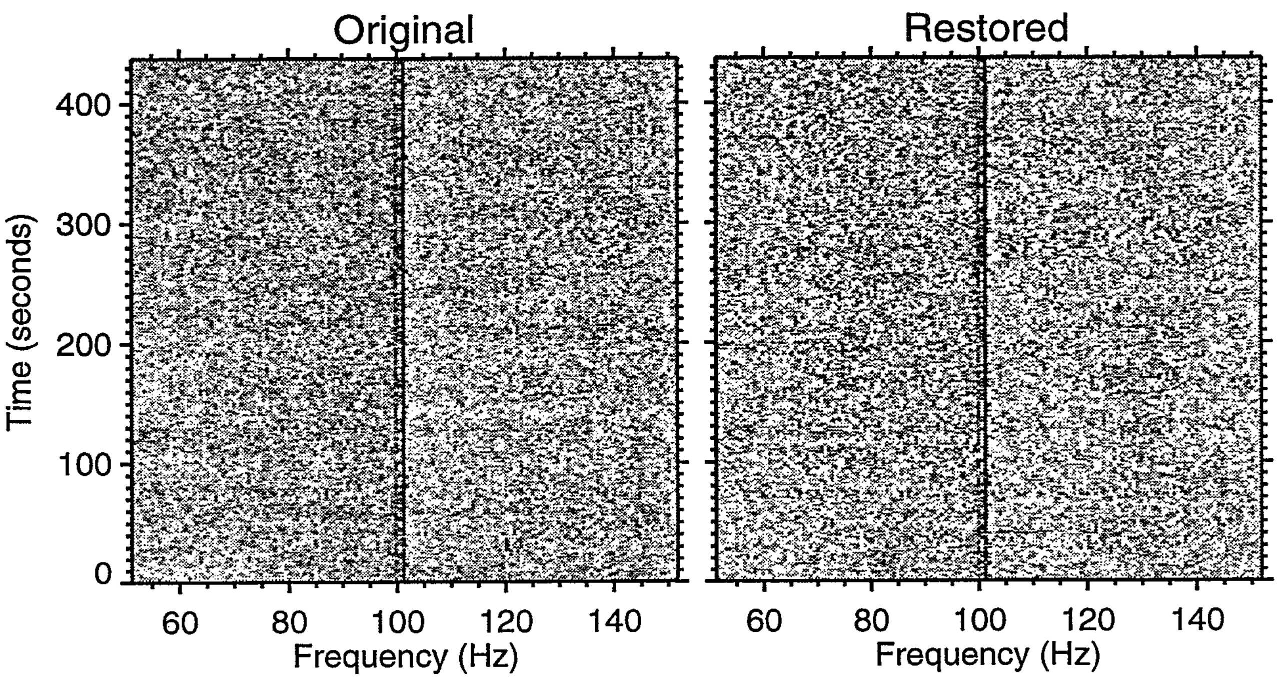

FIG. 4. Example lofargram for simulated weak ($SNR = -23$ dB) 100 Hz tone and a relatively strong ($SNR = -17$ dB) 101 Hz tone in Gaussian white noise. Tonals are separated by only 1 pixel.

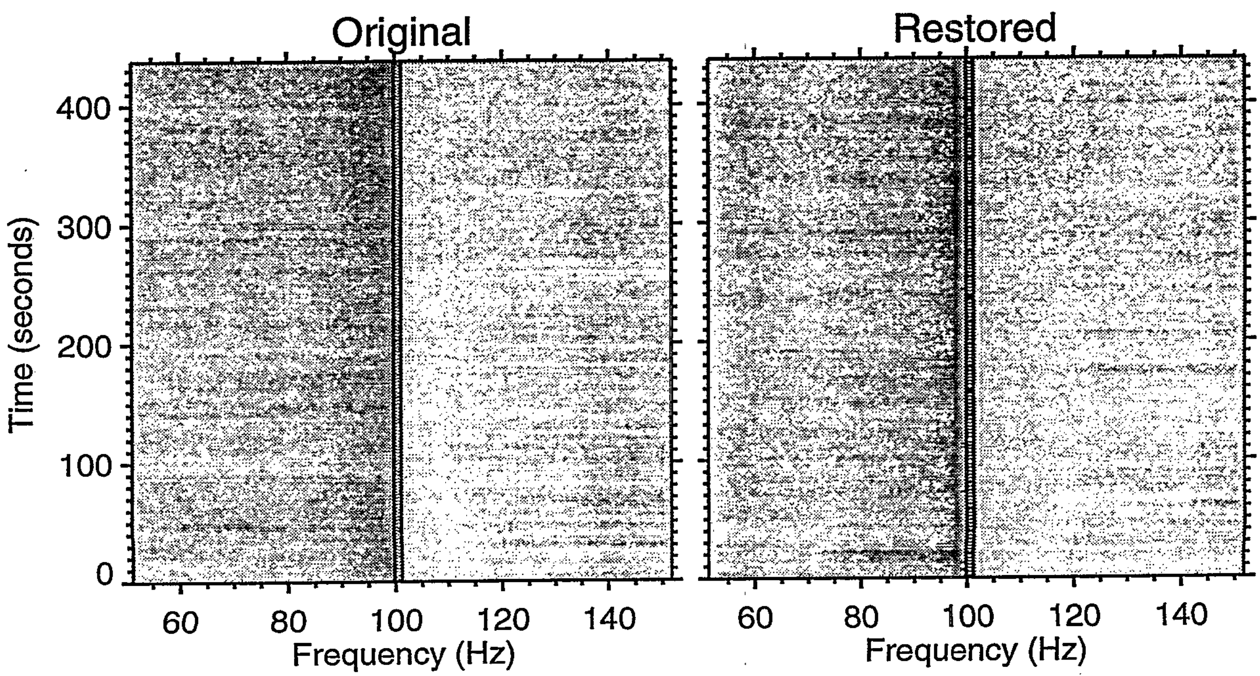

FIG. 5. Example lofargram for simulated strong equal strength ($SNR = -3$ dB) 100 Hz and 101 Hz tones in Gaussian white noise. Tonals are separated by only 1 pixel.

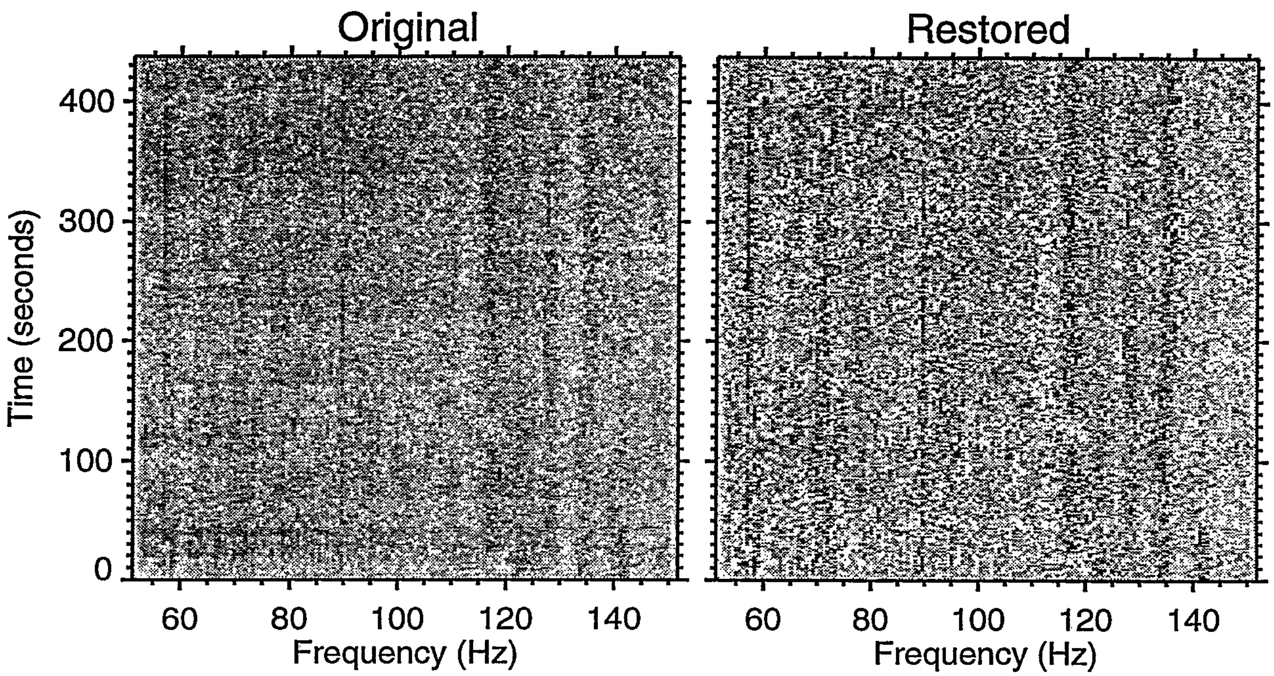

FIG. 6. Real data example lofargram created using Heard Island Feasibility Test data.

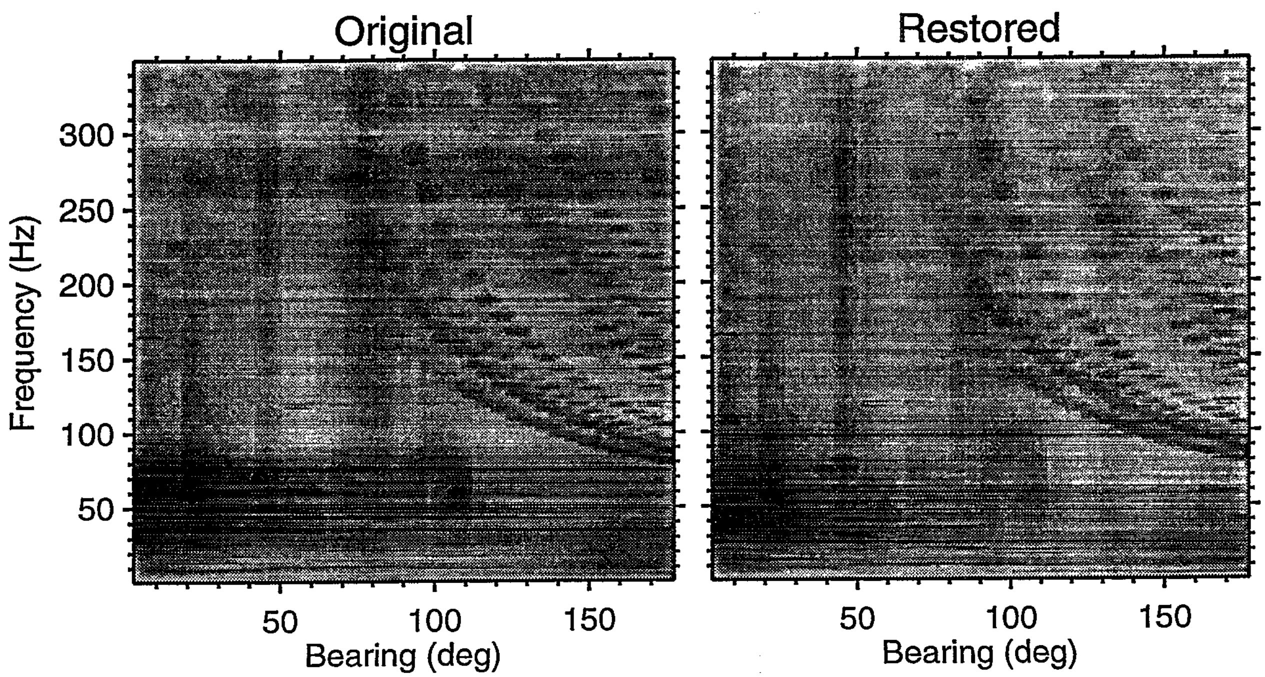

FIG. 7. Real data example FRAZ diagram, again created from Heard Island Feasibility Test data.

Distribution List

REPORT NO.: DREP TECHNICAL MEMORANDUM 94-~~112~~ 146

TITLE: BLOCK TRUNCATION COMPRESSION OF SONAR IMAGES

DATED: DECEMBER 1994

AUTHOR: GARRY J. HEARD

SECURITY GRADING UNCLASSIFIED-UNLIMITED

Distribution:

2 - DSIS

4 - DREA
- 2 - library
- 1 - Dr. J.B. Farrell
- 1 - Dr. I.A. Fraser

3 - DREP
- 2 - Dr. G.J. Heard
- 1 - Dr. N.R. Chapman

1 - CO ADAC, CFB Stadacona, Halifax

1 - CO ADAC, CFB Esquimalt, Victoria

UNCLASSIFIED

SECURITY CLASSIFICATION OF FORM
(highest classification of Title, Abstract, Keywords)

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (the name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Establishment sponsoring a contractor's report, or tasking agency, are entered in section 8.) Defence Research Establishment Pacific CFB Esquimalt, Victoria, BC, V0S 1B0	2. SECURITY CLASSIFICATION (overall security classification of the document, including special warning terms if applicable) unclassified

3. TITLE (the complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S,C or U) in parentheses after the title.)

Block truncation compression of sonar images (U)

4. AUTHORS (Last name, first name, middle initial)

Garry J. Heard

5. DATE OF PUBLICATION (month and year of publication of document) December 1994	6a. NO. OF PAGES (total containing information. Include Annexes, Appendices, etc.) 19	6b. NO. OF REFS (total cited in document) 10

7. DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Technical Memorandum

8. SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development. Include the address)

DREP

9a. PROJECT OR GRANT NO. (if appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant) DRDM-08	9b. CONTRACT NO. (if appropriate, the applicable number under which the document was written) n/a
10a. ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity. This number must be unique to this document.) DREP TM94-146	10b. OTHER DOCUMENT NOS. (Any other numbers which may be assigned this document either by the originator or by the sponsor) n/a

11. DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

(x) Unlimited distribution
( ) Distribution limited to defence departments and defence contractors; further distribution only as approved
( ) Distribution limited to defence departments and Canadian defence contractors; further distribution only as approved
( ) Distribution limited to government departments and agencies; further distribution only as approved
( ) Distribution limited to defence departments; further distribution only as approved
( ) Other (please specify):

12. DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in 11) is possible, a wider announcement audience may be selected.)

no limitation

UNCLASSIFIED

SECURITY CLASSIFICATION OF FORM

UNCLASSIFIED

SECURITY CLASSIFICATION OF FORM

13. ABSTRACT (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), or (U). It is not necessary to include here abstracts in both offical languages unless the text is bilingual).

This paper describes the Block Truncation Compression algorithm and applies it to the problem of storing sonar lofargram and frequency-azimuth (FRAZ) images. Four simulation test cases are presented that demonstrate the algorithm's capabilities, and two real data examples are shown that illustrate the use of the method with actual lofargram and FRAZ images. Only qualitative comparisons of the original and restored images are given, as the technique greatly changes the picture details, but maintains the original image's appearance. This compression algorithm has desirable features such as data independence, minimal computation, and fixed compress/decompress processing times that suit it ideally to real-time applications. Compression ratios of 4:1 are achieved with the basic implementation of the algorithm.

KEYWORDS: Compression, Coding, Sonar, Lofargram, Displays, Image Processing, BTC

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus. e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

COMPRESSION
CODING
SONAR
LOFARGRAM
DISPLAYS
IMAGE PROCESSING
BTC (Block truncation compression)