

O

AR-009-931

DSTO-TR-0443

F

Spatial Filtering of Sidescan
Sonograms to Remove Streaking due to
Signal Level Fluctuations

Stuart Anstee

S

DTIC QUALITY INSPECTED 4


APPROVED FOR PUBLIC RELEASE

19970307 092

© Commonwealth of Australia

F

DISCLAIMER NOTICE



THIS DOCUMENT IS BEST QUALITY AVAILABLE. THE COPY FURNISHED TO DTIC CONTAINED A SIGNIFICANT NUMBER OF PAGES WHICH DO NOT REPRODUCE LEGIBLY.

Mar-09

THE UNITED STATES NATIONAL
TECHNICAL INFORMATION SERVICE
IS AUTHORIZED TO
REPRODUCE AND SELL THIS REPORT

Spatial Filtering of Sidescan Sonograms to Remove Streaking due to Signal Level Fluctuations

Stuart Anstee

**Maritime Operations Division
Aeronautical and Maritime Research Laboratory**

DSTO-TR-0443

ABSTRACT

A spatial filtering technique originally devised to filter streaking from geographically-scaled sidescan sonar images is applied successfully to smaller-scale sidescan sonar images generated during mine warfare exercises. Implementation of the algorithm to minimise computation time is discussed and appropriate filter parameter settings are established. Related techniques for image contrast and noise control are also discussed.

RELEASE LIMITATION

Approved for public release

D E P A R T M E N T O F D E F E N C E

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION

Published by

*DSTO Aeronautical and Maritime Research Laboratory
PO Box 4331
Melbourne Victoria 3001*

Telephone: (03) 9626 8111

Fax: (03) 9626 8999

© Commonwealth of Australia 1996

AR No. AR-009-931

November 1996

APPROVED FOR PUBLIC RELEASE

Spatial filtering of sidescan sonograms to remove streaking due to signal level fluctuations

Executive Summary

This report demonstrates the use of a filter to automatically remove streaking from sidescan sonograms collected during mine warfare exercises, particularly route surveys. Algorithms are given which would allow it to be implemented as a near real-time option in a "waterfall" sidescan sonar data review program running on existing, fast personal computer hardware. Examples of programs which could support such a modification are Triton's "Isis" and Elics' "Delph-Sonar" sonar data-analysis packages, which run on fast Pentium PC's. Algorithms are also given for related techniques for contrast enhancement and noise reduction which can also be implemented in real time on PC platforms.

Sidescan sonograms collected by the RAN during route surveillance operations are frequently streaked due to variations in signal level between successive sonar pings. The sonograms are intended to record the presence of any small, minelike objects along designated shipping routes. Streaks on the sonograms can degrade their usefulness and it is desirable to remove them. This report examines a process devised by Chervenka and de Moustier (1993) to reduce similar streaking observed in sidescan sonograms collected with a much larger sidescan sonar during geographical seafloor mapping operations. The filtration technique devised by Chervenka and de Moustier for sonograms of the seafloor kilometres wide is shown in this report to work well for sonograms 200 m wide collected during route surveillance. The main streak filtration technique consists of a two-stage low-pass filter implemented using a Chebychev expansion of the original sonar records. In a departure from the original work, the filter is implemented using a discrete cosine transform technique which allows the filtration to be performed sufficiently rapidly that it can be implemented as a real-time process on a cheap, standard PC platform (Intel Pentium 133 MHz in this case.)

A related technique demonstrated by Chervenka and de Moustier for the correction of systematic variations in image intensity with range is examined and found to be of limited use for the processing of route surveillance sonograms due to limitations in the quality of the data.

The use of techniques for contrast enhancement and noise reduction is also discussed. Histogram equalisation is shown to give useful contrast enhancement in some cases, and to be useful for the pre-processing of extremely noisy images in other cases. A simple and rapid thresholding filter is shown to reduce speckle noise in the sonogram and remove interference from a depth-sounder entirely, without reducing the clarity of the image.

Many example images are included in the report, showing the effects of the various filters on sonograms of shipwrecks and on two sonograms incorporating minelike objects.

Author

Stuart Anstee

Maritime Operations Division

Stuart Anstee is a Research Scientist in the Sydney branch of MOD. He holds a BSc in Applied Mathematics and Physics, an Honours degree in Physics and a PhD degree in Physics from the University of Queensland. His research concerns underwater acoustics, image processing and sonar design issues related to mine warfare.

Contents

1. INTRODUCTION	1
2. STREAK FILTRATION THEORY.....	6
2.1 Rationale for the process	6
2.2 Filtering with Chebychev polynomials.....	11
2.3 Spatial filtering across-track.....	12
2.4 Spatial filtering along track.....	12
2.5 The final processed image	14
2.6 The systematic correction.....	14
3. PRE- AND POST-PROCESSING TECHNIQUES.....	16
3.1 Degradation of low-pass filtration by outliers	16
3.2 Removal of localised outliers	17
3.2.1 A thresholding filter	17
3.2.2 Other outlier filters	20
3.3 Histogram equalisation	21
3.3.1 Processing to enhance contrast.....	23
3.3.2 Processing to reduce noise	24
3.4 Systematic correction	29
3.5 An example using both histogram equalisation and systematic correction	31
4. COMPUTATIONAL CONSIDERATIONS.....	31
4.1 Initial sampling.....	31
4.2 Calculation of Chebychev coefficients.....	32
4.3 Filtering of coefficients.....	33
4.4 Evaluation of the filtered image	33
4.5 Control of filter characteristics.....	34
4.6 Control of filter "over-correction".....	35
5. FILTRATION CHARACTERISTICS.....	36
5.1 Performance	36
5.2 Towards intelligent filtration.....	39
6. CONCLUSIONS	40
7. REFERENCES	40
8. ACKNOWLEDGMENTS	40

1. Introduction

A paper by Cervenka and de Moustier (1993), hereafter referred to as Paper I, describes the application of spatial filtering techniques to remove striping artefacts from sidescan sonograms collected with a large, low-frequency sidescan sonar with a swath-width of some 10 kilometres. The striping is most likely due to beam motions caused by towfish pitch and yaw, to inconsistencies in the triggering of time-varying gain and to cable degradation. Artefacts similar to those described in Paper I are seen in sidescan sonograms collected with a Klein 590 sidescan sonar which has a total swath width of only 200 m. The causes are the same and the effects are often particularly noticeable at the higher "500 kHz" frequency used for object classification.

Figures 1 and 2 show examples of streaked images containing objects of interest. In Figure 1, a 500 kHz sonogram, the object of interest is clearly a shipwreck and the attendant debris. The streaking effect is not strong enough to obscure significant details, however it detracts from the quality of the image. In Figure 2, a 100 kHz sonogram, the image is heavily streaked due to noise generated by a moving cable interconnection (a winch slip-ring) and it is not obvious to casual inspection that a mine is present in the image. In the case of the shipwreck image, filtration should ideally remove the streaking in the image without obscuring the smallest observable details, such as the two small objects to the right of the wreck and the scouring of the bottom on the other side of the track. It must also preserve the contrast. At the other extreme, the quality of Figure 2, which shows the starboard side of the swath only, is so bad that any process which allows the mine to be seen *at all* is worthwhile. Finally, Figure 3, a 100 kHz sonogram showing the port side of the swath, shows an image of a mine collected in good conditions. This image, which contains some strong but localised areas of noise due to interference from an echo-sounder, is intended to act as a "control image" for the filtration process, that is, the quality of the image must not be degraded by the filter.

The process described in Paper I is attractive in the context of minewarfare operations, because the spatial filtering method may be controlled so as not to introduce "foreign" artefacts into the processed sonogram. This means that an operator can be reasonably certain that mine-like objects appearing in the processed image are genuine.

In common with normal practice, the filtration process described in Paper I has three stages: pre-processing to remove outliers, filtration proper and post-processing to restore contrast. All three stages are discussed here, since all are important, although the pre- and post-processing techniques in this report are slightly different to those used in Paper I. An additional technique mentioned in Paper I which gives even image density across the image, the so-called "systematic correction" is also discussed, using both the technique of Paper I and another simpler technique.

In its original form, the streak-removal process described in Paper I is computationally intensive and can in some circumstances be too slow for "real-time" application on a PC platform, that is, it would slow the scrolling rate of a waterfall sonogram display to the point of uselessness. However, with some minor approximations, the computation time can be decreased by a large factor. For example, the filtration time for a screen-

full of data comprising 768 pings, which corresponds to approximately 107 seconds of real-time data collection, is of order 6 to 12 seconds when a PC with an Intel Pentium 133 MHz processor is used, depending on the nature of the data. This would therefore be a tolerable load on a "real-time" post-processing system, even running three times faster than the original data collection. The nature of the algorithm given in Paper I is such that it could be implemented as a continuous process, able to be toggled on or off with the press of a button, subject to an initial waiting period.




Figure 1. The original sonogram of a shipwreck recorded in Sydney Harbour using the 500 kHz frequency of a Klein 590 sidescan sonar. The total swath width is 200 metres. In this and subsequent images, the along-track direction is from right to left and up the page is across-track to starboard. The centre of the track is in this case halfway up the image.




Figure 2. The starboard side of an original sonogram recorded in Jervis Bay using the 100 kHz frequency of a Klein 590 sidescan sonar. The image is 100 m across, and contains a mine, which is obscured by streaking caused by either a bad electromechanical connection (a slip ring) or a cable breakdown.




Figure 3. The port side of an original sonogram of a sandy region containing a mine, collected in Jervis Bay in good conditions, using the 100 kHz frequency of a Klein 590 sidescan sonar. The image is 100 m across-track.

2. Streak Filtration Theory

2.1 Rationale for the process

The details of the filtering technique given in the following sections only describe a particular method of producing the desired filtering effect. In order to give an overview of the idea behind the process, images showing the effects of the steps in the process on the original shipwreck image are given here.

Considering Figure 1, the streaks appear in the image because each of the records from separate pings has relatively large-scale differences from its neighbours, for example the whole record may appear lighter or darker than its neighbours, or large stretches of the record may do so. The major streaking artefacts are relatively large, while the features in the image which are of most interest are generally much smaller. Thus, a low-pass spatial filtration along the length of each record should isolate the factors causing the streaking without affecting the finer details which make the image useful. Figure 4 shows the effect of low-pass filtration across the track on Figure 1. It is clear that most of the streaking in the image is contained in the low frequencies. When the low frequency image is subtracted out of the original image, the result is Figure 5, the "fine details" image, which is shown here with contrast artificially increased. In the absence of contrast enhancement, Figure 5 is faint because the higher spatial frequencies constitute a relatively small part of the total density of the image. Figure 5 is of little use by itself, since the visual cues and contrast of the original image are absent. In order to restore them, the low frequencies must be added back into the image, but without the streaking. In order to do this, the low-frequency image is low-pass filtered a second time, but this time in the along-track direction. The result is Figure 6, the "base" image, from which most of the streaking has been removed. When the "base" and "fine details" images are added together, the result is Figure 7, the "processed" image, which retains the detail of the original image, while most of the streaking is gone. The most important feature of this filtering technique is that, when implemented *with care*, it does not distort or obscure the objects in the original image.




Figure 4. The shipwreck sonogram (Figure 1), after low-pass filtering across-track. Filter parameters (see section 4) were $M=512$, $N=20$.



Figure 5. The "fine details" image, corresponding to the subtraction of the across-track filtered image (Figure 4) from the original sonogram (Figure 1).




Figure 6. The "base" image, corresponding to the across-track filtered sonogram (Figure 2) after low-pass filtration in the along-track direction.




Figure 7. The "processed" image corresponding to the original sonogram in Figure 1, which is the sum of the "fine details" image (Figure 5) and the "base" image (Figure 6).

2.2 Filtering with Chebychev polynomials

Following Paper I, denote the sidescan sonar image as a succession of records (represented as rows of pixels on the screen) described by *continuous* functions $s_y(x)$ where y , an *integer* variable, refers to the y th record (giving its relative vertical position on the screen) and the horizontal range x , a *continuous* variable, extends over the interval $[x_{\min}, x_{\max}]$. In general, there are two sides to the sidescan image, with each side described by a separate record. In this work, the two sides are assumed to be independent and to be processed independently, so only one side is considered. Note that in Section 4, which discusses the implementation of the algorithms in this section, the discrete nature of an actual sidescan image is taken into account via various approximations and modifications.

In principal, the filtering technique described in Section 2.1 could be implemented with any type of low-pass filter. Filtration by Chebychev expansion was chosen by Cervenka and de Moustier because of the well-known robustness and elegance of the technique. In order to perform the filtration, each record is expanded as a Fourier series over Chebychev polynomials

$$s_y(x) = \frac{a_0[y]}{2} + \sum_{j=1}^{\infty} a_j[y] T_j(X) \quad (2.1)$$

where square brackets $[]$ denote an integer index, round brackets $()$ denote a function evaluation, T_j is the j th Chebychev polynomial and the evaluation locations $X(x)$ for the Chebychev polynomials are calculated from the relation

$$X(x) = \frac{2x - (x_{\max} + x_{\min})}{x_{\max} - x_{\min}}. \quad (2.2)$$

The Chebychev (Fourier) coefficients $a_j[y]$, $j=0,1,\dots$ for the record are calculated from the inverse relation

$$a_j[y] = \frac{2}{M} \sum_{i=1}^M s_y(x_i) T_j(X_i) = \frac{2}{M} \sum_{i=1}^M s_y(x_i) \cos(j\theta_i) \quad (2.3)$$

noting the relation $T_j(X_i) = \cos(j\theta_i)$ and where M is the number of sample locations taken across the record and the x_i and θ_i correspond to the zeros of the polynomials, optimising the fit of the polynomial expansion to the original function (record). They are given by

$$\theta_i = \frac{i-1/2}{M} \pi \quad (2.4)$$

$$X_i = \cos(\theta_i) \quad (2.5)$$

$$x_i = 1/2[(x_{\max} + x_{\min}) + X_i(x_{\max} - x_{\min})] \quad (2.6)$$

for $i = 1, 2, \dots, M$.

2.3 Spatial filtering across-track

The advantage in using a Chebychev expansion like (2.1) is that the Chebychev polynomials T_j are an orthogonal series whose spatial frequency (number of oscillations in the interval $[x_{\min}, x_{\max}]$) increases with their index j . Thus, as j increases, the polynomial coefficients $a_j[y]$ correspond to progressively smaller details of the image and typically form an exponentially-decreasing sequence. Truncation of the expansion is equivalent to the application of a low-pass spatial filter to the function $s_y(x)$, and the exponential decay of the coefficients helps to ensure that the filter is well-behaved - a "windowing function" is not needed in this case.

Figure 4 shows the image obtained by calculating the Chebychev coefficients using equation (2.3) with $M=512$ and truncating the expansion at $N=20$. The filtered records shown in Figure 4 are obtained from the truncated expansion

$$\sigma_y(x) = \frac{a_0[y]}{2} + \sum_{j=1}^{N-1} a_j[y] T_j(X). \quad (2.7)$$

Figure 5 shows the "fine details" of Figure 1, that is, it shows the features which remain in the image when the low-pass filtered image Figure 4 has been subtracted out. The records shown in the image are defined by

$$\phi_y(x) = s_y(x) - \sigma_y(x). \quad (2.8)$$

Note that the features of Figure 5 are faint. As was noted previously, the magnitudes of the Chebychev coefficients decrease exponentially with index j , so that Figure 5, which includes only spatial frequencies corresponding to j values higher than 20, is a dim image.

2.4 Spatial filtering along track

The next stage of the filtering process is the construction of a "base" image which is low-pass filtered both across- and along-track, thereby ridding it of the streaking, which has a high spatial frequency in the along-track direction. One technique for achieving this might be to start from the low-pass filtered image as represented by Figure 4 and expand it as a Fourier series in the same manner as equation (2.1), but in the along-track direction. Low-pass filtration could then be achieved by truncation of the series in the same manner as equation (2.7). Such a technique is inefficient, however, since it requires two sets of forward and reverse transformations, and is impossible to implement as a continuous process. The technique described in Paper I for the along-track filtration therefore acts directly on the truncated Chebychev coefficients $a_j[y], j = 0, 1 \dots N$, rather than on the across-track filtered image. The low-pass filtration is achieved using a variable-width parabolic convolution mask applied along-track. This means that the Chebychev coefficients for the record are replaced by an along-track average with their nearest neighbours. The "parabolic" shape of the

convolution mask means that the coefficients nearest to the original record are given most weight, scaling down to zero weight at the edges of the mask. The smooth taper of the parabola is a windowing function, intended to stop filter truncation artefacts from appearing in the filtered image. The width of the mask is adjusted so as to decrease with increasing spatial frequency. Thus, a wide convolution mask is applied at low spatial frequencies (small j) corresponding to features which will probably be large along-track as well as across-track, whereas a progressively narrower mask is applied at higher spatial frequencies corresponding to features which would be expected to have smaller extent along-track. The Chebychev coefficients of the filtered image are given by the convolution

$$\bar{a}_j[y] = a_j[y] * b_j[y], \quad 0 \leq j < N \quad (2.9)$$

where the asterisk denotes the convolution operator. The convolution mask is given by

$$b_j[y] = \frac{6}{p_j(p_j+1)(p_j+2)} [(q_j+1)^2 - y^2], \quad -q_j \leq y \leq q_j \quad (2.10)$$

$$q_0 = N \quad (2.11)$$

$$q_j = (N-1) \operatorname{div} j, \quad 1 \leq j \leq N-1 \quad (2.12)$$

$$p_j = 2q_j + 1 \quad (2.13)$$

where "div" represents integer division. Note that the width of the mask scales inversely with the order of the polynomial, so it is proportional to the size of features which will typically be represented by the polynomial. Note also that equation (2.11) represents an arbitrary choice for the convolution acting on the zero-order coefficients, which appears to function correctly. An alternative, which is now used routinely in the filtration, is to replace all the zero-order coefficients by the same value, that is,

$$\bar{a}_0[y] = A_0 \quad (2.14)$$

where A_0 is arbitrary and is typically selected to be slightly below the midpoint of the available grey-scale values.

The records forming the "base image" are reconstructed from the filtered Chebychev coefficients $\bar{a}_j[y]$ by the mapping

$$\bar{\sigma}_y(x) = \frac{\bar{a}_0[y]}{2} + \sum_{j=1}^{N-1} \bar{a}_j[y] T_j(X) \quad (2.15)$$

and, as noted, an example of the result of the filtering is shown in Figure 6.

2.5 The final processed image

The final processed image is the sum of the "base image" (2.15, Figure 6) and the "fine-details" image (2.8, Figure 5), that is, the processed records $\bar{s}_y(x)$ are given by

$$\bar{s}_y(x) = s_y(x) - \sigma_y(x) + \bar{\sigma}_y(x) \quad (2.16)$$

or more realistically, by

$$\bar{s}_y(x) = s_y(x) + \Psi_y(x) \quad (2.17)$$

where the "processed part" Ψ_y is given by

$$\Psi_y(x) = \frac{\bar{a}_0[y]}{2} - \frac{a_0[y]}{2} + \sum_{j=1}^{N-1} (\bar{a}_j[y] - a_j[y]) T_j(X). \quad (2.18)$$

Considering Ψ_y , we see that the truncation of the expansion should prevent the filter from interfering with those features of the original image which are dependent on the very highest spatial frequencies, that is, the "fine details". Hence, edges and boundaries in the image should retain definition. The truncation of the expansion also implies that small errors in the value of x and hence X resulting from interpolation or decimation, should not seriously reduce the quality of the final image, because they will result in relatively small movements and distortion of relatively large features in the image. As noted, Figure 7 gives an example of the final processed image.

2.6 The systematic correction

In Paper I, Cervenka and de Moustier include an additional "systematic correction" which is designed to correct gain variations across the track which are consistent along the entire sonar run, for example, the black areas typically seen in sonograms close to the centre of the swath which may be due to an uncorrected lobe of the transducer beam pattern. Such variations are typically caused by a mismatch between an exponential time-varying gain function and a more complicated functional form for the fall-off of echo intensity with range.

The first step of the correction for the gain variation is to calculate a running average along-track of the pixel level at each range cell (position across-track), effectively,

$$f(x) = \frac{1}{2P+1} \sum_{y=-P}^P s_y(x). \quad (2.19)$$

The simplest correction possible (applied hitherto by the author) applies a smoothing filter to $f(x)$ (for example, a second running average going across-track) and subtracts it from $s_y(x)$. Finally, an arbitrary constant value is added to the record to raise the average pixel level and avoid negative values. This can be expressed as

$$s_y^{SC}(x) = \bar{f}(x) - F_0 \quad (2.20)$$

where F_0 is an arbitrary constant with a typical value half-way along the range of grey-scale values and $\bar{f}(x)$ is the across-track smoothed version of $f(x)$. The final form of the processed records is then given by

$$\hat{s}_y(x) = \bar{s}_y(x) - s_y^{SC}(x) \quad (2.21)$$

A more refined technique is given in Paper I, using a Chebychev approximant. In this case, a second set of Chebychev coefficients

$$\tilde{a}_j[y] = \frac{2}{M'} \sum_{i=1}^{M'} f(x_i) \cos(j\theta_i), \quad j=0,1\dots N'-1 \quad (2.22)$$

is calculated from the running average, after which a parabolic taper is applied to give the filtered coefficient set

$$\hat{a}_j[y] = \tilde{a}_j[y] \left(1 - \left[\frac{j-1}{N'-1} \right]^2 \right), \quad j=1,2\dots N'-1. \quad (2.23)$$

Finally, the systematic correction for the filtered record is given by

$$s_y^{SC}(x) = \sum_{j=1}^{N'-1} \hat{a}_j[y] T_j(X_i') \quad (2.24)$$

where obvious definitions apply and the final form of the records is given by (2.21). This correction is superior to subtraction of an arbitrarily-smoothed form of the running average as per (2.20), in that the zeroth-order Chebychev coefficient is not involved, so that the base level of the record is not affected. This means that an arbitrary constant need not be added to each record to avoid negative pixel values.

3. Pre- and post-processing techniques

In general, the pre-processing stage of the filtration process is meant to remove "outlier" data values which reduce the effectiveness of the streak filtration process. The processes described here do not assume any knowledge of the mechanism used to generate the image; rather, they act to remove artefacts from the image so that it is a better input to the next stage of filtration. The image is treated as a rectangular array of pixel values $S_y[k]$, $k = 1, 2, \dots, M_0$, where the integer y represents the row of the image and the integer k represents the column.

All the techniques described here for pre-processing can also be used for post-processing, since the processes which make data more suitable for filtration also make it more accessible to a human observer.

3.1 Degradation of low-pass filtration by outliers

Areas of very intense or very low signal level, whether isolated or continuous, can reduce the effectiveness of the streak-removal filter, or in fact any filter which uses low-pass filtration.

In the case of the streak-removal filter, when the coefficients of the Chebychev expansion are calculated from (2.3), the values of the low-order Chebychev coefficients are intended to represent the "background" on which other objects are perceived. If sufficiently many extreme pixel values are present, the low-order frequency components will be higher or lower than the "true background" that a human observer would perceive, since the human observer would ignore the extreme values whereas the expansion takes them into account. In an image reconstructed from the full Chebychev expansion, the distortion of the low frequency components to conform to extreme values does not matter, because the higher frequency components correct it. However, low-pass filtration retains only the low frequency components in the expansion. Visible effects in the reconstructed image include faint patches in the neighbourhood of strong reflectors like mines, spurious objects appearing in shadow regions and loss of contrast in regions with large amounts of isolated noise pixels. Pre-processing to remove extreme values can sometimes remove such effects.

It is also interesting to consider the "two-sigma" filter devised by Lee (1983) to reduce the incidence of isolated noise in synthetic aperture radar applications. The filter classifies a pixel as "noise" if its value is more than two standard deviations of the background pixel level away from the mean background pixel level. This criterion, applied globally, would also classify a typical mine signature as noise, along with other strong reflectors. It is thus clear that "outlier" values are a necessary feature of sidescan sonar images collected for mine hunting purposes, and that restraint in the removal of noise is a virtue.

3.2 Removal of localised outliers

3.2.1 A thresholding filter

Figure 3, which shows the image of a mine collected in good conditions, also demonstrates two kinds of isolated noise which can reduce the effectiveness of the streak-removal filter. Close examination of the image reveals a series of short, across-track "dashes", which all register at high intensity. They are the response of the sonar to the ship's echo-sounder. "Speckle" noise is also present in the image. Such noise is probably the result of constructive acoustic interference and is extremely localised. Figure 8 shows the normalised pixel level across a single sonar record from Figure 3. The record was selected because both the mine and depth-sounder signals register on it. The mine appears as a small clump of pixels at maximum intensity near range bin 230. The echo-sounder noise is a "spike" near range bin 310, to the right of the mine. Speckle noise is present as sharp, isolated peaks. Most speckle noise is restricted to a single pixel width along-track and at most a few pixels across-track. The echo-sounder noise has longer extent across-track, but remains one pixel wide in the along-track direction. Experience has shown that for the two types of noise shown, and for "line drop-outs" which also often affect only a single ping and therefore a single row of the image, a simple filtration procedure is very effective.

The thresholding filter, which is sometimes referred to as "1-dimensional, non-linear", is described by a simple algorithm:

- (i) Consider the pixel location (y,k) , which has level $S_y[k]$.
- (ii) Consider the two nearest neighbour pixels along-track, which have levels $S_{y-1}[k]$ and $S_{y+1}[k]$.
- (iii) Define an "action threshold" pixel level S^A . If either $S_{y-1}[k] > S^A$ or $S_{y+1}[k] > S^A$ then move to the next pixel, since the filter may be on the edge of an extended object.
- (iv) Define an "outlier margin" S^0 . If the difference of the central pixel away from the mean of its neighbours is greater than S^0 , replace it by their mean. That is, if $|0.5(S_{y-1}[k] + S_{y+1}[k]) - S_y[k]| > S^0$ then $S_y[k] = 0.5(S_{y-1}[k] + S_{y+1}[k])$.




Figure 8. Relative pixel values for a single sonar record from Figure 3 which was selected because it intersected the mine, giving a clump of pixels at maximum value near range bin 230, and a depth-sounder record, giving the vertical-appearing "spike" of pixels near range bin 310.

Figure 9 shows the result of the application of the thresholding filter to the Figure 3 image. The action threshold is half the maximum pixel level and the outlier margin is one quarter of the pixel level range. The definition of the image is unchanged, but the echo-sounder noise has been acceptably removed. Figure 10 shows the same single sonar record as Figure 8, after filtration. Note that the clump of pixels registering the mine is not affected, but the spike due to the depth sounder has gone, as have many of the higher-lying pixels. This is because the mine extends further than one sonar record, but the depth-sounder and speckle noise does not. The background level of the record remains unaffected. The thresholding filter is often applied by the author to sonar images before other processing and appears to have no detrimental effects.




Figure 9. The result of the application of the thresholding filter to Figure 3. The short across-track "dashes" of noise caused by interference from the ship's echo sounder have been removed. Many isolated high-intensity pixels have also been removed.




Figure 10. Relative pixel values from the same single sonar record as was shown in Figure 8, after the application of the thresholding filter, that is, as it would appear in Figure 9. After filtration, the clump of pixels at maximum value corresponding to the mine near range bin 230 is unaffected, but the high-value pixels near range-bin 310 corresponding to depth-sounder noise are mostly gone. Note that many high-lying pixels have been filtered out, but the background level is the same as in Figure 8.

3.2.2 Other outlier filters

In Paper I, Cervenka and De Moustier (and references therein) describe a median filter for the reduction of outliers which they claim works well. In his honours thesis on automatic classification of sea bottom types, Overkov (1995) reviewed many types of outlier reduction filters, including the median filter, for use with multibeam echo sounder data and concluded that the two-sigma filter described by Lee (1983) was the best suited to his application. Neither type of filter was used in the present case, since the thresholding filter gave adequate results. Trials of the effect of the two-sigma filter on images including speckle and minelike objects showed that, if the parameters of the filter were set correctly, it was very effective at reducing speckle and did not interfere with the appearance of the minelike objects. However, the filtration time for the two-sigma filter is much greater than for the thresholding filter, without resulting in an obviously superior image.

3.3 Histogram equalisation

"Histogram equalisation" is a technique based on the assumption that an image with high contrast will typically have an approximately uniform distribution of pixel intensities across the entire range without a peak at any particular value. Assuming this is true, then for a quantised grey-level display system, with for example 255 grey levels, the "level occupation distribution", that is, the total number of pixels in a given range of grey levels, should be approximately constant or "uniform". For example, there should be approximately the same number of pixels in the image registering from 0 to 19 and from 150 to 169. If most of the pixels in an image are clustered around a single value, then the histogram of pixel level occupation (the number of pixels at each discrete level) will be strongly peaked and there will be little contrast in the image.

For mine hunting purposes, a uniform pixel level distribution is not usually desirable. Useful contrast must be preserved on two scales: there must be sufficient contrast in the background pixel levels to be able to distinguish geological features on the sea floor, but the level typical of the background must also be much lower than is typical for the very small area of a minelike object. If the pixel levels of the background are of the same order as the levels of a minelike object, then that object will be very difficult to see.

The sonogram shown in Figure 3 is a "good" mine hunting image, collected in favourable conditions with the sonar properly "tuned", that is, with the gain controls properly adjusted. It does not have a uniform pixel level distribution, since the sonar is tuned such that the background registers at low intensity and *rare* minelike objects register much more strongly. The relative pixel level occupation histogram of Figure 3, which is shown in Figure 11, peaks at a low level and decays rapidly at higher levels. In Figure 11 and following plots of relative pixel level occupations, a value of "1" corresponds to the relative occupancy of a completely uniform distribution, that is, if the pixel values in the image were equally divided between all possible levels. In Figure 11 and following histogram plots, the relative occupancy of the 255th level is not shown, since "clipping" during digitisation means that the highest level may be over-represented. The histogram shown in Figure 11 is typical for sidescan sonograms, except when the image is cluttered with lots of strongly reflecting features or strong noise sources. If clutter is present, histogram equalisation can be useful, since it tends to increase the differentiation between pixel levels which are originally close together.




Figure 11. Relative pixel level occupation histogram of Figure 3. A value of 1 corresponds to the (constant) relative occupancy of a completely uniform distribution. The "zero offset" evident in the histogram is the result of a hardware voltage offset adjustment in the sonar gain control.

The histogram equalisation algorithm works by reassigning the grey levels so that the new grey-level distribution is more uniform, that is, it has a more evenly distributed histogram than the original one. Full histogram equalisation produces the most evenly distributed final histogram; partial equalisation broadens peaks in the distribution and is often more useful for sonograms.

Assuming an image quantised to n levels from 0 to $n-1$, if the number of pixels in the image with level i is $N[i]$, then the normalised histogram for the image is

$$h[i] = \frac{N[i]}{\sum_{i=0}^{n-1} N[i]}, \quad i=0,1,\dots,n-1. \quad (3.1)$$

Partial histogram equalisation uses the transformation

$$v[0] = 0 \quad (3.2)$$

$$v[n-1] = n-1 \quad (3.3)$$

$$v[u] = \text{INT} \left(u + \alpha \left[n \sum_{i=0}^u h[i] - u \right] \right) \quad (3.4)$$

between old pixel values u and new pixel values $v[u]$ where the parameter α controls the degree of equalisation: if $\alpha = 0$ then no equalisation takes place, while $\alpha = 1$ gives full equalisation. Depending on the circumstances, more useful values tend to be of order 50%, that is, $\alpha = 0.5$. Note that both the original and new sets of pixel levels are quantised to the same resolution and that this technique does not attempt to interpolate values. For example, if the pixel level "m" in the old image maps to pixel value "n" in the new image, whereas "m+1" maps to "n+3" then the pixel values "n+1" and "n+2" will never be represented in the new image.

Alternatives to histogram equalisation include (a) truncation, where a certain subset of the full range of pixel values in the original image is reassigned to cover the whole range of pixel values in the final image, using a mapping like

$$\begin{aligned} v[u] &= 0 & u < u_{\text{lower}} \\ v[u] &= (n-1) \frac{u - u_{\text{lower}}}{u_{\text{upper}} - u_{\text{lower}}} & u_{\text{lower}} \leq u \leq u_{\text{upper}} \\ v[u] &= (n-1) & u > u_{\text{upper}} \end{aligned} \quad (3.5)$$

where the notations are obvious, and (b) non-linear mappings, of which the most used is the logarithmic mapping, given by

$$v[u] = (n-1) \frac{\log(1+u)}{\log(n)} \quad (3.6)$$

for example. Both of these methods are widely used in image processing, but offer less control than the histogram equalisation technique.

3.3.1 Processing to enhance contrast

In practice, the degree of histogram equalisation which should be used, from none to full, is very much dependent on the nature of the image. The image shown in Figure 3 does not benefit significantly from histogram equalisation since there is already useful contrast in the background and good differentiation between the background and the minelike object.

In contrast, the left third of Figure 12 shows a very faint image of a shipwreck recorded at 500 kHz with a Klein "Multiscan" multiple-beam sidescan sonar, in a form which is unsuitable for further processing. In this case, 50% partial histogram equalisation produces a more useful image for further processing, as is shown in the middle third. The right third of the image has been further processed and is discussed later in the report. Figure 13 shows the relative pixel level histograms before (top) and after (bottom) histogram equalisation, respectively. Note that after histogram equalisation,

the separation between the pixel levels with greatest occupation has increased. In effect, the grey-level mapping has been "stretched" to increase the contrast.

3.3.2 Processing to reduce noise

The top part of Figure 14 shows the opposite extreme to Figure 13, the relative level occupation histogram of Figure 2, which is dominated by slip-ring noise to the extent that approximately one third of the pixels in the image register at the maximum pixel level, (which is therefore excluded from the histogram, since it gives a level of 94.) Such a high level of maximal values in the image means that the streak removal algorithm is guaranteed to fail when applied to the original image, since it will try to compensate for the noise population and will hence over-compensate for any detail in the rest of the image. In this case, full histogram equalisation can be considered to be a pre-processing technique for filtration purposes. The relative level occupation histogram of the image after full histogram equalisation is shown in the bottom part of Figure 14. Note that the histogram is more like the one shown in Figure 11, in that most pixel values are concentrated at the lower end of the range. The large range of levels with no occupation is the result of the discrete nature of the algorithm, which has separated the unseen peak at level 255 from the rest of the populated levels in an attempt to make the distribution more uniform. The doubled values occurring in the lower levels are also a result of the discrete algorithm. Anticipating the following sections, Figure 15 shows the result of streak filtration of Figure 2. The left third of Figure 15 is the subset of Figure 2 which contains the mine, and the middle and right thirds are the same subset of Figure 2, after streak filtering and respectively without pre-processing (middle) and with full histogram equalisation and threshold filter pre-processing (right.) All parts of Figure 15 are low-quality images, but in the right image, the mine is more readily visible than in the other parts. (The mine is the dark, oval blob lying near the centre of each section horizontally and one quarter of the way from the bottom of the image, that is, about 25 m from the centre of the track.)




Figure 12. Three views of a section of the starboard side of an original sonogram recorded offshore of Sydney Heads. The faint, left image is the original form, the middle image has had histogram equalisation and the right image has had systematic correction and further histogram equalisation.




Figure 13. Relative pixel level occupation histograms corresponding to the left section of Figure 12 (top histogram) and middle section of Figure 12 (bottom histogram.)

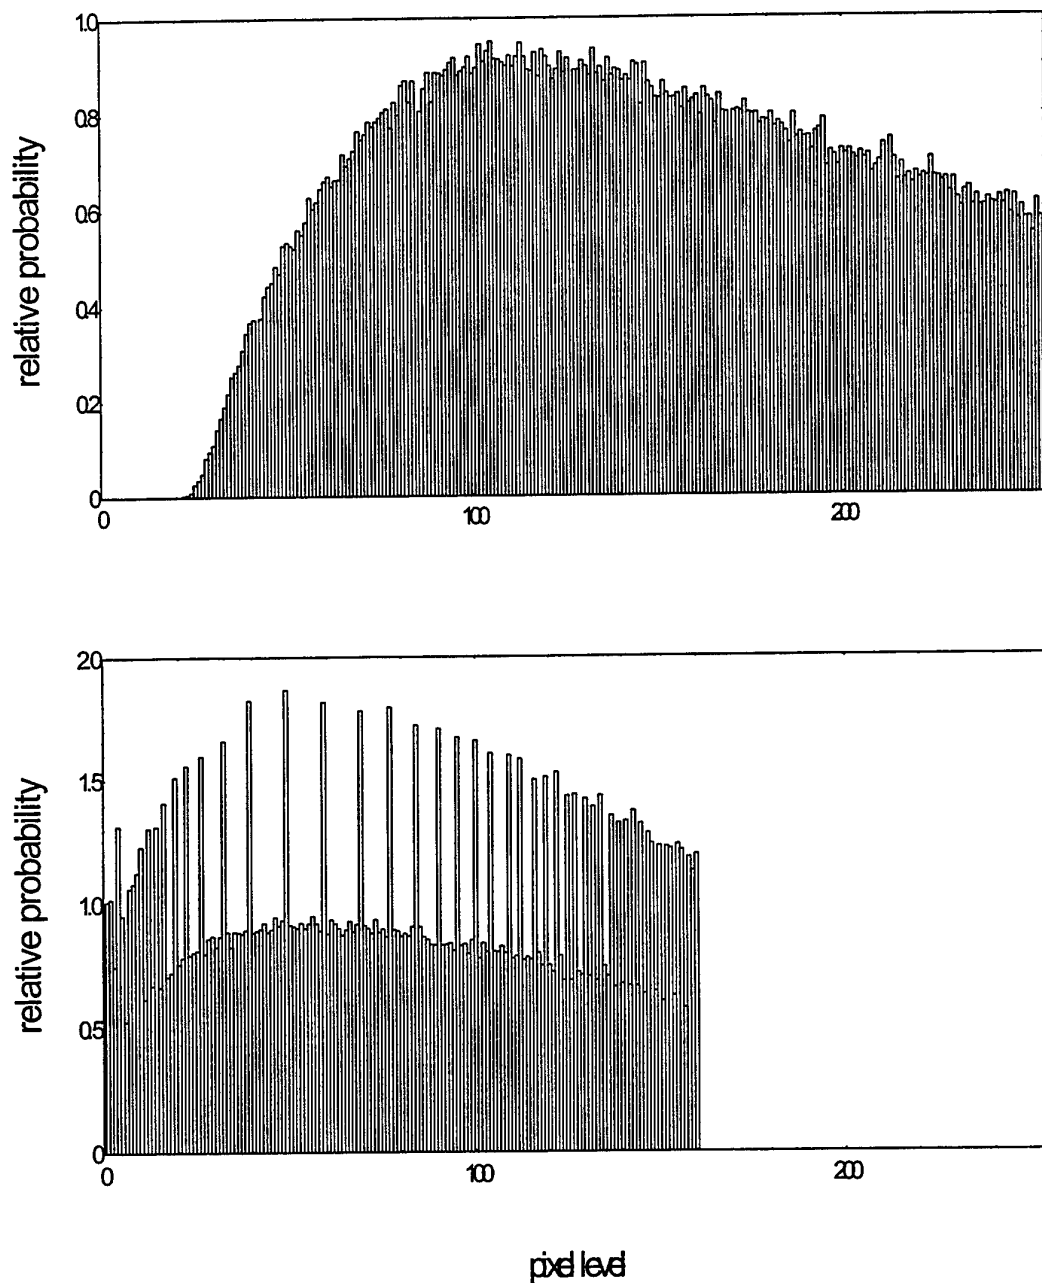


Figure 14. Relative pixel level occupation histograms corresponding to Figure 2. The top histogram corresponds to the unprocessed image, and the bottom histogram corresponds to the image after the application of full histogram equalisation.




Figure 15. Three images comparing the effect of streak removal on a section of Figure 2 containing a mine. The left third is original, the middle third has been streak filtered and the right third has been streak filtered after full histogram equalisation and thresholding.

3.4 Systematic correction

Although the systematic correction described in Section 2.6 is not strictly a pre-or post-processing technique, it requires a completely separate calculation to that used for streak removal, so it can be considered as such. It is designed to correct for systematic variation of image density with range, assuming that it exists. Many sidescan sonograms, such as Figure 1 for example, do not have a systematic variation with range in any case. The main problem with the technique in situations where a systematic variation is present, is that while it can restore the mean level of the image, for example correcting for intensity fall-off with range, it cannot restore the data which is lost due to lack of dynamic range and "clipping". As an example, consider the left third of Figure 16, which shows an area of sandy sea bottom in Jervis Bay. The data was recorded with the 500 kHz frequency of a Klein 590 sidescan sonar during trialing of a prototype sonar baffle and contains obvious striping in the along-track direction due to the interference of the baffle with the sound reaching the sonar transducers. The middle third of Figure 16 shows the same section of sonogram after application of the systematic correction from Paper I with $M'=512$ and $N'=80$. The right third of Figure 16 shows the same section of sonogram after application of the arbitrary correction technique described by equations (2.20) and (2.21), using a running average 40 pixels wide across track to smooth the along-track average, which was taken over the whole image. The value F_0 in (2.20) was set to the average level of the entire image. Comparing the two corrected sections, the middle section has the least variation across-track, but has the lowest contrast. The arbitrary technique used for the right section is less successful at reducing large intensity variations and in some cases has worsened them. Both corrected images would benefit from contrast enhancement by histogram equalisation, for example. Also, while it is true that the range variations in the corrected sections have been reduced, the total amount of information available to the observer has not been much increased.

In general, the systematic correction (unless assisted by contrast enhancement, for example by histogram equalisation) is of limited benefit for images in which the dynamic range of the sonar data used to create the image (here, 256 levels) is the same as the dynamic range of the display. It may be of greater benefit in the case of sonar systems storing sonar data at greater resolutions (say 12-bit with 4096 levels) where clipping is not prevalent and the contrast remaining after correction of the systematic gain variations remains useful. In the present case, most of the contrast has been lost to clipping of underflow and overflow.




Figure 16. Three images comparing of the effect of systematic corrections on a sonogram of a sandy area. The left third of the image is original. The middle third has had systematic correction as given by Cervenka and de Moustier and the right third by another algorithm.

3.5 An example using both histogram equalisation and systematic correction

Returning to the faint sonograph of the shipwreck, Figure 12, the right third of the image shows an example of the benefits of pre-processing with both histogram equalisation and the systematic correction. Starting from the original data, a 50% partial histogram equalisation was performed to bring pixel levels up to values at which they became useful input to other processes. Next, the systematic correction was applied with $M'=1024$ and $N'=50$, which corrected for the fall-off of image density with range. Finally, 50% partial histogram equalisation was applied again, to restore the contrast lost to the systematic correction. The result is markedly more useful than the original image, showing some detail of the wreck and also demonstrating the range at which the sonar is noise-limited.

4. Computational considerations

The spatial filtering algorithms mentioned above can be implemented as a continuous process relatively easily, given sufficient memory capacity. As each sonar record is read, the Chebychev coefficients (2.3) are calculated and stored. The convolution filter (2.9) can be applied to the stored coefficients (when sufficiently many are available) on a record by record basis, and reconstitution of the final image involves a single inverse transformation of type (2.7) using the difference of the filtered and unfiltered Chebychev coefficients, added to the original image, again on a record-by-record basis. The processing time associated with each stage depends on the number of initial Chebychev sampling points M and the number of Chebychev coefficients N retained in the expansion. In general, the intermediate convolution stage of the filter takes relatively little time to compute and the duration of the calculation is determined by the initial calculation of the Chebychev coefficients and the final evaluation of the filtered records. The times required for the initial and final steps are heavily dependent on which algorithms are used.

4.1 Initial sampling

All expressions to date have assumed a continuous functional form for the input and output data, $s_y(x)$ and $\bar{s}_y(x)$. In practice, each sonar record consists of a fixed number of samples representing equal spacings in either horizontal range or time (equivalently, slant range.) These samples are represented in a sonogram "waterfall image" as equally-spaced horizontal rows of pixels which typically correspond to some fraction of the original sonar records and some fraction of the original samples, dependent on the resolution of the display device. The values of the discrete, equally-spaced pixels in the image usually form the input to the filtration process, although the original sonar records could equally well be used. In either case, the continuous range variable x used in Section 2 must be replaced by some discrete (equivalently integral), variable representing pixel or sampling location.

In contrast to pixel locations, the optimal sampling locations (2.6) for the calculation of Chebychev coefficients are unequally spaced. The re-sampling which is therefore required can be achieved by either an interpolation or by decimation. In view of the fact that sonar records are typically (a) noisy and (b) consist of a large number of samples, it is not clear that interpolation during the initial calculation of Chebychev coefficients is particularly beneficial - decimation is probably sufficient - however, the final calculation of the "processed part" (2.18) of the record may also require mapping of function values at unequally-spaced locations to pixels with equally-spaced locations, in which case interpolation may be of value.

The decimation scheme used in this work assumes that the set of original samples is given by

$$S_y[k], \quad k = 1, 2, \dots, M_0 \quad (4.1)$$

where M_0 is the original number of samples and the decimation to a set of M "approximately optimal" (in the sense of a Chebychev expansion) samples is achieved by the mapping

$$s_y(x_i) \cong S_y[k(i)], \quad i = 1, 2, \dots, M \quad (4.2)$$

where

$$k(i) = \text{NINT} \left(1 + \frac{x_i - x_{\min}}{x_{\max} + x_{\min}} (M - 1) \right) \quad (4.3)$$

using the optimal sampling points x_i defined in equation (2.6). "NINT" means "the nearest integer number". The mapping $k(i)$ need only be calculated once and the decimation process is extremely rapid. Two or three-point interpolation schemes can also be readily derived from the same mapping.

4.2 Calculation of Chebychev coefficients

The Chebychev coefficients $a_j[y]$ may be calculated in at least two distinct ways, the utility of which depend on the values of M and N .

The first way, described in Paper I, invokes equation (2.3) directly. Such a calculation can be partially optimised by pre-calculation of the values of x_i , θ_i and so on, however the calculation is dominated by the requirement for MN multiplications to produce the full coefficient set for each record. Pre-calculation of the entire set of values $\cos(j\theta_i)$, $i = 1, 2, \dots, M$, $j = 0, 1, \dots, N - 1$ does not reduce the duration of the calculation to a useful extent and the process becomes extremely time-consuming for moderate values of M and N . For example, for $M=512$ and $N=200$, the calculation for a screen-full of data took approximately half an hour on an Intel 486, 33-MHz PC.

A second option for the evaluation of equation (2.3) is noted by Press, *et al* (1992) in their book "Numerical Recipes in C", p 193, where they note that expressions (2.3-2.6) for the Chebychev coefficients are identical to the discrete cosine transform (DCT) of the second kind (p 519). Thus, so long as M is set to a power of two, a fast (order $M \log_2 M$) algorithm can be used to generate the first M Chebychev coefficients. Implementation of the coefficient calculation using the routine "cosft2" from Numerical Recipes, which is based on the fast Fourier transform (FFT), resulted in a process which is faster for useful values of N than optimisation of equation (2.3). In addition, experimentation with different values of N was particularly easy.

4.3 Filtering of coefficients

Across-track filtration of each record is achieved simply by truncation of the Chebychev series at term $N-1$. The next stage of filtration, application of a convolution mask to produce the "base image" filtered across- and along-track, acts on the block of coefficients $a_j[k]$, $j=0,1\dots N-1$ and $k=j-N, j-N+1 \dots j+N$, or rather on the subset of that block selected by the variable-width convolution mask. Calculation time increases as N^2 , however, in practice, the convolution is a relatively minor part of the total calculation time. Implementation of the convolution for continuous filtration is a straight-forward process which can be achieved using a standard "circular buffer".

4.4 Evaluation of the filtered image

After the calculation of the coefficients $a_j[i]$ and $\bar{a}_j[i]$, the remaining process is the reconstitution of the final filtered image as per equation (2.17). As was the case with the evaluation of Chebychev coefficients, at least two methods of calculation are possible.

The equal sample spacing of the original image can be incorporated into a direct evaluation of equation (2.17) with appropriate modifications to the function evaluation points. Assuming that the original equally-spaced sample points are given by

$$\Xi_i, \quad i = 1, 2, \dots, M_0 \quad (4.4)$$

then the equivalent X values are

$$X_i = \frac{2\Xi_i - (x_{\max} + x_{\min})}{x_{\max} - x_{\min}} \quad (4.5)$$

and finally,

$$T_j(X_i) = \cos(j\Theta_i) \quad (4.6)$$

where

$$\Theta = \cos^{-1}(X_i). \quad (4.7)$$

Pre-calculation of the arc cosine values is a partial optimisation in this case. The truncated Chebychev polynomial expansion (2.7) is then best evaluated by the Clenshaw recurrence formula given by Press *et al* (1992) in Numerical Recipes, p193.

The second method uses the inverse DCT to calculate (2.18). This has the advantage of returning all the function values after one calculation, however the spacing of the output values from the DCT is uneven and must be corrected by an inverse decimation or interpolation mapping. To apply the inverse DCT, an input vector of length M_1 is prepared, where M_1 must be a power of two greater than N and preferably greater than or equal to M_0 , the original number of pixels in the record. The first N values of the vector are set to the Chebychev differences $\bar{a}_j[y] - a_j[y]$, $j = 0, 1, \dots, N-1$ and the rest are set to zero. After application of the inverse DCT, the M_1 function values must be mapped to equal spacings. This is achieved with the decimation mapping

$$\kappa(i) = \text{NINT} \left(0.5 + \frac{M_1}{\pi} \cos^{-1} \left\{ \frac{2i - (M_0 + 1)}{M_0} \right\} \right) \quad (4.8)$$

where $i = 1, 2, \dots, M_0$ is the index of the position in the original (displayed) record and $\kappa(i)$ is the index of the closest corresponding position in the output of the DCT. Equation (4.8) is derived by inverting equations (2.4-2.6) as follows: we require that the output of the DCT be mapped to an equally-spaced grid whose samples fall on the locations

$$x_i = x_{\min} + \frac{x_{\max} - x_{\min}}{M_0} (i - 1/2), \quad i = 1, 2, \dots, M_0, \quad (4.9)$$

that is, at the centres of M_0 equal sub-intervals. Inverting equations (2.5) and (2.6), we see that these sampling locations correspond to angles θ_i which are given by

$$\theta_i = \cos^{-1} \left(\frac{2i - (M_0 + 1)}{M_0} \right) \quad (4.10)$$

after some cancellation. As there are M_1 coefficients in the input to the inverse DCT, we invert equation (2.4) with $M = M_1$ and substitute from (4.10) to give equation (4.8). This can also be used as the basis for interpolation.

4.5 Control of filter characteristics

The parameters M and N defining the truncation of the exact, infinite Chebychev expansion (2.1) to the approximate "low-pass filtered" version (2.18) control not only the appearance of the filtered image, but also the calculation time and the storage required for the calculation. Without optimisation, the transformations (2.3) and (2.18) require order MN separate calculations for each record, which makes the process

prohibitively slow if more than a few tens of coefficients are required. In general, the DCT implementation of the process has proven to be acceptably rapid, involving order $M \log_2 M$ calculations, and flexible. The restriction of M values to powers of 2 required by the DCT can be removed by zero-padding or decimation. Comparisons of the results of filtration using interpolated, directly-calculated and decimated, DCT-based implementations of the filter revealed little obvious loss of quality by decimation: although the directly-calculated image was of marginally superior quality, the differences caused by other factors such as selection of the values M and N were more important.

4.6 Control of filter "over-correction"

Recalling Section 3.1, inspection of Figure 1, the shipwreck image, and Figure 3, the mine, reveals that the focal points of each image contain areas of high intensity and shadow which will result in an over-reaction from the low-pass filtering process used in the streak-removal filter. If the filter as described is used on these images without any allowance for the over-correction, then when low values of N , the truncation parameter, are used, regions to the left and right of the high-intensity patches will be rendered with lower intensity than they should be, while the shadow regions will be corrupted by streaks.

There are two ways to correct the problem. The first way is to use a large value for the filter truncation parameter N , say up to one half of M . This effectively means that low-frequency "over-correction" is removed by the retention of the higher frequencies, while the streaking is removed by the wide smoothing convolution mask acting on the Chebychev coefficients. However, this results in loss of contrast in the image, which may mean that crucial features are missed.

The second method of removing over-correction is to recognise that the regions in which the over-correction is most obvious are shadow regions and regions of very high intensity which do not normally contain obvious streaking before processing, due to clipping and to the limitations of the display. Thus, a "mask" may be applied to the filter, such that pixels in the original image with levels above a "high" threshold or below a "low" threshold are passed over and left unchanged by the filter. This technique has proved effective when used with a high threshold level of 90% of the maximum intensity and a low threshold level of 25% of the maximum intensity.

A potentially better approach, which may be investigated in future, would involve the temporary replacement of regions of maximum and minimum intensity in the original image with some mean value, which would not disrupt the streak-removal processing on the rest of the image. The subtracted regions would be restored in the processed image after streaks had been removed from the rest of the image. The ramifications of such a technique are however not clear at present.

5. Filtration characteristics

5.1 Performance

Tests of the behaviour of the streak removal filter were made using various images, including those shown by Figures 1, 2 and 3.

Initially, most of the effort was concentrated on Figure 1, the shipwreck. "Filter over-correction" problems mentioned in Section 4.6, at that time uncorrected, meant that the best output from the filter occurred when filter parameters $M=512$ and $N=200$ were used. The value of M corresponded to the number of pixels in each record. Values of N much below 200 resulted in shadow regions corrupted by streaking and values higher than 200 were unacceptably blurry due to the action of the convolution filter on the finest details. The value $N=200$ was a surprising result, since it implied that the smoothing given by the convolution operator (2.9-2.13) was required at spatial frequencies corresponding to almost pixel-level detail. It was also discouraging from the point of view of implementation of the filter as a component of a continuous sonar data-processing package. If the filter were invoked by a button-press, the operator would have to wait while the Chebychev coefficients for 200 records before and after the "current" record were calculated and processed. Finally, the filter caused a considerable loss of contrast.

When an equivalent filter (with $M=1024$ and $N=400$) was later applied to the mine image, Figure 3, the averaging inherent in the convolution operation caused the sharp edges of the mine to become blurred and caused the contrast to drop. It was thus clear that high values of N could be actively harmful in a mine hunting context, even where they gave pleasing results in other circumstances. Finally, the "masking" technique described in Section 4.6 was implemented and this allowed filtration with values of N as small as 10 or 20 to give acceptable results. The streak filter with such parameters has the required property that, when it is applied to an image of good quality such as Figure 3, it has relatively little effect, that is, the image of the mine retains its contrast against the background and its sharp edges.

Experiments made with various values of M , the sampling number, showed that the best filtered image quality was always obtained when M was set greater than or equal to the number of pixels M_0 in the record. The computation time for the filtering process was relatively insensitive to lower values of M , particularly since reconstruction of the image from the truncated Chebychev coefficients was most effectively done by an inverse DCT with output array length at least M_0 . Figures 17 and 18 show comparisons of different filter parameters acting on the same section of sonogram. Figure 17 shows the shipwreck which is part of Figure 1 and Figure 18 shows the mine which is part of Figure 3. The filter parameters and approximate execution times for the filtrations are shown in Tables 1 and 2 following.




Figure 17. A mosaic of identical sections of the Figure 1 shipwreck image after streak filtration with different filter parameters. The lower left image is unprocessed. Filter parameters and approximate filtration times for each of the cases are given in Table 1 in the text.




Figure 18. A mosaic of identical sections of the Figure 3 mine image after streak filtration with similar filter parameters to those used in Figure 17. The lower left image is unprocessed. Filter parameters and approximate filtration times for each of the cases are given in Table 2 in the text.

Table 1. Filter parameters and execution times for Figure 17

	<i>Left</i>	<i>Middle</i>	<i>Right</i>
Top	M=512, N=100, t=9 s	M=512, N=200, t=11 s	M=128, N=20, t=5 s
Bottom	Unprocessed	M=512, N=5, t=6 s	M=512, N=20, t=7 s

Table 2. Filter parameters and execution times for Figure 18

	<i>Left</i>	<i>Middle</i>	<i>Right</i>
Top	M=1024, N=100, t=7 s	M=1024, N=200, t=9 s	M=256, N=20, t=4 s
Bottom	Unprocessed	M=1024, N=5, t=6 s	M=1024, N=20, t=7 s

Comparing the images in Figure 17, the image with $M=512$ and $N=200$ has the most pleasing background, however this is probably due to the smoothing of features through the filtration process. The image with $M=128$ and $N=20$ clearly shows less contrast than the equivalent image at the normal filter settings, with $M=512$ and $N=20$, showing that it is desirable to use a value for M which is as high as practical. Figure 18 compares the effect of the filter on the mine, using filter settings which are somewhat equivalent to those used in Figure 17, although the sonar record is twice as long in Figure 3 as it is in the two-sided sonogram Figure 1, so the filter parameters are not exactly equivalent. Depth-sounder noise was not pre-processed out of the sonogram before the streak filter was applied. Close inspection of the image with $M=1024$ and $N=200$ shows that the resolution of the mine has deteriorated and that some of the features in the background are exaggerations of what is seen in the raw image. Hence, it seems that high values of N are not generally useful.

As noted previously, the calculation time for the process with a Pentium 133MHz based PC is of order 6 to 12 seconds per screen-full. In comparison, the systematic correction and histogram equalisation calculations both take a few seconds on the same machine and the thresholding filter takes of order 1 second.

5.2 Towards intelligent filtration

The techniques demonstrated in this report are generally not universally applicable, with the possible exception of the thresholding filter, because the nature of sidescan sonar images is extremely variable. It would be preferable that the data processor itself monitored the data before it came on-screen in order to determine whether filtration should be applied, and then applied it automatically. There are indications in the data given in this report as to how such a process could be implemented, with a small delay between the reading of sonar records and their display. Continuous monitoring of the pixel level histogram of the records not yet displayed would reveal if histogram equalisation were required, and the degree to which it would be needed

could be calculated from the statistical moments of the histogram. Monitoring of the lowest few Chebychev coefficients of the incoming records would likewise reveal if streaking were occurring, so that it could be corrected before it came on-screen.

6. Conclusions

This report has shown that techniques discussed by Cervenka and de Moustier in their 1993 paper on the processing of large, low-frequency sidescan sonar images translate well to images from the smaller, higher-frequency devices used by the RAN Route-Survey System and that the techniques can be implemented in a fashion in which real-time image processing is possible on PC platforms available now.

7. References

- Cervenka, P. and De Moustier, C. (1993). Sidescan sonar image processing techniques. *IEEE Journal of Oceanic Engineering*, **18**, 2, 108-122.
- Lee, J. -S. (1983). A simple speckle smoothing algorithm for synthetic aperture radar images. *IEEE Transactions on Systems, Man and Cybernetics*, **13**, 85-89.
- Overkov, N. C. (1995) *Honours thesis*. Dept. of Earth Sciences, University of Sydney, Australia.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T. and Flannery, B. P. (1992). *Numerical Recipes in C (2nd. Ed.)*. Cambridge University Press

8. Acknowledgments

The author would like to acknowledge the efforts of Dr Roger Neill in proof-reading this report and providing useful advice and appropriate sonar data. He would also like to acknowledge the efforts of Dr John Best in twice reviewing the report and identifying weaknesses in the original approach.

DISTRIBUTION LIST

Spatial Filtering of Sidescan Sonograms to Remove Streaking due to Signal Level
Fluctuations

Stuart Anstee

AUSTRALIA

DEFENCE ORGANISATION

Task Sponsor Director General Force Development (Sea)

S&T Program

Chief Defence Scientist	} shared copy
FAS Science Policy	
AS Science Industry and External Relations	
AS Science Corporate Management	
Counsellor Defence Science, London (Doc Data Sheet)	
Counsellor Defence Science, Washington (Doc Data Sheet)	
Scientific Adviser to MRDC Thailand (Doc Data Sheet)	
Director General Scientific Advisers and Trials/Scientific Adviser Policy and Command (shared copy)	
Navy Scientific Adviser	
Scientific Adviser - Army (Doc Data Sheet and distribution list only)	
Air Force Scientific Adviser	
Director Trials	

Aeronautical and Maritime Research Laboratory

Director

Electronics and Surveillance Research Laboratory

Director

Chief of Maritime Operations Division
Research Leader Operational Environment
Group Leader Minehunting Sonar
Roger Neill, MOD
Stuart Anstee
Kurt Benke, MOD
John Best, MOD
Ross Dawe, MOD
Tony Parkinson, MOD
Ranjit Thuraisingham, MOD
David Blair, MOD

DSTO Library

Library Fishermens Bend
Library Maribyrnong
Library DSTOS (2 copies)
Australian Archives
Library, MOD, Pyrmont (2 copies)

Forces Executive

Director General Force Development (Land) (Doc Data Sheet only)

Director General Force Development (Air) (Doc Data Sheet only)

Navy

SO (Science), Director of Naval Warfare, Maritime Headquarters Annex,
Garden Island, NSW 2000. (Doc Data Sheet only)

Army

ABCA Office, G-1-34, Russell Offices, Canberra (4 copies)

S&I Program

Defence Intelligence Organisation

Library, Defence Signals Directorate (Doc Data Sheet only)

B&M Program (libraries)

OIC TRS, Defence Central Library

Officer in Charge, Document Exchange Centre (DEC), 1 copy

*US Defence Technical Information Centre, 2 copies

*UK Defence Research Information Center, 2 copies

*Canada Defence Scientific Information Service, 1 copy

*NZ Defence Information Centre, 1 copy

National Library of Australia, 1 copy

UNIVERSITIES AND COLLEGES

Australian Defence Force Academy

Library

Head of Aerospace and Mechanical Engineering

Deakin University, Serials Section (M list), Deakin University Library, Geelong, 3217

Senior Librarian, Hargrave Library, Monash University

Librarian, Flinders University

OTHER ORGANISATIONS

NASA (Canberra)

AGPS

OUTSIDE AUSTRALIA

ABSTRACTING AND INFORMATION ORGANISATIONS

INSPEC: Acquisitions Section Institution of Electrical Engineers

Library, Chemical Abstracts Reference Service

Engineering Societies Library, US

American Society for Metals

Documents Librarian, The Center for Research Libraries, US

INFORMATION EXCHANGE AGREEMENT PARTNERS

Acquisitions Unit, Science Reference and Information Service, UK

Library - Exchange Desk, National Institute of Standards and Technology, US

SPARES (10 copies)

Total number of copies: 65

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA				1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)		
2. TITLE Spatial filtering of sidescan sonograms to remove streaking due to signal level fluctuations			3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)			
4. AUTHOR(S) Stuart Anstee			5. CORPORATE AUTHOR Aeronautical and Maritime Research Laboratory PO Box 4331 Melbourne Vic 3001			
6a. DSTO NUMBER DSTO-TR-0443		6b. AR NUMBER AR-009-931		6c. TYPE OF REPORT Technical Report		7. DOCUMENT DATE November 1996
8. FILE NUMBER 510/207/0638	9. TASK NUMBER ADS 94/193	10. TASK SPONSOR DGFD (Sea)	11. NO. OF PAGES 40		12. NO. OF REFERENCES 4	
13. DOWNGRADING/DELIMITING INSTRUCTIONS None			14. RELEASE AUTHORITY Chief, Maritime Operations Division			
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <i>Approved for public release</i> OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE CENTRE, DIS NETWORK OFFICE, DEPT OF DEFENCE, CAMPBELL PARK OFFICES, CANBERRA ACT 2600						
16. DELIBERATE ANNOUNCEMENT No limitations						
17. CASUAL ANNOUNCEMENT Yes						
18. DEFIEST DESCRIPTORS Acoustic Imaging; Seaborne Mine Hunting; Route Surveys; Towed Sonar; Active Sonar; Scanning Sonar						
19. ABSTRACT A spatial filtering technique originally devised to filter streaking from geographically-scaled sidescan sonar images is applied successfully to smaller-scale sidescan sonar images generated during mine warfare exercises. Implementation of the algorithm to minimise computation time is discussed and appropriate filter parameter settings are established. Related techniques for image contrast and noise control are also discussed.						