Australian Government
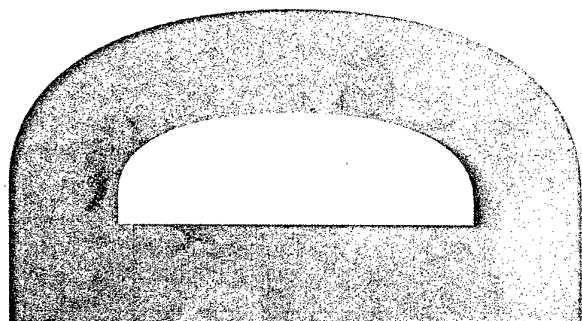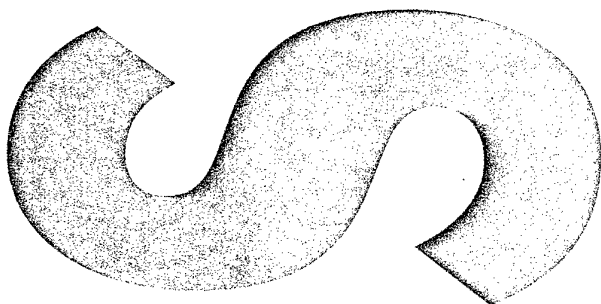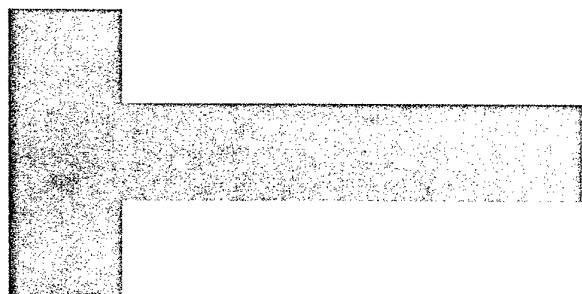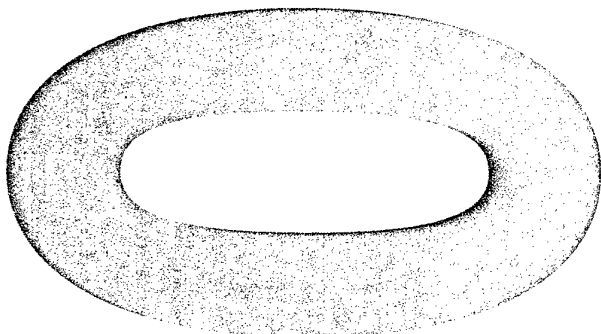
Department of Defence

Defence Science and
Technology Organisation

Implementation of the Radon Transform Using Non-equispaced Discrete Fourier Transforms

Ronald Jones, Tristrom Cooke and Nicholas J. Redding

DSTO-TR-1576

20040915 050

Implementation of the Radon Transform Using Non-equispaced Discrete Fourier Transforms

Ronald Jones, Tristrom Cooke and Nicholas J. Redding

Intelligence, Surveillance and Reconnaissance Division
Information Sciences Laboratory

DSTO–TR–1576

ABSTRACT

This report discusses the implementation of the Radon transform in the Analysts' Detection Support System (ADSS) environment using non-equispaced Discrete Fourier Transforms (DFTs). It provides an analysis and experimental results for discretisation error and the use of matched filtering to enhance peaks in the transform.

APPROVED FOR PUBLIC RELEASE

AQ F04-11-1309

APPROVED FOR PUBLIC RELEASE

Implementation of the Radon Transform Using Non-equispaced Discrete Fourier Transforms

EXECUTIVE SUMMARY

The Radon transform may be used to identify linear features in an image and has proven to be a useful tool for extracting roads and faint trails in Synethetic Aperture Radar (SAR) imagery. In particular, it is robust to the large amount of background clutter and speckle noise associated with such images. This document reports on an implementation of the Radon transform within the Analyst's Detection Support System (ADSS). In particular, a novel implementation method has been used that is based on non-equispaced discrete fourier transforms, which are leading edge. The report offers an analysis and experimental results for two paths of study: the discretisation errors that arise from the approach and a matched filtering method that may be used to enhance detection results.

Authors

Ronald Jones
SYDAC Pty Ltd

Ronald Jones received a B.Sc. and Ph.D. in Physics from Victoria University of Wellington and Monash University, Melbourne, in 1990 and 1994, respectively. Since completing his PhD, he has worked as a Research Scientist at the CSIRO, Division of Mathematical and Information Sciences, and Proteome Systems Ltd, a BioTech company based in Sydney. His research interests include linear and nonlinear image processing techniques.

Tristrom Cooke
CRC for Sensor Signals and Information Processing

After completing a B.Eng in Electrical Engineering from the University of South Australia in 1992 and a B.Sc in Applied Mathematics from the University of Adelaide in 1996, Tristrom Cooke completed a Ph.D. relating to the thermal and mechanical deformation of mechanical structures in 1999. Since this time, he has been employed by CSSIP on projects related to automatic target detection and classification in various types of radar imagery, comparison of bullet scratches for use in forensics, and electronic document classification. His research interests include machine learning and image processing.

Nicholas J. Redding
Intelligence, Surveillance and Reconnaissance Division

Nicholas Redding received a B.E. and Ph.D. in electrical engineering all from the University of Queensland, Brisbane, in 1986 and 1991, respectively. From 1988 he received a Research Scientist Fellowship from the Australian Defence Science and Technology Organisation (DSTO) and then joined DSTO in Adelaide as a Research Scientist after completing his Ph.D. in artificial neural networks in 1991. In 1996 he was appointed as a Senior Research Scientist in the Microwave Radar Division (now Intelligence, Surveillance and Reconnaissance Division) of DSTO. Since joining DSTO he has applied image processing techniques to the automatic classification of ionospheric data, and more recently has researched target detection (both human and algorithmic) in synthetic aperture radar (SAR) imagery. He has recently returned from a one and a half year posting to the UK's Defence Evaluation and Research Agency where he continued the development of a suite of target detection algorithms for SAR imagery and researched new algorithms in SAR image forming using the circular Radon transform.

Contents

1 Introduction

The Radon transform may be used to extract the parameters of linear features in an image [1]. Lines in the input image are realised as peaks in the Radon transform image at positions corresponding to the parameters of the line. In this way, the Radon transform can render a difficult global linear detection problem into a more easily solved peak detection problem. This report details our implementation of a **Radon** module in the Analysts' Detection Support System (ADSS) framework [7] using non-equispaced DFTs. We describe efforts to control discretisation effects and apply matched filtering to enhance peaks.

An example of the Radon transform is shown in Fig. 1 using the ADSS module **radon**. At the top, a sample of a SAR image is shown with two faint horizontal trails, one light and one dark. With no prior assumptions of width, length and direction, detecting these trails reliably in the image domain is a difficult task. The bottom left of the figure shows a portion of the result for the Radon transform. A bright peak corresponding to the light trail is clearly visible at the top of the image, along with a similar dark object representing the dark trail in the image. A number of peak detection algorithms might be used to extract these objects in a straight forward manner. For example, the absolute value of the radial derivative of the Radon transform (RDRT; see Section 2.3) has proven effective at finding the edges of roads and trails in noisy SAR images. It is also handy because a single threshold can be used to pull out both light and dark trails. The bottom right of the figure shows a portion of the result for the RDRT which may now be thresholded to extract the line parameter information.

This report will proceed as follows. In the following section, we provide some background theory on the Radon transform before detailing our Radon transform implementation using non-equispaced DFTs. We then look at discretisation errors in Section 3, an issue that has been the focus of a lot of attention and work. In Section 4 we look at the analytical form of a matched filter that may be used to accentuate peaks in noisy Radon transforms. We conclude with some remarks regarding future work in Section 5.

2 Background Theory and Implementation

The Radon transform of the function $f(x, y), x, y \in R$ is typically defined as a path integral along a straight line of the function,

$$g(\rho, \theta) = (\mathcal{R}f)(\rho, \theta) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y)\delta(\rho - x\cos\theta - y\sin\theta)\, dx\, dy,$$

where $\delta(.)$ is the Dirac delta function and (ρ, θ) are the parameters of a normal equation for the line of integration. The Projection Slice Theorem [1] states that the 1-D transform of any projection $p_\theta(\rho) = (\mathcal{R}f)(\rho, \theta)$ is equal to the 2-D FFT of the image $f(x, y)$ with respect to polar coordinates, *i.e.*

$$(\mathcal{R}f)(\rho, \theta) = \mathcal{F}_\omega^{-1}\left[(\mathcal{F}_\mathbf{x}f)(\omega, \theta)\right]. \tag{1}$$
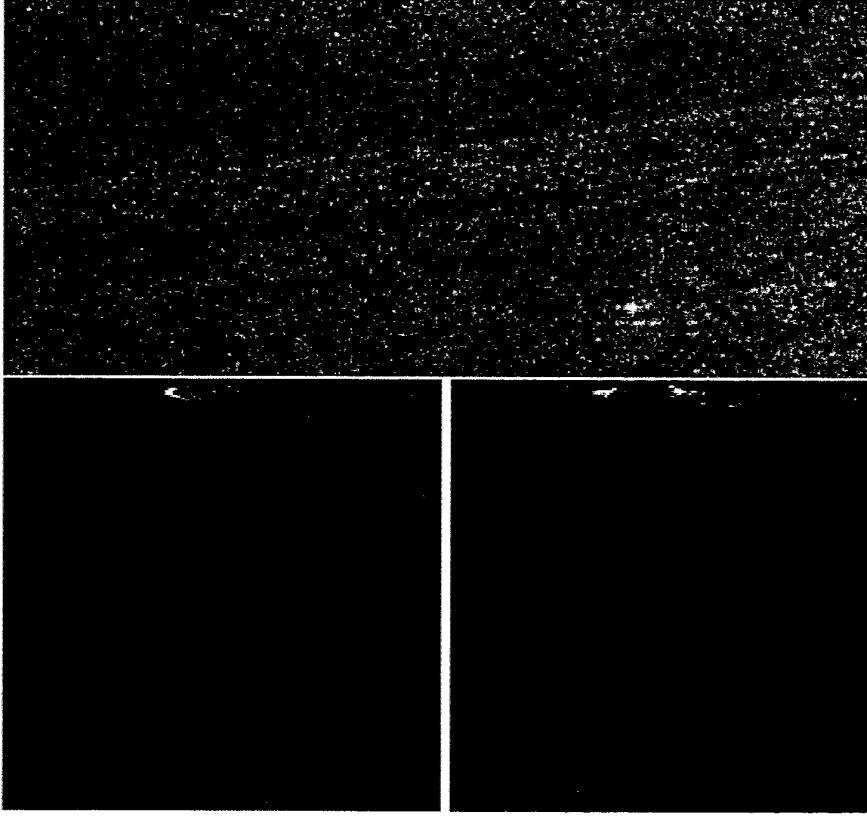
1

Figure 1: Example Radon transform of SAR image. Top: Sample of input image showing faint light and dark horizontal trails. Bottom Left: Result from standard Radon transform. Bottom Right: Result from RDRT.

Here $\mathcal{R}f$ denotes the Radon transform of f, $\mathcal{F}_\omega^{-1}$ denote the 1-D inverse FFT with respect to the variable ω (the frequency counterpart to ρ) and $\mathcal{F}_\mathbf{x}f$ denotes the standard 2-D FFT of the image f with respect to the (2-D) variable $\mathbf{x}$.

2.1   Implementation using Non-equispaced FFTs

The relationship in Eq. (1) allows a comparatively fast and direct implementation of the Radon transform by utilising non-equispaced DFT algorithms [6]. All variants of non-equispaced Fourier algorithms utilise normal equispaced DFTs. As detailed in [6], the input image is first scaled and zero padded before the 2-D DFT is applied. Illustrated in Fig. 2 is the result $\mathcal{F}_\mathbf{x}f$, represented by a square with extent Ω. The non-equispaced algorithm uses this result to calculate the DFT at the non-equispaced polar coordinates $(\omega_p \cos\theta_q, \omega_p \sin\theta_q)$, where

$$\theta_q = q\pi/Q, \qquad q = 0, \ldots, Q-1$$

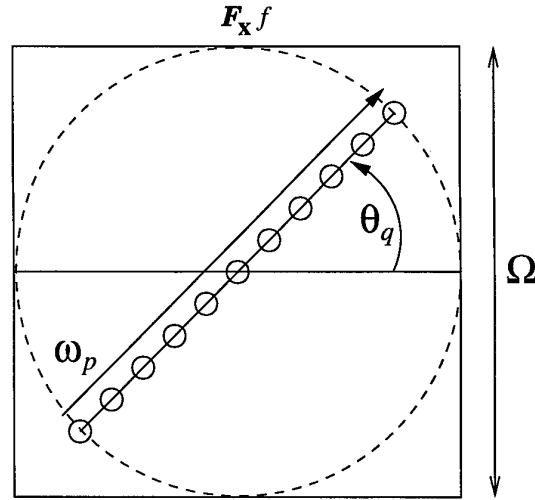$$\omega_p = p\,\Omega/N, \qquad p = 0, \ldots, N-1 \tag{2}$$

Figure 2: Implementation of Radon transform using a non-equispaced Fourier transform.

From here, it is a simple matter to compute the Radon transform result $(\mathcal{R}f)(\rho, \theta_q)$ for each angle θ_q by applying a 1-D inverse DFT to the set of points $(\mathcal{F}_\mathbf{x}f)$ $(\omega_p \cos\theta_q, \omega_p \sin\theta_q)$, where $p = 0, \ldots, N-1$. In our implementation, this yields a rectangular image where θ is in the y direction and ρ is in the x direction.

Note that as shown the set of polar coordinates $(\omega_p \cos\theta_q, \omega_p \sin\theta_q)$ does not cover the entire extent of the transform image $\mathcal{F}_\mathbf{x}f$; high frequency components in the corners of the image are not sampled. Moreover, points near the centre of the image will be more densely sampled than points out towards the edges. This sampling strategy was necessary in order to generate a Radon transform on a grid. However, there are certainly other possibilities, some of which are discussed in Section 3.

2.2   Interpreting the Radon Transform Image

Figure 3 shows an example result illustrating the coordinate system of the Radon transform. A single point, parameterised by (ρ, θ) and of greylevel one, shown on the left of Fig. 3, yields the expected sinusoidal curve in the Radon domain, shown on the right. The angle θ, in the vertical direction of the image, has a range $[0, 180)$, where 0 is at the top of the image and the number of rows in the vertical direction is specified by the Radon parameter theta-resolution. The length ρ, in the horizontal direction, has a range $(-N/2, N/2)$, where N is specified by the Radon parameter rho-resolution and $\rho = 0$ at the pixel location rho-resolution / 2. Negative values for ρ occur for points that lie in the lower half of the input image, as this is where θ values in the range $[180, 360)$ wrap into the range $[0, 180)$. The amplitude of the sinusoidal curve is given by ρ, the distance from the origin to the point.

Note that the parameter rho-resolution is actually the number of points N that are sampled in the frequency domain along the line at each angle θ_q, as per Eq. (2). Increasing N, for example, will reduce the distance Δw between points in the frequency
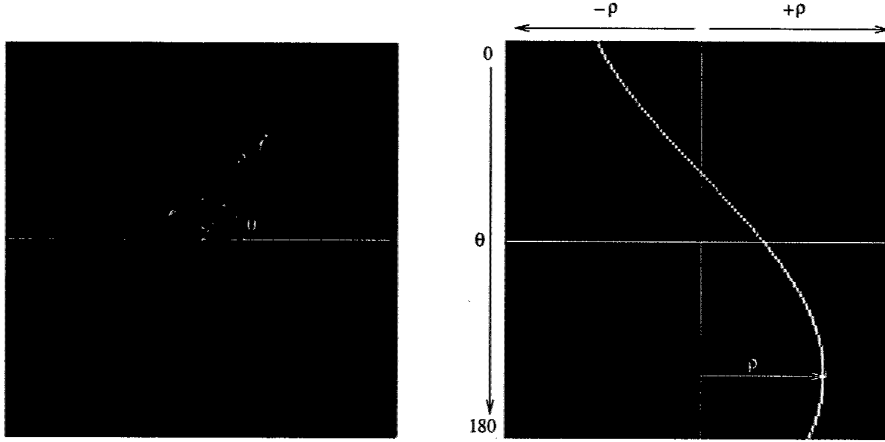
Figure 3: Example illustrating Radon domain coordinate system. Left: Input image with single point parameterised by (ρ, θ). Right: Resulting Radon transform with ρ in the horizontal direction and θ in the vertical direction.

domain while the distance Δx in the spatial domain remains constant. Thus the range of ρ increases but the spatial resolution remains the same. Values of ρ that range outside the `rho-resolution` specified will wrap to the other side of the Radon transform. It is important to avoid this wrap around in peak detection, so the parameter `rho-resolution` should be chosen so as to cover the total range of possible values for ρ in the image. This can be achieved by setting `rho-resolution` to the length of the diagonal of the input image.

Figure 4 shows an example using a line parameterised by the same values of ρ and θ used for the point in the example above. The resulting Radon transform, shown in the middle of the figure, is a superposition of sinusoids that reinforce at coordinates corresponding to the parameters (ρ, θ) for the line. A coordinate (ρ', θ') representing a peak in the Radon transform image will correspond to a line in the image domain with parameters

$$\rho = \rho' - \frac{\texttt{rho-resolution}}{2}, \qquad \theta = \frac{180.\theta'}{\texttt{theta-resolution}}.$$

The characteristic peak, which we will refer to as a *butterfly shape*, is shown magnified on the right of the figure. Curvilinear detection using the Radon transform seeks to locate such shapes in order to yield parameters for curvilinear features. In Section 4 we look at efforts to enhance this feature using a matched filter.

2.3   The Radial Derivative of the Radon Transform

In our current implementation, the radial derivative of the Radon transform (RDRT) is carried out in accordance with its original implementation in [2]: A simple pointwise difference in the horizontal direction (corresponding to ρ) of the Radon transform. This generates positive and negative edges at locations corresponding to line edges in the input image. We use the absolute value of the RDRT in order to turn (negative) valleys into (positive) peaks and thus be able to apply a single threshold to extract peak locations. The
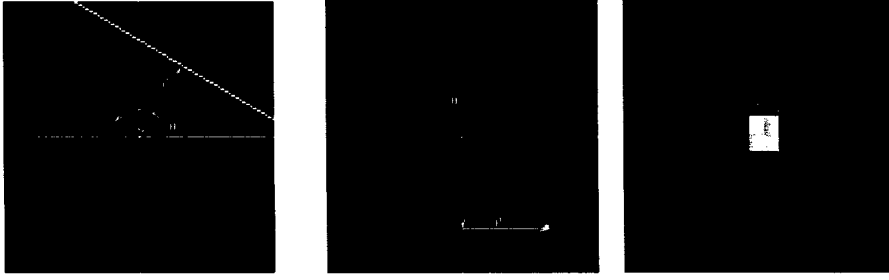
Figure 4: Radon Transform for a line showing characteristic peak or "butterfly" shape.

use of the absolute value has consequences however when it comes to applying a matched filter, as discussed in Section 4.

As pointed out in [5], using pointwise subtraction is known to be numerically unstable and it is desirable to implement the derivative in the Fourier domain if possible, where the derivative is realised as a simple multiplication, *i. e.*

$$\mathcal{F}_\rho\left(\frac{\partial(\mathcal{R}f)(\rho,\theta)}{\partial\rho}\right) = 2\pi i\omega F(\omega,\theta).$$

Consequently, the derivative could be achieved by applying the multiplication factor to the 2D Fourier transformed data before the final Fourier Transform of the Projection Slice Theorem. However, for our purposes we found that this led to a number of problems that emerged when we used the Radon transform as the basis for a curvilinear feature extractor [2]. In particular, for practical implementations, it is important to normalise the Radon transform to remove bias towards the longer lines that pass through the centre of the image. A normalised Radon transform $\mathcal{R}_N f$ is constructed by pointwise dividing the original transform $\mathcal{R}f$ by a Radon transform $\mathcal{R}I$ that captures the distance bias,

$$\mathcal{R}_N f = \mathcal{R}f \,/\, \mathcal{R}I.$$

Here I is an image that has the same domain as the input image f but with values all set to one. Importantly, it is only after this normalisation step that the radial derivative is taken. The problem with applying the derivative in the Fourier domain is that the derivative is necessarily applied *before* normalisation can be carried out.

In order to reconcile this problem, a partial derivative approach was investigated that combined the partial derivatives of both the regular Radon transform and the transform $\mathcal{R}I$ used for normalisation,

$$\frac{\partial(\mathcal{R}_N f)}{\partial\rho} = \frac{(\mathcal{R}I)\frac{\partial(\mathcal{R}f)}{\partial\rho} - (\mathcal{R}f)\frac{\partial(\mathcal{R}I)}{\partial\rho}}{(\mathcal{R}I)^2}.$$

Unfortunately, we found the approach in practice did not work very well, as it seemed to introduce new errors that outweighed any potential benefit gained from using differentiation in the Fourier domain. Moreover, it reduced speed and increased memory usage each by a factor of two.

3   Discretisation Considerations

In certain situations, the result for the Radon transform can exhibit strong evidence of "noise". Examples are shown in Fig. 5 taken from the sinusoidal Radon transform result in Fig. 3. Points along the line are of greylevel approximately 1.0, which corresponds to the greylevel of the single point in the input image. The noise on either side of the line is of order ±0.1 and fades to a background value of less than ±0.001. The noise oscillates regularly between positive and negative values, hence generating a noticeable "checkerboard" pattern. At certain key values for θ, for example zero or ninety degrees (the top row in the image on the left), the noise reduces to zero. It was important for us to find out the reason for this noise, as it seemed too large to be merely rounding errors, or some other acceptably small error. We should note here however that this example is quite a stringent test for the Radon transform because the input image consists of just a single point. When applying the Radon transform to any typically image such as a SAR image, the noise would be not be noticeable at all (it would tend to be reduced in variance by the square root of the number of neighbouring non-zero points in the image).
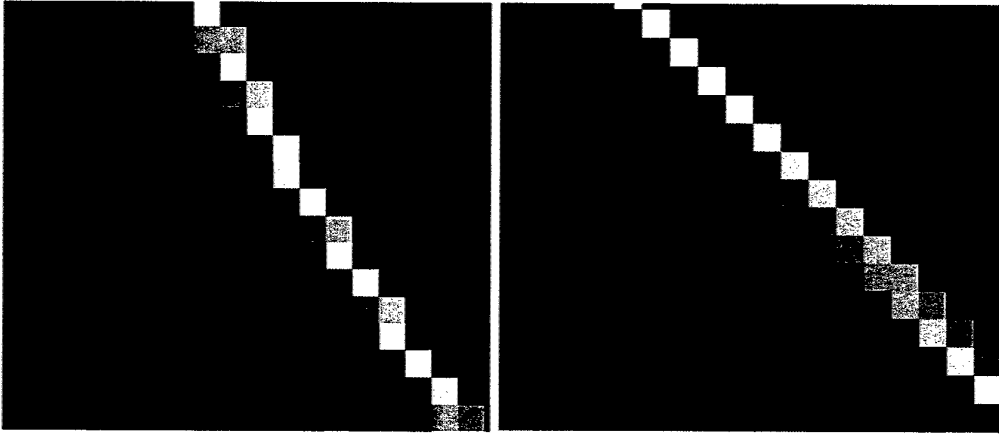
Figure 5: Examples of noise in the Radon transform.

Newsam [5] notes that the Radon transform in the continuous domain is approximated in the discrete domain by the following formulation,

$$(\mathcal{R}f)(\rho_s, \theta_q) \approx \sum_p w_p e^{2\pi i \rho_s \omega_p} F(\omega_p, \theta_q).$$

In order to minimise discretisation error, the task is to choose the quadrature points ω_p and associated weights w_p so as to minimise the error in the approximation. A solution for the weights w_p is proposed and properties for the underlying grid points ω_p recommended. However, as the noise seems highly structured and dependent on the amplitude of the signal (the noise is much diminished in the background), we sought another (perhaps related) explanation for the noise. We thought the quadrature errors must play a part in generating noise, but perhaps at some more acceptable low level. We detail our approach in this section.

3.1 Convolution by Sinc Function

We propose that the noise so obvious in Fig. 5 may in fact be due to the finite extent Ω of the 1-D slice through the frequency image $\mathcal{F}_{\mathbf{x}}f$ (as shown in Fig. 2). The finite slice extent could be considered as a multiplication of the actual underlying data $\mathcal{F}_{\mathbf{x}}f$ by a box car function $\Pi_{\Omega}(\omega)$ of width Ω. The inverse Fourier transform $\mathcal{F}_{\omega}^{-1}$, which operates on 1-D slices through $\mathcal{F}_{\mathbf{x}}f$, is then given by

$$\mathcal{F}_{\omega}^{-1}\left[\Pi_{\Omega}.\mathcal{F}_{\mathbf{x}}(\omega,\theta)\right] = \Omega \text{ sinc } (\Omega\pi\rho) * F(\rho) = \text{ sinc } (\pi\rho) * F(\rho).$$

So, our desired result $F(\rho)$ in the spatial domain is effectively convolved by the function sinc $(\pi\rho)$ (note that in our implementation, the spatial extent $A = N$, the number of sample points, and so the reciprocity relationship $A\Omega = N$ dictates that $\Omega = 1$).

On the left of Fig. 6 is shown a plot of the sinc function that is centered on a pixel. If a point in the Radon transform that is also centered directly on pixel is convolved by this function, we would expect the contribution from the sinc function to disappear. This is because the zero crossings of the sinc function, at integral values of ρ, coincide with the grid of points in ρ. On the other hand, if a point in the Radon transform is not centered on a pixel, we would expect the contribution from the sinc function to emerge, as shown to the right of Fig. 6.
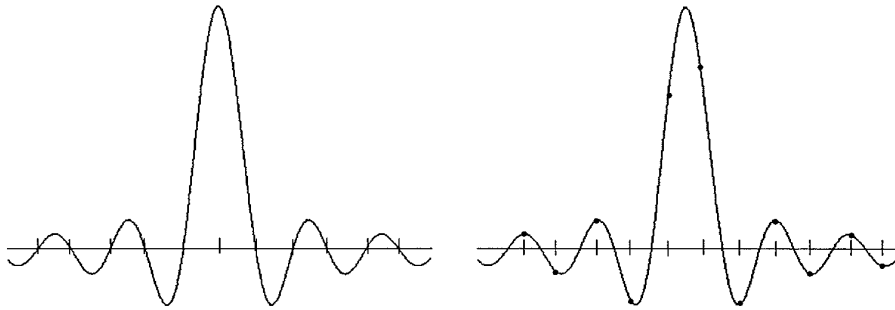

Figure 6: Sinc function and underlying grid spacings. Left: Sinc function is centered on a pixel; effectively the contribution is nullified. Right: Since function is not centered and contribution becomes apparent.

We consider that the sinc function is always present in our result for the Radon transform, but only becomes apparent when values in the continuous Radon transform are not centered on points in the discrete grid of the Radon transform. For angles such as zero and ninety degrees, the results in the Radon transform are always centered on the grid, because the distance ρ in the input image is always integral in the x ($\theta = 0$) and y ($\theta = 90$) direction. Hence, we see no evidence of the sinc function for these angles. In contrast, when $\theta = 45$ or 135 degrees for example, distance in the input image is measured in steps of $1/\sqrt{2}$ and the results in the Radon transform are never centered on the grid. However, scaling Ω by a factor of $\sqrt{2}$ sets the size in the spatial domain to $1/\sqrt{2}$. Thus the results in the Radon transform will be centered on the grid at 45 degrees and evidence of the sinc function will disappear. At the same time though, it will emerge at 0 and 90 degrees. This was borne out by experimental results and would seem to support our theory regarding the sinc function.

3.2 Different Sampling Strategies

We have experimented with the two sampling strategies shown in Fig. 7, amongst others. The sampling strategy on the left was obtained through a scaling of Ω by the factor $\sqrt{2}$ discussed above. In order to obtain values for $\mathcal{F}_\mathbf{x} f$ outside its domain, we tried both zero padding and periodic repetition of data (in x and y directions). Both sets of results showed the expected shift in noise patterns; the noise was present for all angles except 45 and 135 degrees. We note here that though we are utilising the entire domain of $\mathcal{F}_\mathbf{x} f$ (including the corners that were hitherto ignored), this does not seem to improve the noise in the rest of the Radon transform appreciably.
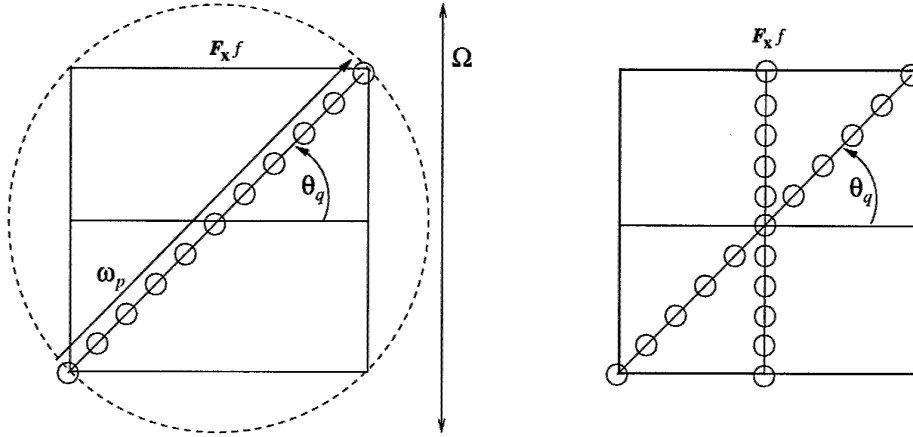

Figure 7: Different sampling strategies used in the Radon transform. Left: Sampling over a circle with diameter equal to the diagonal of the image. Right: Sampling up to the border of the domain of $\mathcal{F}_\mathbf{x} f$.

The sampling strategy on the right of Fig. 7 was obtained by scaling Ω to the borders of the domain of $\mathcal{F}_\mathbf{x} f$ using the function:

$$\alpha(\theta) = \begin{cases} 1/\cos(\pi/2 - \theta) & \text{for } \theta > \pi/4 \text{ and } \theta < 3\pi/4 \\ 1/\cos(|\theta|) & \text{otherwise.} \end{cases}$$

The results show that for a point placed at random in the input image, the scaling removes the sinc function for the angles 0, 45, 90 and 135, but not for any other angles. From a number of experiments we have performed, it does not seem possible to choose a simple and practical angle-dependent scaling measure in order to make the sinc function disappear over the entire Radon transform.

3.3 The Hanning Window

A technique that has worked well is to apply a weighting function in the frequency domain to remove the effect of the sinc function. Specifically, we define a Hanning window

$$h(n_\omega) = \frac{1}{2}(1 - \cos(2\pi n_\omega/N)), \quad n_\omega = 0, 1, \ldots N - 1,$$

and apply it as a multiplication to the set of points $(\mathcal{F}_\mathbf{x}f)(\omega, \theta)$ before the inverse 1-D DFT it taken. An example is shown in Fig. 8, where we get a visually pleasing image with no noise evident. The Hanning window does however have the unavoidable effect of smoothing out or broadening the response, as is evident in this example with the line spreading out over several pixels.
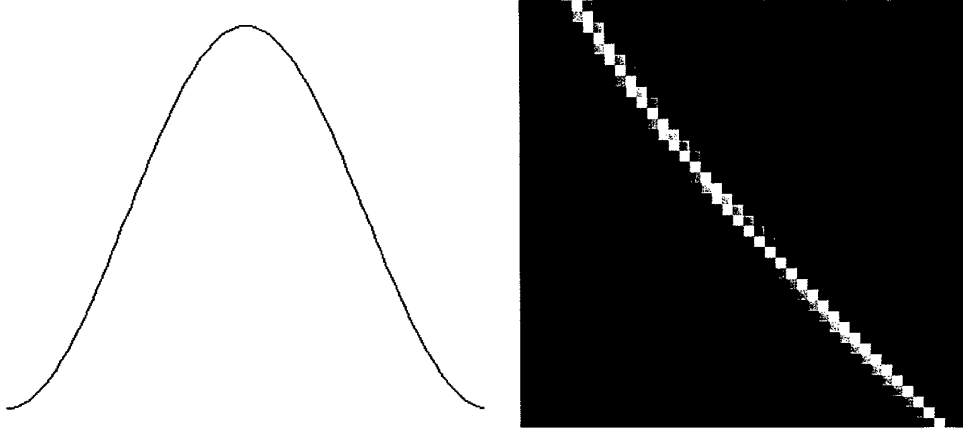

Figure 8: Hanning window used in the frequency domain to dampen effects of box car function. Left: Shape of curve used. Right: Result using window.

Another simple technique that worked well, at least visually, is to take the intensity (square root of the magnitude) of the (complex) Radon transform, rather than just the real component. We expect the Radon transform to be purely real, but there is a small contribution from the imaginary component which would normally be of value around the noise level. Taking the intensity flips negative contributions from the sinc function to positive, thus producing an apparently smoother result. One drawback however is that we lose genuinely negative information in the Radon transform; although the Radon transform should be real there is no requirement that it should be positive valued (negative values would come from negative values in the input image). An example result is shown in Fig. 9. The result now looks much more acceptable: the sinc noise normally present at points that do not fall on the discrete grid now appears to be an acceptable smoothing of values into adjacent bins.

4   Matched Filtering

In this section, we describe the use of matched filtering [4] to accentuate the characteristic butterfly shape that represents a line in the input image (as previously shown in Fig. 4). In its simplest form, matched filtering is given by the correlation f between an image g and filter h:

$$f(x) = g * h = \int_{-\infty}^{\infty} g(u - x)h(u)du. \tag{3}$$

This section details an analytic form for a discrete version of h which is derived from the expected shape of the butterfly in the discrete Radon transform. A similar analysis of the
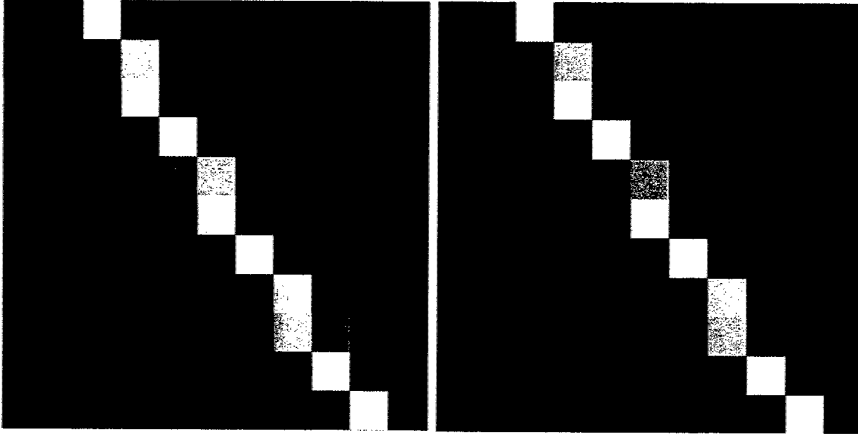
Figure 9: Left: Radon transform without using intensity. Right: Radon transform using intensity.

distribution of line peaks was detailed in [3], for the continuous domain Hough Transform (qualitatively very similar to the Radon transform). Our analysis differs in that we study lines passing right through the image and we are interested in working in the discrete domain where we have noticed marked digitisation effects.

4.1 Matched Filter Shape

We begin by looking at the form of the Radon transform for single points in the image, as we can study the bounds of the butterfly shape by examining the two end points of the line. We can simplify the study somewhat because we are ultimately interested in detecting long lines that pass right through the image under study (*i.e.* parts of long curvilinear trails). Figure 10 shows how the points are restricted. We consider points lying on the circle with radius ρ_M shown in the figure, where a pair of points A and B separated by an angle Δ on the circle define the extremes of a line passing through the image.
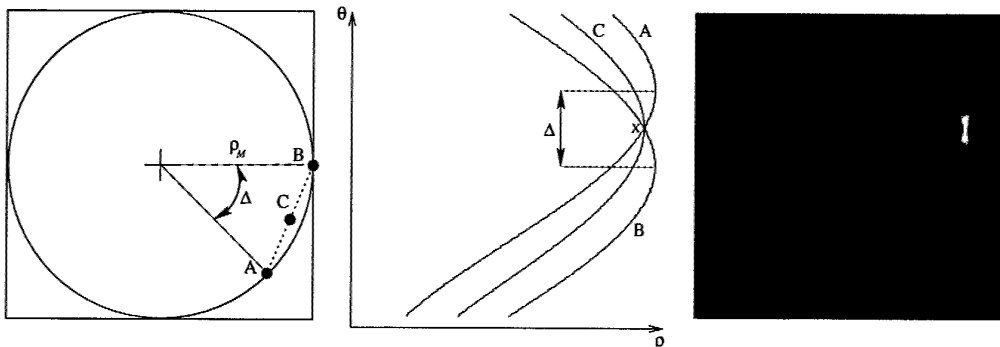

Figure 10: Generating butterfly shape in Radon transform. Left: Three points in the image space. Middle: corresponding Radon transform. Right: Actual result for line between two points A and B shown.

The Radon transform for the three points A, B and C is shown in the middle of the figure. We can consider the two curves A and B as the extremes of the butterfly shape generated for the line $\vec{AB}$, because all points along the line generate a sinusoid that falls within their bounds (for example line C). We know this is true because all points along the line generate a sinusoid that is coincident with the point X indicated (the point we wish to detect in order to extract the line parameters) with magnitue less than ρ_M. An actual experimental result for all the points along the line is shown to the right of the figure; this result is zero except between the two extreme sinusoidal curves.

So, a line passing through the image may be represented by its two extreme points that cross a circle of radius ρ and these points are separated by Δ. In the Radon space, this line will have a butterfly shape bounded by two sinusoidal curves of amplitude ρ_M and separated in phase by Δ. This situation is represented in Fig. 11 (here the axes have been inverted because for the time being we wish to consider ρ as a function of θ). It is immediately clear then that the butterfly shape is not constant but will change shape depending on the value of Δ, which in turn is proportional to the amount of line visible within the circle. As our matched filtering approach is based on the use of a single representative filter, we must accept then that there will be a degree of inherent mis-match error due to variations in the amount of visible line. In the figure, curve A is represented by the sinusoid $\rho = \rho_M \cos(\theta)$ and curve B by $\rho = \rho_M \cos(\theta - \Delta)$. If Δ is measured in radians, at the point of intersection the slope of A is given by $m_A = -\rho_{\max} \sin(\Delta/2)$ and the slope of B is given by $m_B = \rho_{\max} \sin(\Delta/2)$. On the right of Fig. 11 is plotted the slope m_B as a function of Δ.
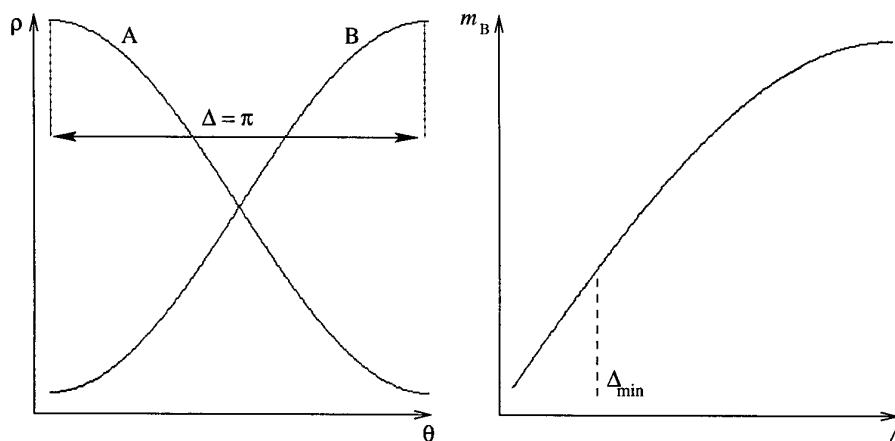
Figure 11: Left: Ideal butterfly shape in Radon transform, corresponding to a line that passes through the centre of the image. Right: Graph of slope m_B versus Δ.

For simplicity, we have decided to set Δ to π radians, which corresponds to a line that passes directly through the centre of the image (this produces the "ideal" butterfly shape shown to the left of Fig. 11). Other values for Δ could be used, for example the "most likely" value could be derived based on the probability distribution for lines in the window. We note that in practical implementations, we can put a lower limit $\Delta_{\min}$ (as shown on the right) on the range of Δ by acknowledging the use of overlapping windows to detect lines. Any matched filter we define that is insensitive to lines with $\Delta < \Delta_{\min}$ would be

picked up by the next overlapping window. In order to implement this, we can show that the spacing between windows must be less than equal to $2\rho_M \cos(\Delta_{\min}/2)$, where ρ_M is the radius of the circular window used.

For practical purposes, we are only interested in defining the matched filter in a small region focussed around the peak of the butterfly shape. If this region is small enough we may reasonably approximate the bounds of the matched filter with straight lines, as shown to the left of Fig. 12. The actual slope values may be derived by setting $\Delta = \pi$ in the derivative formulations $\pm\rho_{\max}\sin(\Delta/2)$, yielding slopes of $\pm\rho_{\max}$. In our discrete setting, $\theta = [0, \pi)$ is distributed over `theta` pixels and the slope is therefore given by $\rho_{\max}\pi/$`theta`. We are at liberty to set $\rho_{\max}$ and `theta` however we choose, and by setting `theta`$/\rho_{\max} = \pi$, the slope becomes one (or 45 degrees). From now one, we assume that the slope has been set to this value because it simplifies the discussion with no loss of generality.
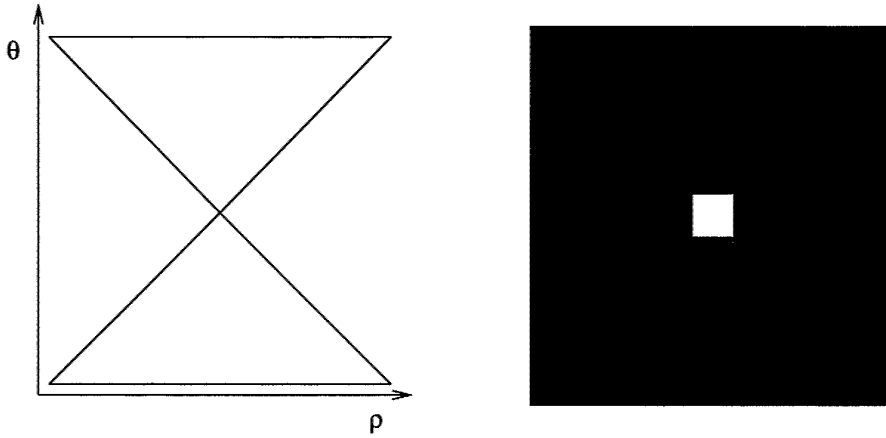

Figure 12: Left: Approximation of butterfly shape in central region using straight lines. Right: Actual data from image showing butterfly shape in central region.

4.2 Matched Filter Kernel Values

Using the bounds of the butterfly shape as suggested above, the next step is to compute the actual kernel values for the matched filter. To the right of Fig. 12 is shown some actual data where the slope at the central region is ±45 degrees (through appropriate choice of `theta` and $\rho_{\max}$). This was generated from an input image consisting of a single line of width one and intensity 128 passing through the centre of the image. As discussed in Section 3, the checkerboard pattern of noise is present (we have not applied a Hanning window or used intensity as the output), but we shall ignore it in the following discussion. As expected, the central pixel value is 128, corresponding to the length of the line, and pixels outside the butterfly shape are zero. For other pixels within the butterfly shape, we find the value 128, or line length, has been distributed over the pixels in each row of the butterfly shape. That is, the sum of values in any given row is equal to the length of the line. This observation accords with what we expect from the Radon transform (which is in essence a line integral formulation).

However, an interesting feature we also observe is that for any given row of values in the butterfly shape, the end points have half the value of the central points in the row. Figure 13 illustrates how this comes about, for three separate cases. The line in the input image is represented by the horizontal rectangular box and the black circles represent (ρ, θ) points where the line is sampled. The top case represents the situation when the angle θ coincides with the orientation of the line. Here the entire volume of the line is put into a single bin (ρ, θ); this is the central location of the butterfly shape. In the next case, the volume of the line is divided into three samples. Recall that we have set up our butterfly shape so that we have a sample point at either end of the line. We can see then that because the two end samples are at the very ends of the line, only a quarter of the volume of the line will be attributed to these points, whereas the central point will be get half the volume. The same follows for the third case, where there are five samples along the line. Here we have the end points are attributed with half as much line volume as all the other samples in the line.
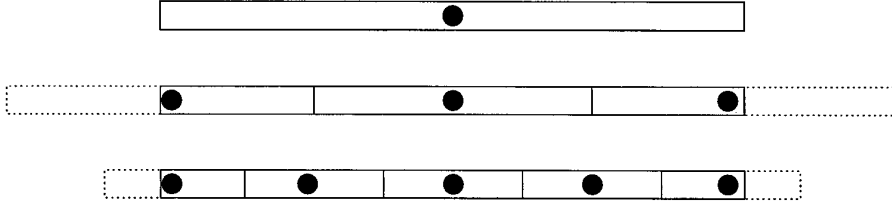
Figure 13: Division of line weight over (ρ, θ) in discrete setting. Top: Entire line volume goes to single value of (ρ, θ). Middle: Line volume is distributed over three (ρ, θ) values. Bottom: Line volume is distributed over five values of (ρ, θ). pixels.

To formalise the above observations, we consider first the bottom right hand side of the match window, where the centre of the match window is at coordinates $(0, 0)$. By considering $x, y >= 0$ we have the matched filter $M(x, y)$ kernel values given by:

$$M(x, y) = \begin{cases} 1 & \text{for } x = y \text{ and } y = 0 \\ 1/4y & \text{for } x = y \text{ and } y > 0 \\ 1/2y & \text{for } x < y \\ 0 & \text{for } x > y. \end{cases}$$

To get the full window $M(x, y)$, we simply reflect this result in the x and y axes. The final result is shown in Fig. 14, where the end points of each row have exactly half the value of the central points (apart from the central row). In practice, the match window can be normalised by subtracting off the mean value, so that the sum is zero, or dividing by the mean value, so that the sum is one. We currently implement the former of these normalisations. It should be pointed out that we do not use the more sophisticated normalised correlation function [4], because the resulting filter is non-linear and we would not be able to use any linearity assumptions (as we do below).

The actual extent of the matched filter is limited by our approximation of the bounds of the filter with straight lines, as this approximation breaks down as the extent increases. A simple measure of the limit of the extent can be defined as the point where the straight line approximation departs significantly, say one pixel, from the underlying sinusoidal form. This point depends on the dimensions theta and ρ_{max} specified; for our purposes we have found a matched filter of extent 11 by 11 to be suitable.

1/12	1/6	1/6	1/6	1/6	1/6	1/12
0	1/8	1/4	1/4	1/4	1/8	0
0	0	1/4	1/2	1/4	0	0
0	0	0	1	0	0	0
0	0	1/4	1/2	1/4	0	0
0	1/8	1/4	1/4	1/4	1/8	0
1/12	1/6	1/6	1/6	1/6	1/6	1/12

Figure 14: Values defined within the match window.

Finally, we should point out that the above observations hold for a butterfly shape that is actually centered on a pixel, *i.e.* for which the parameters (ρ, θ) fall on the discrete grid. Naturally most lines will not have this convenient form. Moreover, as pointed out before, the visible line length in the window is variable and this varies the bounding sinusoids of the butterfly shape as parameterised by Δ. Ultimately, we are really relying on the inherent robustness of the matched filter to allow for these variations in shape while still providing sufficient matching power.

4.3   Lines of Arbitrary width

When considering lines of arbitrary width, we are fortunate in that the Radon transform and matched filter (in its simple unnormalised form, as in Eq. 3) are both linear operators. We consider the simple scenario that a line of width w can be represented as a sum of parallel lines of width one. The Radon transform for the whole line is then simply the sum of Radon transforms for the individual lines. As each individual line is manifest as a butterfly shape in the Radon transform, the result for the line of width w is simply a sum of these butterfly shapes. Moreover, we can apply a matched filter to this resulting shape and expect the result to be simply a sum of matched filters. An example is shown in Fig. 15, where the input image on the left contains a line of width 10 and the Radon transform on the right is a sum of adjacent butterfly shapes for a single line.

An example illustrating the application of the matched filter is shown in Fig. 16. Here the input image is a faint vertical line of width 3 and intensity of one added to a distribution of random Gaussian noise. The middle figure shows the result for the Radon transform. The peak corresponding to the line is faintly visible in the centre of the figure. The peak has an intensity of approximately 16 greylevels and the background ranges from approximately 14 to 15 grey levels. Shown on the right is the result for the matched filter. Here the peak has a value of approximately 5, all other peaks have a value of approximately 3 and the background value is less than zero. So clearly for this example, the convolution filter has managed to bring out the peak from the background values.
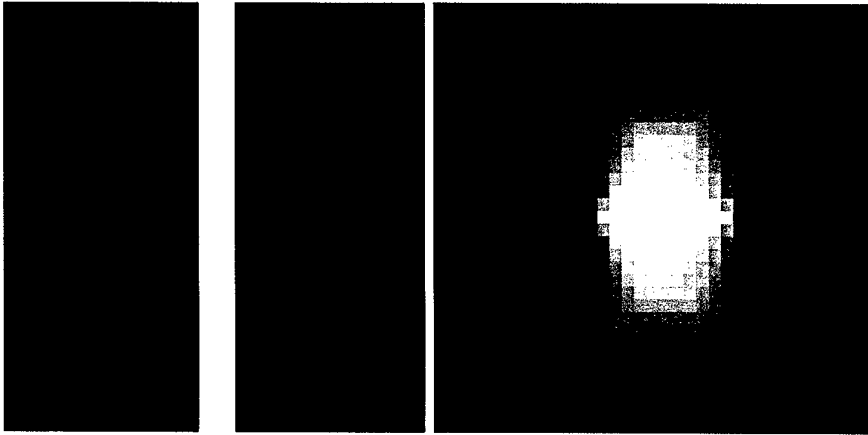
Figure 15: Radon transform of wide line. Left: input image. Right: Radon transform.
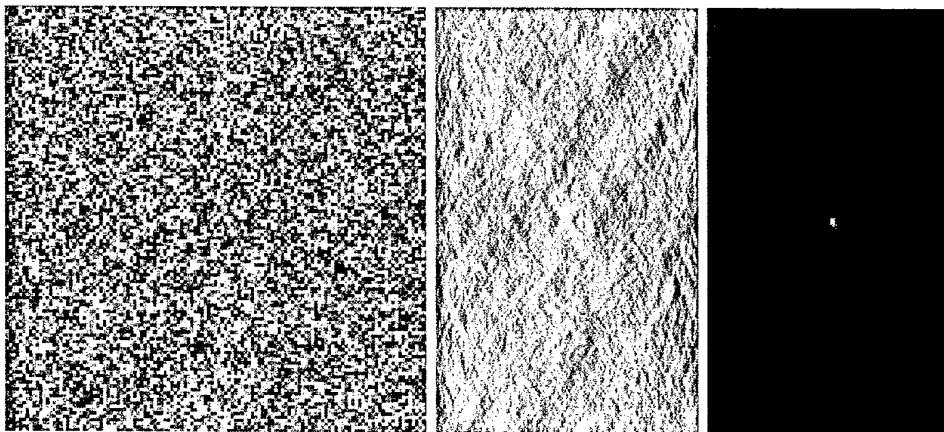

Figure 16: Applying a matched filter to test image. Left: Input image with strong Gaussian noise. Middle: Result from Radon transform. Right: Result from matched filter.

The matched filter discussed above was derived for the case of a light line on a dark background, yielding a butterfly shape that is of value greater than the background. Figure 17 shows a counter example. Here the image has a dark line on light background, and the corresponding Radon transform shows the dark butterfly shape that is produced. Applying the positive valued matched filter produces the result to the right of the figure, which is essentially the inverse of the result shown in Fig. 16. The bright peak in the radon transform has now become a dark hole of value around -10. This too can be thresholded to find the parameters (ρ, θ) for the line. Both light and dark lines can be detected by combining threshold results using two thresholds. Alternatively, the results for two separate matched filters, one with positive kernel values and one negative kernel values, can be combined using, for example, a maximum operator.
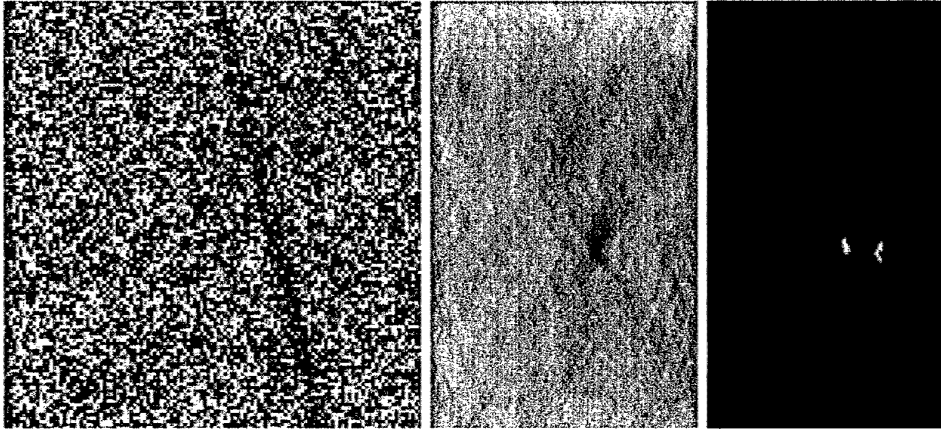
Figure 17: Applying a positive matched filter to image with dark line. Left: Input image. Middle: Result from Radon transform. Right: Result from matched filter.

4.4 Matched Filtering the RDRT

A matched filter may also be applied to the RDRT to enhance peak detection. As we might expect, a suitable matched filter for RDRT is given by the derivative of the matched filter defined above. Use of the RDRT would seem to circumvent any inconvenience in dealing with both light and dark lines, because the RDRT takes the absolute value of the radial derivative. In this case, light and dark lines look qualitatively the same in the RDRT and we can use a single matched filter to perform matching. However, we note here that the use of the absolute value renders the RDRT a non-linear operator, as in general $|f + g| \neq |f| + |g|$. Therefore, we can not use a linearity property to generalise the treatment of the RDRT to lines of arbitrary width as we have done for the Radon transform. In practical terms, what this means is that there is not a single consistent shape we can use to match the RDRT for lines of arbitrary width; in Fig. 18 is shown RDRT butterfly shape results for lines of width 1, 3 and 5, shown left to right respectively. One matching process we have tried is to match each side of the RDRT butterfly shape with separate matched filters while ignoring the central region. This process finds the leading and trailing edge of lines in the image and is independent of line width. However, the results we obtained were somewhat disappointing.
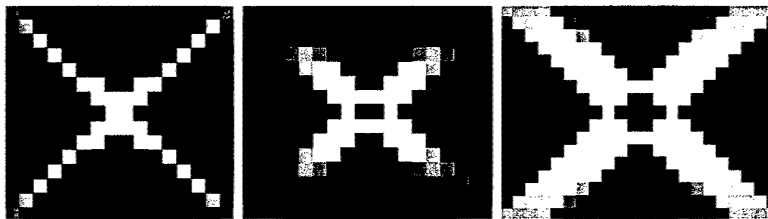

Figure 18: RDRT for lines of varying width. Left to Right: Line width 1, 3 and 5.

An alternative approach is to omit the absolute value in the calculation of the RDRT. We can show that the result for the RDRT then consists of the sum of two butterfly

shapes of opposite phase, one corresponding to the leading edge and one to the trailing edge of the line. Figure 19 shows an example, where the RDRT result without the absolute value is shown in the middle for the input image on the left. The matching process must now either combine two matched filters with opposite phase or two thresholds in order to enhance and locate the peaks. We note here that the use of the RDRT has the advantage of yielding both of the edges of the line. In contrast, the Radon transform, shown to the right of the figure, will yield only a single central peak for the line.
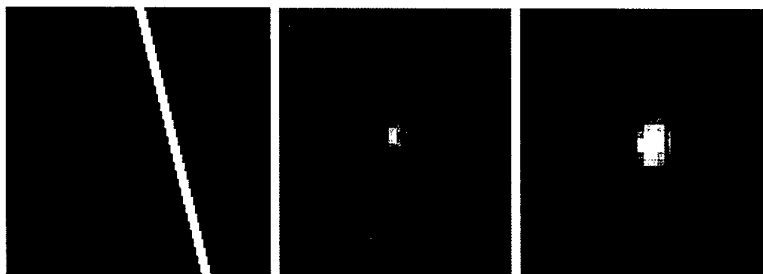

Figure 19: RDRT results for light line. Left: Input image with line of width 5. Middle: RDRT without absolute value, showing leading and trailing edges. Right: Standard Radon transform result.

A final method we have looked at is to apply the radial derivative *after* matched filtering of the Radon transform. This approach offers the edge detecting ability of the RDRT while still allowing the use of the absolute value to yield positive and negative phases simultaneously. An example is presented in the following section.

4.5   Some Examples

Shown in Fig. 20 is a result for the image previously shown in Fig. 1. Here the Radon transform, at top, shows the expected evidence of lines in the image. The peak has a value of approximately 0.6, the dark hole approximately 0.2 and the background varies in the range 0.3 to 0.5. The result from the matched filtering, using a positive value matched filter, is shown at the bottom of the figure. Here, the peak now has a value around 0.37, the hole (which is not visible against the dark background) has a value around -0.35 (*i.e.* opposite in phase to the peak) and the background now sits in the range ±0.1. From this example at least we can see that the matched filter is able to enhance peaks in the Radon transform (and holes using a second threshold). It could perhaps best be described as a mechanism for background normalisation that sets the background to zero and lifts brights peaks up while pushing dark peaks down in value.

Shown in Fig. 21 is a more challenging example. A portion of the 512 by 512 image used is shown at the top, where the background noise is stronger and larger in size and the road edges are more diffuse. Directly below the image is shown the result for the Radon transform. The result shows two large responses corresponding to the two obvious linear objects in the input image. Interestingly, the features in the Radon transform look rather like edge responses, with light then dark features side by side. This is because each line in the image consists of a bright line of several pixels wide followed by a much wider dark
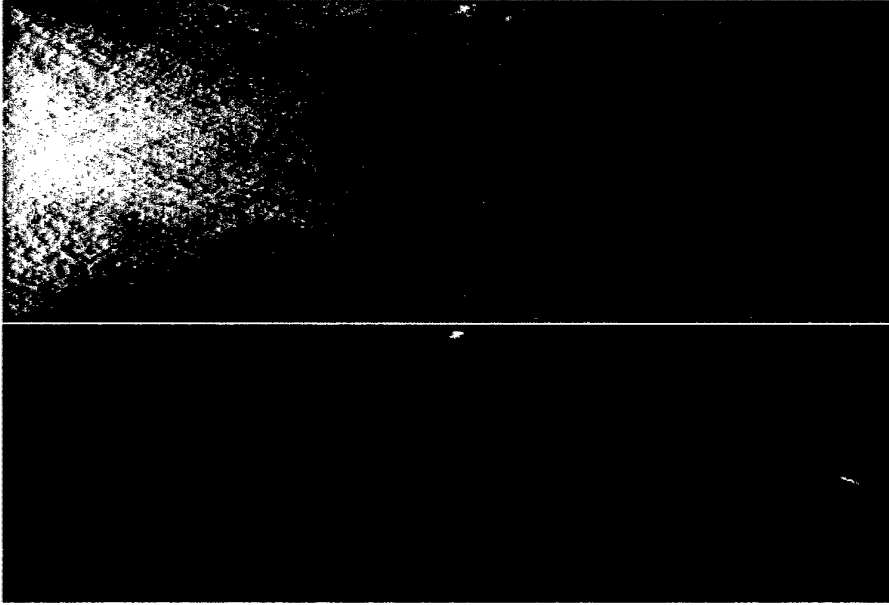
Figure 20: Applying the matched filter to real image data. Top: Result from Radon transform. Bottom: Results after applying matched filter.

line that fades into the background; effectively two linear features of opposite phase side by side. Shown below the Radon transform is the RDRT, where we have used the absolute value. The use of the absolute value has the visual effect of fortifying peaks by combining strongly positive and strongly negative values into a single strong peak; without using the absolute value the result would be (visually at least) unsatisfactory. Below this is shown the matching result on the RDRT, and we can see here that the matched filter has not really improved the situation, but rather seems to have made it worse. We have found matching results to be generally unsatisfactory when used on the absolute value of the RDRT. Finally, at the bottom of the figure is the absolute value of the radial derivative applied to the matched filtering of the Radon transform. When compared to the RDRT, we have a better result with background noise much reduce and peaks highlighted as desired.

4.6 Analysis

To measure the effectiveness of the matched filter technique for detecting curvilinear features, a quantitative measure of signal to noise was used, based on a two-parameter adaptive threshold prescreener for detecting targets in SAR imagery. This filter measures target signal to noise using the formula

$$k = (x - \mu)/\sigma$$

where x is the intensity of the point being tested, and μ and σ are the estimated mean and standard deviation of the background pixels. This quantity k was calculated for both of the curvilinear features seen at the top of Fig. 21 for five different methods of processing
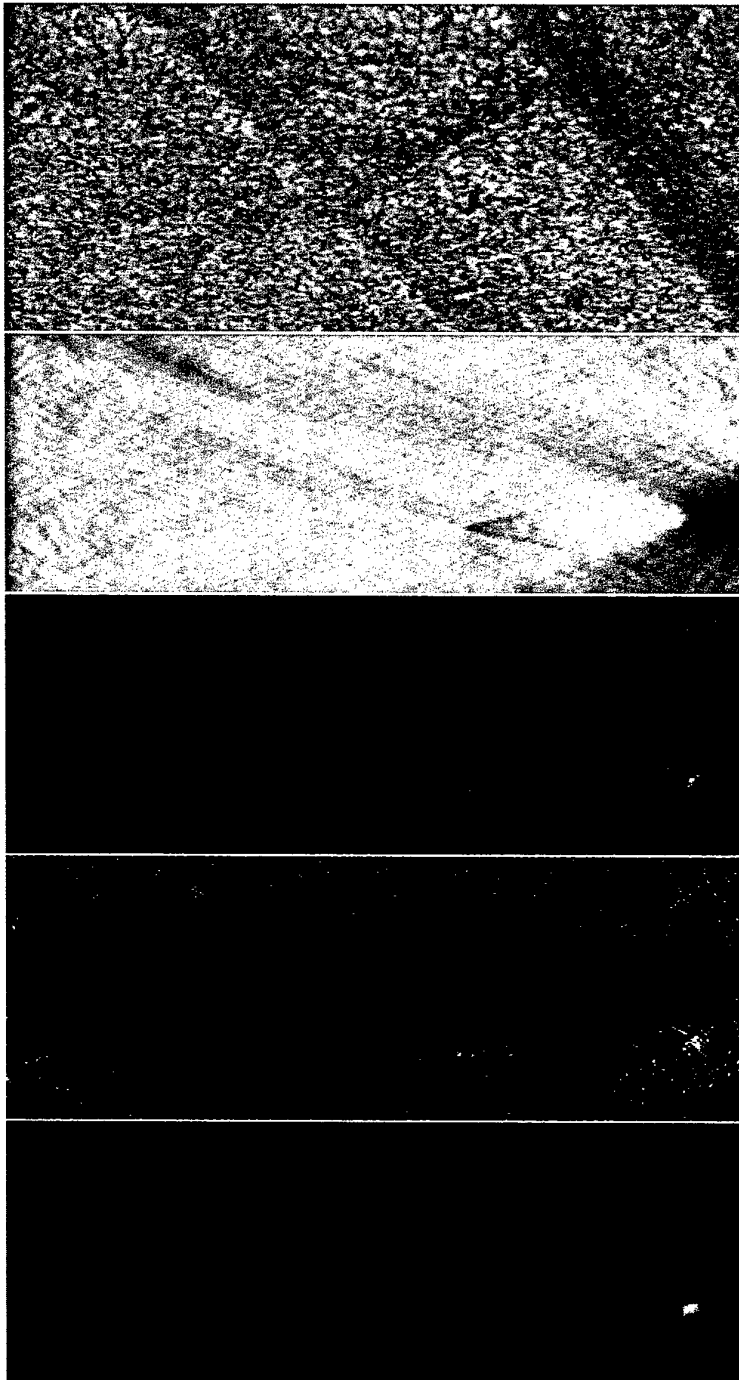
Figure 21: Top to Bottom: Portion of input image, Radon transform, RDRT, Matched filter applied to RDRT, radial derivative of matched filter applied to Radon transform.

Method	SNR Feature 1	SNR Feature 2
Standard RDRT	7.72	10.5
RD of matched filtered RT (37×9)	13.5	30.1
RD of mean filtered RT (37×9)	6.68	29.5
RD of median filtered RT	19.7	29.8
Median filtered RD of median filtered RT	26.8	53.9

Table 1: *Signal to noise ratios for five different methods of processing the RDRT.*

the RDRT. The first method was the default RDRT as shown in Figure 21, while the second included the matched filtering stage prior to the calculation of the derivative, as shown at the bottom of Fig. 21. Since the matched filter acts as a smoothing operator, it is also useful to determine what the effect of the butterfly shape of the filter has on the detection SNR, as opposed to a filter of the same size with uniform values. To do this, a filter of the same size (37×9) with equal pixel weights was used to smooth the radon transform, and then first differences were calculated. The SNR results are shown in the third line of Table 4.6.

The above table shows that the matched filter (in the second row) is doing a better job of detecting curvilinear features than a uniformly weighted filter (third row). For a second comparison, the matched filter result was compared against that using a standard image processing noise reduction technique; the median filter. In this case, the median filter was applied prior to the calculation of the radial derivative. Since the derivative was required only in the radial direction, a 1D median filter was applied in this direction. The performance of the median filter was found to increase with window size, and the results for a 20 pixel median filter are shown in the fourth row of Table 4.6 (although higher SNRs were obtained for even larger window sizes). Here, the weaker linear feature had an improved SNR, while the strong feature's SNR was mostly unaffected, which indicates that the median filter may be more useful than the matched filter for faint trail detection. It should be pointed out however that this is only a single image, and that statistically meaningful results would require a larger set of images to be tested. This, however, is outside the scope of this report.

The matched filter method produced a more blurry estimate for the RDRT than did the median filter, which tends to preserve edges. As a quick test whether the SNR could be further improved by blurring the RDRT, a 2D median filter (chosen to be 6×6, since this gave blurring visually similar to the matched filter) was applied to the RDRT estimate used in line 4 of the table. As expected, the SNR was significantly increased and appears to be much better than that achieved for the matched filter, although again, the result is only anecdotal. It is speculated that the matched filter method could be further improved by applying it to a median filtered Radon transform. Again, testing of this nature is outside the scope of the report, although more detailed testing may be attempted in follow-up work.

5 Conclusion

This report has discussed the ADSS implementation **radon** of the Radon transform, which uses non-equispaced DFTs to provide a relatively quick and direct implementation of the transform.

The report has offered an analysis and experimental results for two paths of study: discretisation errors and matched filtering. In terms of future work, there is much that could be done. The most obvious direction in terms of discretisation errors, or more particularly the "checkerboard" noise we observe, is to pursue the recommendations in [5]. The perspective our report has taken is that the noise is inherent to the implementation we are using and can only be ameliorated but not removed altogether. Our most successful approach was to use a Hanning window of weights in the frequency domain. The approach discussed in [5] provides a rigorous framework for quantifying and controlling discretisation errors. Control over the errors is captured in the value and placement of quadrature weights applied in the frequency domain. The use of the Hanning window involves a similar implementation, though our choice of Hanning weights and placements could be considered suboptimal. However, pursuing the approach further could prove quite a challenging project. For example, the choice of weighting placements is still an outstanding theoretical problem.

The matched filter we have proposed in this report would also benefit from further work. The matched filter has only been applied to a few images and, while the results have been encouraging, a validation process is required to determine to what extent the filter improves peak detection. The filter proposed is static and a number of assumptions were made in order to derive a simple form. These could be relaxed in order to provide a more generalised and robust matched filter. Another more challenging direction would be to explore a dynamic matched filtering approach, which might capture variations such as line width in the RDRT and shape variations arising from variable line position.

References

1. Bracewell R. N. *The Fourier Transform and its Applications, Third Edition.* McGraw Hill, 2000.

2. Cooke T. A Radon Transform Derivative Method for Faint Trail Detection in SAR Imagery. pages 31–34, Perth, Australia, 1999. DICTA '99.

3. Furukawa Y. and Shingagawa Y. Accurate and Robust Line Segment Extraction by Analysing Distribution around Peaks in Hough Space. Computer Vision and Image Understanding, 92 (2003), pages 1-25.

4. Gonzalez R. C. and Woods R. E. *Digital Image Processing.* Addison Wesley, 1993.

5. Newsam G. N. Controlling Errors when Calculating the Radon Transform by Discrete Fourier Transforms on Non-Uniform Grids, DSTO Note, August 2003.

6. Potts D., Steidle G. and Tasche M. Fast Fourier transforms for nonequispaced data: A tutorial, In *Modern Sampling Theory: Mathematics and Applications, J.J. Benedetto and P. Ferreira (Eds.)*, pages 249-274, 2000.

7. Redding N. J. Design of the Analysts' Detection Support System for Broad Area Aerial Surveillance. Technical Report DSTO–TR–0746, DSTO Technical Report, 1998.

DISTRIBUTION LIST

Implementation of the Radon Transform Using Non-equispaced Discrete Fourier Transforms

Ronald Jones, Tristrom Cooke and Nicholas J. Redding

	Number of Copies
DEFENCE ORGANISATION	
Task Sponsor	
DGISREW	1
S&T Program	
Chief Defence Scientist	
FAS Science Policy	
AS Science Corporate Management	1
Director General Science Policy Development	
Counsellor, Defence Science, London	Doc Data Sheet
Counsellor, Defence Science, Washington	Doc Data Sheet
Scientific Adviser to MRDC, Thailand	Doc Data Sheet
Scientific Adviser Joint	1
Navy Scientific Adviser	Doc Data Sheet and Dist List
Scientific Adviser, Army	Doc Data Sheet and Dist List
Air Force Scientific Adviser	1
Scientific Adviser to the DMO M&A	Doc Data Sheet and Dist List
Scientific Adviser to the DMO ELL	Doc Data Sheet and Dist List
Director of Trials	1
Information Sciences Laboratory	
Chief, Intelligence, Surveillance and Reconnaissance Division	1
Research Leader, Imagery Systems	1
Head, Image Analysis & Exploitation	1
Guy Blucher	1
Dr David Booth	1
David I. Kettler	1
Dr Nicholas J. Redding	4
Merrilyn Fiebig	1
Dr Vittala Shettigara	1
Ray Oermann	1

Bob Whatmough	1
Head, Imaging Radar Systems	1
Dr David Crisp	1
Rodney Smith	1
Dr Mark Williams	1

DSTO Library and Archives

Library, Edinburgh	1 and Doc Data Sheet
Australian Archives	1

Capability Systems Division

Director General Maritime Development	Doc Data Sheet
Director General Land Development	1
Director General Aerospace Development	1
Director General Information Capability Development	Doc Data Sheet

Office of the Chief Information Officer

Deputy Chief Information Officer	Doc Data Sheet
Director General Information Policy and Plans	Doc Data Sheet
AS Information Structures and Futures	Doc Data Sheet
AS Information Architecture and Management	Doc Data Sheet
Director General Australian Defence Simulation Office	Doc Data Sheet

Strategy Group

Director General Military Strategy	Doc Data Sheet
Director General Preparedness	Doc Data Sheet

HQAST

SO (Science) ASJIC	Doc Data Sheet

Navy

SO (SCIENCE), COMAUSNAVSURFGRP, NSW	Doc Data Sheet and Dist List
Director General Navy Capability, Performance and Plans, Navy Headquarters	Doc Data Sheet
Director General Navy Strategic Policy and Futures, Navy Headquarters	Doc Data Sheet

Army

ABCA National Standardisation Officer, Land Warfare Development Sector, Puckapunyal	Doc Data Sheet (pdf format)
SO (Science), Deployable Joint Force Headquarters (DJFHQ)(L), Enoggera QLD	Doc Data Sheet

SO (Science), Land Headquarters (LHQ), Victoria Barracks, NSW	Doc Data Sheet and Exec Summ

Air Force

SO (Science), Headquarters Air Combat Group, RAAF Base, Williamtown	Doc Data Sheet and Exec Summ

Intelligence Program

DGSTA, Defence Intelligence Organisation	1
Manager, Information Centre, Defence Intelligence Organisation	1 (pdf format)
Assistant Secretary Corporate, Defence Imagery and Geospatial Organisation	1

Defence Materiel Organisation

Head Airborne Surveillance and Control	Doc Data Sheet
Head Aerospace Systems Division	Doc Data Sheet
Head Electronic Systems Division	Doc Data Sheet
Head Maritime Systems Division	Doc Data Sheet
Head Land Systems Division	Doc Data Sheet
Head Industry Division	Doc Data Sheet
Chief Joint Logistics Command	Doc Data Sheet
Management Information Systems Division	Doc Data Sheet
Head Materiel Finance	Doc Data Sheet

Defence Libraries

Library Manager, DLS-Canberra	Doc Data Sheet
Library Manager, DLS-Sydney West	Doc Data Sheet

UNIVERSITIES AND COLLEGES

Australian Defence Force Academy Library	1
Head of Aerospace and Mechanical Engineering, ADFA	1
Deakin University Library, Serials Section (M List), Geelong, Vic	1
Hargrave Library, Monash University	Doc Data Sheet
Librarian, Flinders University	1

OTHER ORGANISATIONS

National Library of Australia	1
NASA (Canberra)	1
Government Publications Librarian, State Library of New South Wales	Doc Data Sheet

INTERNATIONAL DEFENCE INFORMATION CENTRES

US Defense Technical Information Center	2
UK Defence Research Information Centre	2
Canada Defence Scientific Information Service	1 (pdf format)
NZ Defence Information Centre	1

ABSTRACTING AND INFORMATION ORGANISATIONS

Library, Chemical Abstracts Reference Service	1
Engineering Societies Library, US	1
Materials Information, Cambridge Scientific Abstracts, US	1
Documents Librarian, The Center for Research Libraries, US	1

INFORMATION EXCHANGE AGREEMENT PARTNERS

National Aerospace Laboratory, Japan	Doc Data Sheet
National Aerospace Laboratory, Netherlands	Doc Data Sheet

SPARES

DSTO Edinburgh Library	5

Total number of copies: 52

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA		1. CAVEAT/PRIVACY MARKING	
2. TITLE Implementation of the Radon Transform Using Non-equispaced Discrete Fourier Transforms	**3. SECURITY CLASSIFICATION** Document (U) Title (U) Abstract (U)		
4. AUTHORS Ronald Jones, Tristrom Cooke and Nicholas J. Redding	**5. CORPORATE AUTHOR** Information Sciences Laboratory PO Box 1500 Edinburgh, South Australia, Australia 5111		

6a. DSTO NUMBER DSTO–TR–1576	6b. AR NUMBER 013–091	6c. TYPE OF REPORT Technical Report	7. DOCUMENT DATE April, 2004

8. FILE NUMBER DSTO-E-IN-2004-69	9. TASK NUMBER JTW 01/218	10. SPONSOR DGISREW	11. No OF PAGES 22	12. No OF REFS 7

13. URL OF ELECTRONIC VERSION http://www.dsto.defence.gov.au/corporate/ reports/DSTO–TR–1576.pdf	14. RELEASE AUTHORITY Chief, Intelligence, Surveillance and Reconnaissance Division

15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT

Approved For Public Release

OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SOUTH AUSTRALIA 5111

16. DELIBERATE ANNOUNCEMENT

No Limitations

17. CITATION IN OTHER DOCUMENTS

No Limitations

18. DEFTEST DESCRIPTORS

synthetic aperture radar
image analysis
radon transform

19. ABSTRACT

This report discusses the implementation of the Radon transform in the Analysts' Detection Support System (ADSS) environment using non-equispaced Discrete Fourier Transforms (DFTs). It provides an analysis and experimental results for discretisation error and the use of matched filtering to enhance peaks in the transform.