Phase Error Estimation for Synthetic Aperture Imagery

Steven A. Fortune
B.E. (Hons I)

A thesis presented for the degree of
Doctor of Philosophy
in
Electrical and Electronic Engineering
at the
University of Canterbury,
Christchurch, New Zealand.

June 2005

ABSTRACT

The estimation of phase errors in synthetic aperture imagery is important for high quality images. Many methods of autofocus, or the estimation of phase errors from the measured data, are developed using certain assumptions about the imaged scene. This thesis develops improved methods of phase estimation which make full use of the information in the recorded signal. This results in both a more accurate estimate of the image phase error and improved imagery compared to using standard techniques.

The standard phase estimation kernel used in echo-correlation techniques is shear-average. This technique averages the phase-difference between each ping over all range-bins, weighted by the signal strength. It is shown in this thesis that this is not the optimal method of weighting each phase estimate. In images where the signal to clutter ratio (SCR) is not proportional to the signal amplitude, shear-average does not meet the predicted error bound. This condition may be met by many image types, including those with shadows, distributed targets and varying surface structure. By measuring the average coherence between echos at each range-bin, it is possible to accurately estimate the variance of each phase estimate, and weight accordingly. A weighted phase-difference estimation (WPDE) using this coherence weighting meets the performance bound for all images tested. Thus an improved performance over shear-average is shown for many image types.

The WPDE phase estimation method can be used within the framework of many echo-correlation techniques, such as phase-gradient autofocus (PGA), phase curvature estimation, redundant phase-centre or displaced phase-centre algorithms. In addition, a direct centre-shifting method is developed which reduces bias compared to the centre-shifting method used in PGA. For stripmap images, a weighted phase curvature estimator shows better performance than amplitude weighted shear-average for images with high SCR.

A different method of phase estimation, known as sharpness maximisation, perturbs an estimate of the phase error to maximise the sharpness of the reconstructed image. Several improvements are made to the technique of sharpness maximisation. These include the reduction of over-sharpening using regularisation and an improvement in accuracy of the phase estimate using range-weighting based on the coherence measure. A cascaded parametric optimisation method is developed which converges significantly

faster than standard optimisation methods for stripmap images.

A number of novel insights into the method of sharpness maximisation are presented. A derivation of the phase that gives maximum intensity squared sharpness is extended from a noncoherent imaging system to a coherent spotlight system. A bound on the performance of sharpness-maximisation is presented. A method is developed which allows the direct calculation of the result of a sharpness maximisation for a single ping of a spotlight synthetic aperture image. The phase correction that maximises sharpness can be directly calculated from the signal in a manner similar to a high-order echo-correlation. This calculation can be made for all pings in a recursive manner. No optimisation is required, resulting in a significantly faster phase estimation.

The techniques of sharpness maximisation and echo-correlation can be shown to be closely related. This is confirmed by direct comparisons of the results. However, the classical intensity-squared sharpness measure gives poorer results than WPDE and different sharpness measures tested for a distributed target. The standard methods of shear average and maximisation of the intensity-squared sharpness measure, both perform well below the theoretical performance bound. Two of the techniques developed, WPDE and direct entropy minimisation perform at the bound, showing improved performance over standard techniques.

The contributions of this thesis add considerably to the body of knowledge on the technique of sharpness maximisation. This allows an improvement in the accuracy of some phase estimation methods, as well as an increase in the understanding of how these techniques work on coherent imagery in general.

ACKNOWLEDGEMENTS

I would like to thank all those that have contributed to this thesis. The multi-disciplinary team nature of the Acoustics Research Group at the University of Canterbury has been especially conducive to good research. Many thanks to my two supervisors, Dr Michael Hayes and Professor Peter Gough. Their broad base of knowledge added essential perspective. Thanks to the Department of Electrical and Computer Engineering for putting up with me for so long and the lack of hassles I have experienced.

Work by my colleagues in the group on autofocus, speckle modelling, time-delay estimation, and interferometry all provided useful insight into my own work. Thank you Hayden Callow, Al Hunter, Ed Pilbrow, and Phil Barclay. Michael's and later Al's synthetic aperture sonar simulators were invaluable tools. Thanks to Phil and Hayden for always being ready to help out and answer my stupid questions and keep my computer running. Peter's and Michael's thorough correction and patient explanation of the basic rules of the English language have made this document significantly more readable.

It is telling that a rapid acceleration in the useful output from this research coincided with a closer working relationship with Michael. His input and exchange of ideas have lead to better understanding and insight than working in isolation. I would like to acknowledge Michael for the development of much of the mathematics leading towards direct sharpness maximisation in Appendix A. Thanks also to Assoc. Prof. Peter Smith for his input in the development of the second order statistics of sharpness in Section 3.2.1.

On a personal level, thanks to all my family, friends, flatmates, lovers, and climbing partners. Thanks for all your distraction and support and apologies for all the times I have been absent. A special thank you to Jo. Your love, support, and patience have meant a lot to me.

PREFACE

The work in this thesis began as an investigation of the use of sharpness maximisation to autofocus stripmap imagery. This followed the successful implementation in spotlight imagery [Gough and Lane, 1998]. Little previous work had been done on implementing sharpness maximisation on stripmap imagery.

It became apparent that both the sharpness and the phase estimate are random-variables due to speckle. Little had been published on the variation of the phase estimate. With a goal of quantifying the accuracy of sharpness maximisation, the statistics of the sharpness measure were developed. This work was published as a paper [Fortune et al., 2004] and forms the basis of Chapter 3.

Sharpness maximisation was found to be surprisingly unsuccessful on stripmap imagery. Much work was done to determine why, before it was found to be due to interpolation errors in the image reconstruction process. This invalidated much of the previous work. A fresh look was taken.

It was observed that if sharpness maximisation is performed on each range-bin separately, the variance of the result is highly dependent on the form of the signal in that range-bin. This made it apparent that some range-bins are useful for estimating phase error, some are not. By weighting each range-bin by the inverse variance of the phase, the best estimate could be made. The estimate variance was linked to the image phase variance, which was set by the signal to clutter ratio. Many methods of estimating the image phase variance were tried and discarded. An insight into the link to time-delay estimation lead to the successful use of coherence to measure image phase variance. This lead to the work presented in Chapter 4, resulting in improvements to both echo-correlation and sharpness maximisation.

There was a desire to reduce the ad-hoc nature of many developments in sharpness maximisation and autofocus. It was felt more important to gain a fundamental understanding of the technique, rather than present a new method or show improved results on a particular image. The work presented in Chapter 6 set out to answer the following questions about sharpness maximisation: Why does it work? What is the limit of performance? What is the optimal sharpness measure? And what is the optimal weighting to use? The limit of performance and best weighting were determined, but the question on what the *optimal* sharpness metric is remains to be answered.

This thesis is written in seven chapters. The first two define the terminology, geometry and processing steps used. A summary of the thesis contributions and an outline of the thesis organisation is found in Chapter 1. New work is presented in Chapters 2-6, with conclusions and recommendations for future work in Chapter 7.

Papers prepared during the course of work on this thesis are listed below in order of presentation. Work on speckle reduction in SAS has been published [Fortune *et al.*, 2003b], but is not included in this thesis. Work on the statistics of the sharpness measure has been published [Fortune *et al.*, 2003a; 2004] and is contained in Chapter 3. Compared to the paper, the second order speckle statistics are expanded on in this thesis. The remaining papers cover work now discussed in Chapter 5. At the time of submission, papers were in preparation for work discussed in Chapters 4 and 6.

S. A. Fortune, P. T. Gough, and M. P. Hayes. A statistical method for autofocus of synthetic aperture sonar images. In *IVCNZ2000, Image and Vision Computing New Zealand*, pages 56–61, Hamilton, New Zealand, November 2000.

P. T. Gough, M. P. Hayes, S. A. Fortune, and H. J. Callow. Autofocussing procedures for high-quality acoustic images generated by a synthetic aperture sonar. In *Proceedings of the 17^{th} International Congress on Acoustics*. ICA, September 2001.

S. A. Fortune, P. T. Gough, and M. P. Hayes. Statistical autofocus of synthetic aperture sonar images using image contrast optimisation. In *International Geoscience and Remote Sensing Symposium Proceedings, 2001*. IEEE, 2001a.

S. A. Fortune, M. P. Hayes, and P. T. Gough. Statistical autofocus of synthetic aperture sonar images using image contrast optimisation. In *Oceans 2001, Marine Technology and Ocean Science Conference*, volume 1, pages 163–169, November 2001b.

S. A. Fortune, M. P. Hayes, and P. T. Gough. Contrast optimisation of coherent images. In *IVCNZ2002, Image and Vision Computing New Zealand 2002*, pages 299–304. IVCNZ, November 2002.

S. A. Fortune, M. P. Hayes, and P. T. Gough. Contrast optimisation of coherent images. In *Oceans 2003, Marine Technology and Ocean Science Conference*, pages 2622–2628, 2003a.

S. A. Fortune, M. P. Hayes, and P. T. Gough. Speckle reduction of synthetic aperture sonar images. In *World Congress on Ultrasonics, Paris, September 2003*, pages 669–672, 2003b.

S. A. Fortune, M. P. Hayes, and P. T. Gough. Statistics of the contrast of coherent images. *Journal of the Optical Society of America A*, 21(7):1131–1139, July 2004.

DEFINITIONS

The notation used to define various quantities, transforms, operations and statistical properties are defined in this section.

Complex numbers

For a complex number

$$z = x + jy,$$

the conjugate is shown by

$$z^* = x - jy.$$

The real part is shown by

$$\text{Re}\,\{z\} = x,$$

the imaginary part

$$\text{Im}\,\{z\} = y,$$

the magnitude

$$|z| = \sqrt{x^2 + y^2},$$

and the phase

$$\angle\,\{z\} = \tan\frac{y}{x}.$$

Transforms

A Fourier transform is notated using,

$$X(f) = \mathcal{F}_{t \to f}\,\{x(t)\}$$
$$= \int_{-\infty}^{\infty} x(t)\exp\left(-j2\pi ft\right)\,\mathrm{d}t.$$

An inverse Fourier transform is given by

$$x(t) = \mathcal{F}_{f \to t}^{-1}\,\{X(f)\}$$
$$= \int_{-\infty}^{\infty} X(f)\exp\left(+j2\pi ft\right)\,\mathrm{d}f,$$

with the Fourier pair shown using

$$x(t) \longleftrightarrow_t X(f).$$

The following notation is used to define a two-dimensional Fourier transform

$$G(f_x, f_y) = \mathcal{F}_{\substack{x \to f_x \\ y \to f_y}} \{g(x, y)\}$$

$$= \iint\limits_{-\infty}^{\infty} g(x, y) \exp\left[-j2\pi(xf_x + yf_y)\right] \, dx \, dy$$

$$g(x, y) \longleftrightarrow_{x,y} G(f_x, f_y),$$

and marginal Fourier transform

$$G(x, f_y) = \mathcal{F}_{y \to f_y} \{g(x, y)\}$$

$$= \int_{-\infty}^{\infty} g(x, y) \exp\left(-j2\pi y f_y\right) \, dy$$

$$g(x, y) \longleftrightarrow_y G(x, f_y).$$

The notation for the convolution operator is

$$x(t) \odot y(t) = \int_{-\infty}^{\infty} x(\tau) y(t - \tau) \, d\tau = y(t) \odot x(t)$$

with correlation shown using

$$x(t) \star y(t) = x(t) \odot y^*(-t) = \int_{-\infty}^{\infty} x(\tau) y^*(\tau - t) \, d\tau.$$

Convolution and correlation have the corresponding Fourier transforms

$$\mathcal{F}_{t \to f} \{x(t) \odot y(t)\} = X(f) Y(f)$$
$$\mathcal{F}_{t \to f} \{x(t) \star y(t)\} = X(f) Y^*(f)$$
$$\mathcal{F}_{t \to f} \{x(t) \star x(t)\} = |X(f)|^2.$$

Statistics

Second order measures of random processes can be defined in different ways. In this work, statistical measures for two stationary complex valued random processes $X(t)$ and $Y(t)$ are defined as follows:

The ensemble (statistical) average:

$$\mathrm{E}\left[X\right] \equiv \int_{-\infty}^{\infty} X \Pr[X] \, dX.$$

The mean of a process:

$$\mu_X \equiv E[X].$$

The variance:

$$\sigma_X^2 \equiv E\left[(X - \mu_X)^2\right]$$
$$= E\left[X^2\right] - \mu_X{}^2.$$

The standard deviation:

$$\sigma_X \equiv \sqrt{\sigma_X^2}.$$

The cross-correlation of $X(t)$ and $Y(t)$:

$$R_{X,Y}(\tau) \equiv E\left[X(t+\tau)Y^*(t)\right].$$

The cross-covariance:

$$C_{X,Y}(\tau) \equiv E\left[(X(t+\tau) - \mu_X)(Y(t) - \mu_Y)^*\right]$$
$$= R_{X,Y} - \mu_X\mu_Y{}^*.$$

The correlation coefficient, a normalised version of covariance:

$$\rho_{X,Y}(\tau) \equiv \frac{C_{X,Y}(\tau)}{\sigma_X\sigma_Y}.$$

The autocorrelation of $X(t)$:

$$R_X(\tau) \equiv E\left[X(t+\tau)X^*(t)\right].$$

The power spectral density (PSD) is the Fourier transform of a signal's autocorrelation, defined as:

$$S_X(f) \equiv \int_{-\infty}^{\infty} R_X(\tau)\exp\left(-j2\pi f\tau\right)\,d\tau.$$

Functions

The following functions are defined:

The rectangular function:

$$\text{rect}(t) = \begin{cases} 1 & |t| < 1/2, \\ 0 & |t| \geq 1/2. \end{cases}$$

The triangular function:

$$\text{tri}\,(t) = \begin{cases} 0 & t \geq 1, \\ 1 - |t| & -1 < t < 1, \\ 0 & t \leq -1. \end{cases}$$

The sinc function:

$$\text{sinc}\,(t) = \frac{\sin(\pi t)}{\pi t}.$$

CONTENTS

Chapter 1

INTRODUCTION

This chapter introduces the imaging geometry and terminology used throughout this thesis. The imaging medium (i.e., radar, sonar, or other) is not an important consideration, although the image formation process (i.e., spotlight or stripmap) is. This chapter places both imaging methods in a general framework, while the following chapter (Chapter 2) looks specifically at the processing used for image formation.

1.1 SIDE-SCAN IMAGING

Fig. 1.1 shows the general three-dimensional (3-D) imaging geometry of a typical side-scan (or side-looking) imaging system. The platform (aircraft or spacecraft for synthetic aperture radar (SAR) systems, towfish for synthetic aperture sonar (SAS) systems) travels along the u-axis. This is commonly referred to as either the along-track, azimuth, or cross-range direction. The dimension of data collection u, is the pulse number, or slow-time co-ordinate. The image plane (x, y) is offset from the data collection as shown. There is a direct correspondence between the pulse number u and the imaging spatial co-ordinate y. The terminology used to describe the separate domains can be ambiguous. This thesis uses *pulse-number* or *ping* to refer to the data collection domain u and *along-track* to refer to the imaging domain y. Likewise, there is a distinction between the delay co-ordinate t corresponding to the polar *range* co-ordinate and the Cartesian *across-track* co-ordinate x. For a usefull approximation, imaging can be considered in the slant-range plane. By considering the height z as zero, the geometry can be considered in two dimensions. This is shown in Figures 1.2 and 1.3.

1.1.1 Range resolution

In radar and sonar imaging systems, range is measured using the time of flight between transmission and reception of pulses. High range resolution is obtained by transmitting dispersed large time-bandwidth pulsed signals that are compressed on reception using standard techniques known as matched filtering or pulse compression [Rihaczek, 1969].
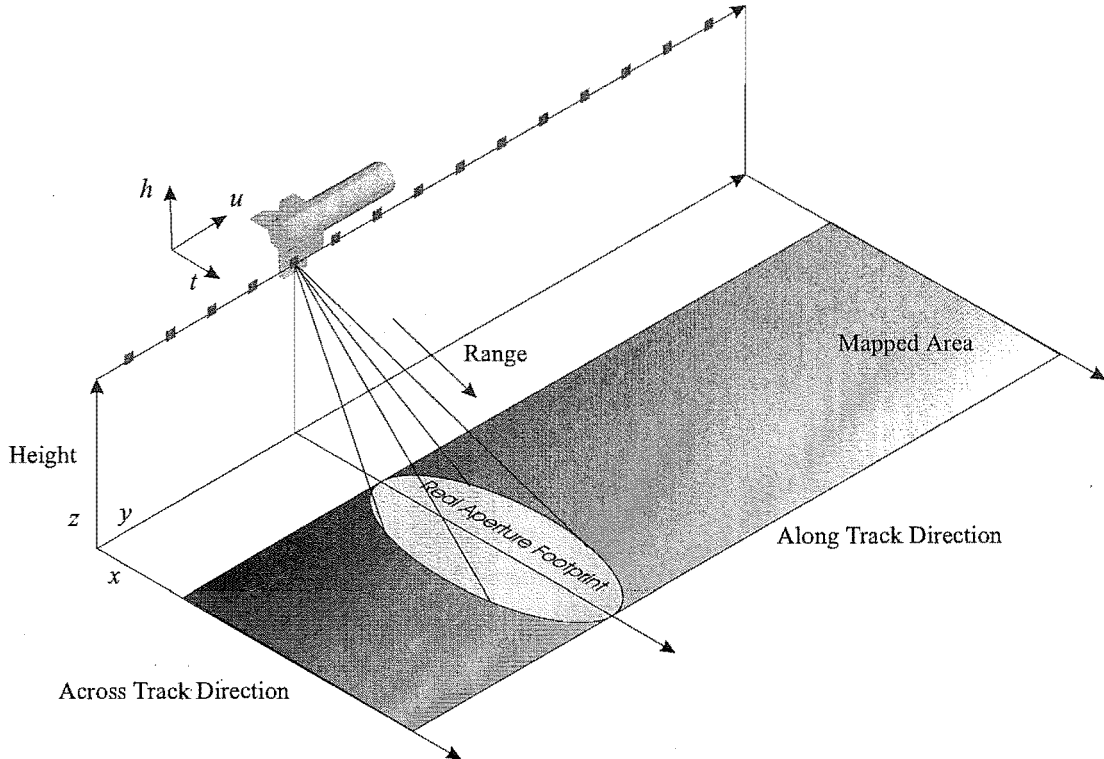
Figure 1.1 Basic imaging geometry for side-looking system.

In this thesis, the data is assumed to be in its pulse-compressed form. The achievable range or across-track resolution is typically referred to as [Hawkins, 1996]

$$\delta x_{3dB} = \alpha_w \frac{c}{2B_c},$$ (1.1)

for a rectangular spectrum of bandwidth B_c and wave propagation speed c. Resolution is defined as the 3 dB width of the range-compressed pulse. α_w is a constant reflecting the effect of any windowing performed to reduce range side-lobes. $\alpha_w = 0.88$ for no (rectangular) windowing and $\alpha_w = 1.30$ for Hamming weighting.

1.1.2 Real aperture

In real aperture imaging, a 2D measure of the *magnitude* of the reflection from an area is built from the returns of a number of pulses. The ground reflectivity in the across-track direction is measured from the magnitude of the return from a single pulse at different delays. The ground reflectivity is sampled in along-track by moving the platform perpendicular to the direction of the pulse and transmitting another pulse. This process builds up a 2D image or *strip-map* of the ground reflectivity.

High along-track resolution is obtained by imaging with a narrow beam-width. This requires a transducer with a large aperture and high-frequencies. The constant angular response of the radiation pattern means that real aperture images have a range

variant along-track resolution given by

$$\delta y_{3dB} = \alpha_w \frac{cx}{f_0 D},$$

(1.2)

where D is the length of the aperture and f_0 is the centre frequency of the signal.

1.1.3   Synthetic aperture

A synthetic aperture imaging system synthesizes the effect of a large aperture by moving a small physical aperture in space. The received echoes are stored, then coherently integrated (summed) in an appropriate manner to produce an image with increased along-track resolution. A limiting factor of synthetic aperture techniques is that the echos must retain phase coherency for the length of the the synthetic aperture. Without echo coherency the reconstructed image is severely degraded. The motion of the platform must be known to a high level of accuracy. The length of the synthetic aperture formed is limited by the beam-width of the transducer. A smaller transducer (or larger beam-width) leads to improved along-track resolution, the inverse of real aperture imaging. The along track resolution of a synthetic apeture image is given by [Curlander and McDonough, 1996; Hawkins, 1996]

$$\delta y_{3dB} = \frac{D}{2}.$$

(1.3)

The data forming a synthetic aperture can be collected using several methods. The two main modes of synthetic aperture imaging discussed in this thesis are *strip-map* and *spotlight*. Other techniques include inverse synthetic aperture radar (ISAR) [Wehner, 1987], consisting of a stationary receiver and a moving target.

Spotlight mode

In spotlight imaging, the real aperture is continuously steered so that it always illuminates the same ground patch as shown in Fig. 1.2. Spotlight allows a simplified image reconstruction process and reduced sampling requirements compared to stripmap imagery. In spotlight imagery there is a simple scaling between the pulse-number u and the along-track spatial frequency co-ordinate k_y or v. The term *azimuth* is used in SAR literature to mean either the along track spatial co-ordinate in the image domain, the pulse number or along track distance in the data domain, or the angle from broadside. This thesis uses azimuth to refer to the along-track spatial frequency co-ordinate v, which relates to both the pulse number and angle from broadside as shown in the diagram.

The measurement system produces *polar* samples of the Fourier transform of the imaged scene. This data is reformatted onto a Cartesian grid. For simplicity and
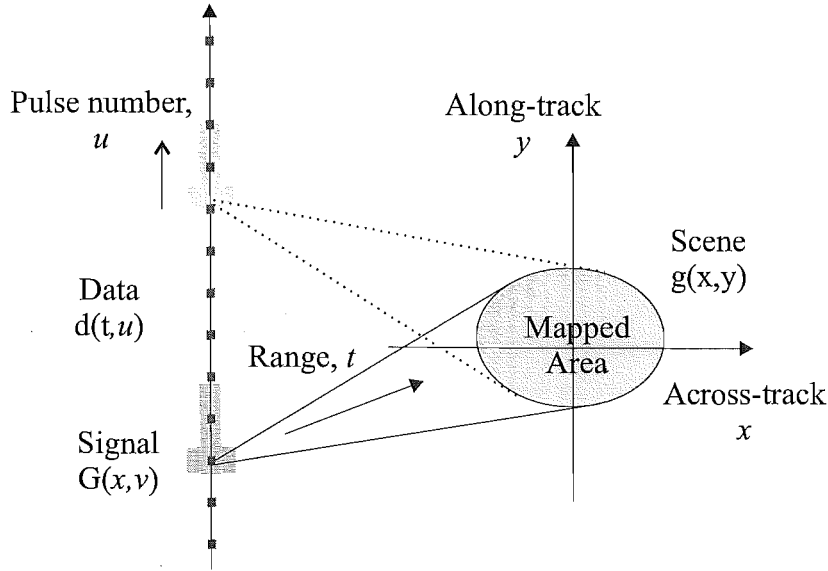
Figure 1.2   Spotlight ground-plane geometry.

reduced processing, this thesis considers data processed into the range-azimuth domain, also known as the range-compressed phase history, referred to as the *signal* $G(x, v)$.

Stripmap mode

In stripmap imaging, the beam always points in the same direction relative to the platform. Fig. 1.3 shows the beam perpendicular to the platform's direction of travel (broadside mode). It is possible to angle the beam off broadside (squint mode). Stripmap is the conventional form of imaging for SAS systems. The image reconstruction process creates the image $g(x, y)$ from the data $d(t, u)$. Image reconstruction is more difficult for stripmap systems as any blurring is point variant.

1.1.4   Multiple receiver system

The slow propagation speed of sound places restrictive limits on the mapping area, or platform speed, for SAS. The mapping area can be increased by using an array of multiple receivers. This is the common system in use today. At the time of writing this thesis, data from a multiple receiver SAS was not available. Thus the algorithms developed have been tested on single receiver systems. Multiple receiver systems are not commonly used in SAR.

1.1.5   Platform motion

Synthetic aperture image reconstruction requires echo phase coherency for the length of the synthetic aperture. Movement of the imaging platform off a straight path reduces
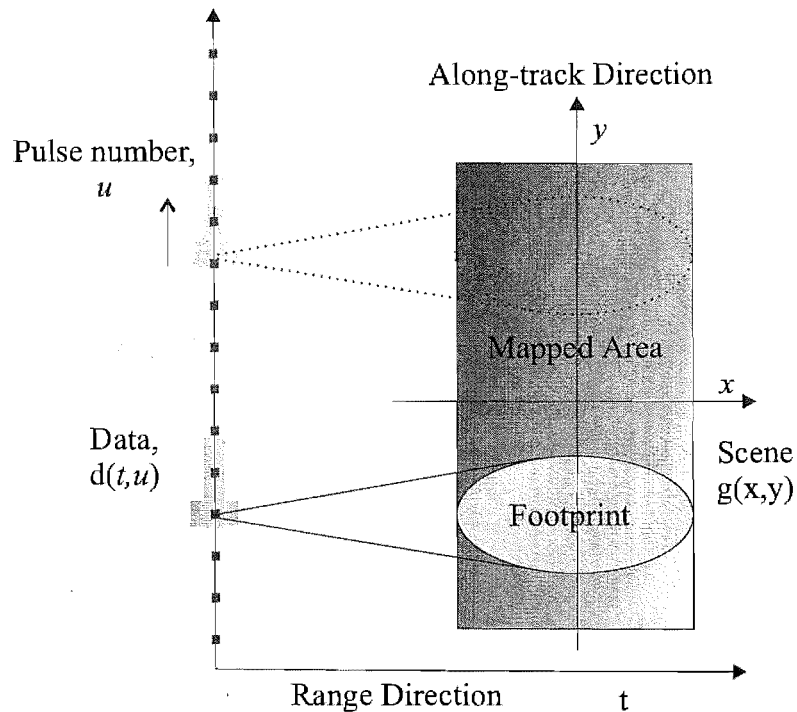
Figure 1.3 Stripmap ground-plane geometry.

phase coherency and degrades the reconstructed imagery. Fig. 1.4 shows the six possible platform motions. For a wide-beam, single receiver SAS, a small rotation does not degrade the image so long as the target area is still contained within the beam-width. Of the translational motions, any motion perpendicular to the slant-range direction does not affect the phase of the returned echo. With the low-grazing angles and long standoff ranges typical for side-scan imagery, sway is the motion that most affects the phase of the recorded echos, causing image degradation. Thus the sway needs to be accurately estimated for quality imagery [Johnson *et al.*, 1995]. Most autofocus methods estimate a phase error, corresponding to the average platform motion in the slant-range direction. Generally, it is not possible to deduce a full 3D platform motion estimate from this measurement. For a multiple hydrophone SAS, a varying yaw effects the phase of each receiver differently. Thus yaw also needs to be accurately estimated [Douglas and Lee, 1993; Christoff, 1998; Gough and Miller, 2004].

Motion compensation

Motion compensation is the process of correcting data for a given estimate of the motion of the platform that collected the data. How the motion is estimated is not important. This thesis uses standard methods of motion compensation, discussed in Chapter 2. Some authors use the term motion compensation to refer to a process of motion estimation, which can be misleading.
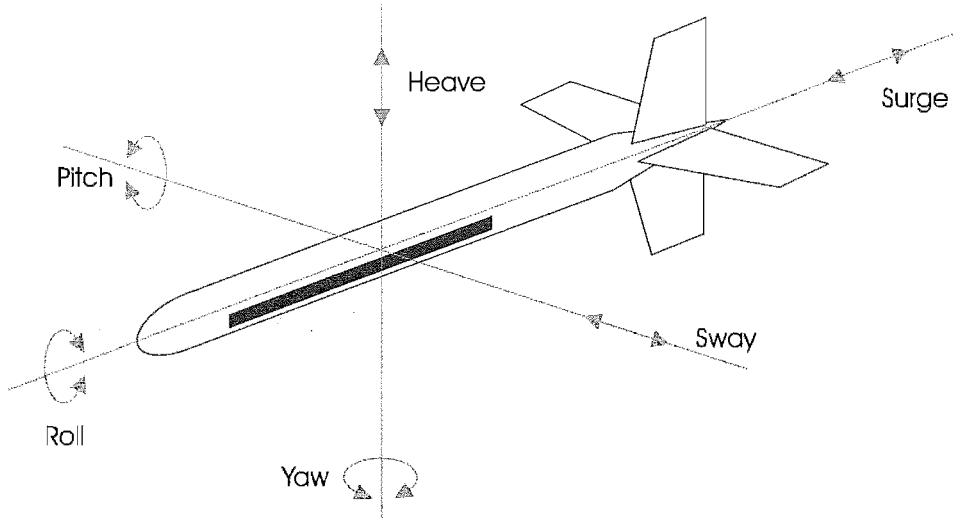
Figure 1.4   Definition of platform motions.

Motion estimation

As its name suggests, motion estimation is the estimation of the 3D motion of the imaging platform. Motion estimation could refer to the estimation of motion using instrumentation as well as using the collected data or imagery. The result of motion estimation can be used for both correcting the imagery and navigation of the imaging platform.

Autofocus

Autofocus is the estimation and correction of platform motion to improve the reconstructed imagery, using only the collected data. Some authors have made a distinction of whether autofocus is performed in the data or image domains. However, this thesis shows that one autofocus method (sharpness maximisation) can be performed in either domain. Thus the distinction between the two domain is not important, and autofocus may be performed in either the image or the data domain.

Motion estimation is a more general term that may include data from instrumentation. Autofocus may be used as one method of motion estimation. However, on its own, autofocus may not give an accurate estimation of motion of the platform. Often autofocus does not estimate any linear component of sway, as this does not degrade the imagery. This may limit the application of autofocus for navigational purposes.

1.2   DATA SETS

There are several different data sets used throughout this thesis. This includes field collected towed SAS imagery, simulated stripmap SAS imagery, and simulated spotlight

test images. At the time of writing this thesis, access to multiple-hydrophone SAS data, or real SAR data was not available.

Field SAS images have been collected using KiwiSAS-III, a free-towed SAS developed by the Acoustics Research Group at the University of Canterbury. The operating parameters of KiwiSAS-III are shown in Table 1.1. In addition, simulated SAS images were used. These were produced using a frequency domain facet based SAS simulator developed by Alan Hunter [Hunter, 2005]. The simulated images used the same parameters as KiwiSAS, shown in Table 1.1.

Parameter	Symbol	Value
Bandwidth	B_c	20 kHz
Centre frequency	f_0	30 kHz
Sampling frequency	f_s	30.03 kHz
Transmitter length (horizontal)	D_t	0.225 m
Receiver length (horizontal)	D_r	0.325 m
Platform velocity (approximate)	v	1.4 ms^{-1}
Wave propagation speed	c	1500 ms^{-1}
Pulse period	T_c	0.0125 s
Pulse repetition frequency	prf	14.6 Hz
Standoff range	r_0	25.57 m

Table 1.1 Operating parameters of KiwiSAS-III.

The images used to test spotlight autofocus come from a variety of sources. Some images are field or simulated stripmap SAS images that have been reconstructed to form an image, which is then considered a spotlight image. This should give negative exponential intensity speckle statistics, with small amounts of correlation of the speckle. One image used (the ship) is a reconstructed field SAS image courtesy of the Norwegian Defense Research Establishment (FFI) [Hansen et al., September 2003]. Another source is simple direct simulation in the image domain. A simple geometric object, for example a block or shadow, is multiplied by a speckle image, consisting of circular-Gaussian uncorrelated noise. Point targets are modelled as additive. To simulate a band-limited image, the image is filtered in the image domain by a Gaussian point-spread-function.

1.3   ASSUMED BACKGROUND

The reader of this thesis should not need a detailed understanding of synthetic aperture processing techniques. The concepts of motion compensation and image reconstruction as well as a model of the effect of motion on synthetic aperture imaging systems is used throughout the thesis. Standard methods of motion compensation and image reconstruction are used, described further in Chapter 2.

A knowledge of signal processing techniques, sampling theory, sampling of complex baseband signals, Fourier theory and stochastic processes is required. This thesis builds on the standard autofocus methods encompassing time-delay estimation, shear-average (echo-correlation), phase gradient autofocus (PGA), and sharpness maximisation. Knowledge of these methods would be useful.

1.4   THESIS CONTRIBUTIONS

This thesis contains major contributions to the understanding and to the practice of estimating the phase error in synthetic aperture imagery, both stripmap and spotlight. The technique of sharpness maximisation is focused upon, with a number of improvements to the method suggested. The method of phase estimation used in echo-correlation is also improved upon. Standard motion compensation and image reconstruction algorithms are used throughout the thesis.

In Chapter 4, it is shown that echo-correlation can be framed as a time-delay estimation problem. If the signal to noise ratio changes with frequency, a generalised correlation method can be used to improve the accuracy of time delay estimation. The same technique can be used for estimating the average phase difference between echos. This leads to a more accurate phase difference estimation if the signal to clutter ratio varies with range. Chapter 4 presents a model of the variance of the image phase. This is determined by the signal to clutter ratio, which can be measured using the coherence between echos. A weighted phase difference estimator (WPDE) is developed, which weights the phase estimate at each range bin by an estimate of its inverse variance. An increase in the accuracy of the phase estimate is shown in many image types.

Strong targets cause a phase bias. Centre-shifting spotlight images can remove the phase difference offset caused by the strongest scatterer in each range-bin. An improved centre-shifting method is developed, which directly removes the mean phase-difference offset from each range-bin. In stripmap imagery, using an amplitude weighting or estimating the phase curvature removes this bias. A weighted phase curvature estimation technique is presented, which shows improved performance over weighted shear average for images of high SCR.

The alternate autofocus method of sharpness maximisation is examined, both practically and theoretically. In Chapter 3, the statistics of the sharpness of a coherent image are developed from the statistics of speckle. The second-order statistics of speckle in a SAS image are modelled and compared to a measured SAS speckle image. A model of the variance of the image sharpness is presented as a function of speckle intensity, speckle correlation length, image size and image sharpness. The sharpness variance is shown to be inversely proportional to the number of independent speckles contained in the image.

A number of improvements to sharpness maximisation are presented in Chapter 5. Sharpness maximisation is presented in a Bayesian framework, using a measure of the path estimate probability to regularise the result. This is shown to reduce oversharpening and increase the accuracy of the path estimate in many images. An optimal (Karhunen-Loeve) decomposition of the path is used for a parametric optimisation. A comparison of different optimisation methods shows an iterated cascaded parametric optimisation is a good choice for minimising the cost function. Results are presented comparing the performance of sharpness maximisation on a range of images. Different sharpness metrics are compared, the effect of different path-weight costs is shown and different range-weightings are compared.

A limit on the minimum variance of the phase estimate from maximising the standard intensity squared sharpness measure is developed. The variance of the phase giving maximum sharpness is shown to be determined by the variance of the image phase, which is set by the signal to clutter ratio. A model of the phase variance using the mean echo coherence is presented and compared to the measured phase variance for a number of images. This is used to form an optimal range weighting and to predict the phase estimate variance for each echo. The performance of sharpness maximisation is compared to the theoretical limit and to the phase estimate from echo-correlation.

A method for the direct calculation of the phase estimate maximising sharpness is developed in Chapter 6. This method may be considered either a high-order echo-correlation method, or an extension to the method of conjugate gradient optimisation. It converges to the same result as a normal sharpness maximisation, without requiring any optimisation. A general sharpness measure may be used. The performance of direct sharpness maximisation is compared on a range of images using different metrics. A negative entropy sharpness measure is shown to perform better than all other phase estimation techniques tested on an extended block target.

1.5 THESIS OUTLINE

Chapter 2 introduces the methods of motion compensation and reconstruction for spotlight and stripmap synthetic aperture imagery used throughout the thesis.

Chapter 3 presents a development of the statistics of a synthetic aperture image. The statistics of speckle, including second order statistics, are developed and compared to measured statistics from field SAS images. The statistics of image sharpness are also developed. A model for the distribution of image sharpness is compared to simulated and field SAS speckle images.

Chapter 4 presents a review and analysis of the phase error estimation method of echo correlation. The method of generalised correlation for time-delay estimation is reviewed. Different methods of complex cross-correlation are compared. Echo

correlation is shown to be a narrow-band equivalent of time-delay estimation. A weighted phase-difference estimation technique, the equivalent of generalised correlation, is presented. Results are shown for a range of image types. Echo correlation of stripmap images is discussed and a new technique developed.

Chapter 5 shows an investigation of the phase error estimation method of sharpness maximisation. A number of practical issues with the method are discussed and extensions presented. This includes reviewing different methods of measuring the image sharpness, measuring the probability of the estimated path, parameterising the platform path and different optimisation methods. Results comparing different metrics and path cost weightings for a range of images are presented.

Chapter 6 presents a theoretical analysis of sharpness maximisation. The nature of the phase estimate that maximises sharpness is analysed. An improved range-weighting method is presented. A direct method of calculating the phase giving maximum sharpness without optimisation is developed. A performance limit to sharpness maximisation is presented and results from a number of images are compared to this limit.

Chapter 7 contains the conclusions drawn from this work and provides recommendations for future research.

Chapter 2

IMAGING TECHNIQUES

This chapter introduces the techniques of spotlight and stripmap synthetic aperture radar (SAR) and synthetic aperture sonar (SAS) imaging that are used throughout this thesis. This is not an exhaustive treatment on the subject. This chapter outlines the techniques required for autofocus, namely motion compensation and reconstruction of both spotlight and stripmap imagery. For more information, readers can refer to sources on SAR processing [Curlander and McDonough, 1996; Carrera et al., 1995; Jakowatz et al., 1996; Soumekh, 1994; 1999; Henderson and Lewis, 1998], or SAS processing [Hawkins, 1996; Callow, 2003; Gough and Hawkins, 1997; 1998].

2.1 EFFECT OF PLATFORM MOTION

Consider a sway, or sideways motion of the imaging platform from the mean path, of $w(u)$ as shown in Fig. 2.1. Assume the transmitter and receiver are co-located (the phase-centres approximation). Assuming the amplitude variation is negligible and that $|w(u)| \ll r(u)$ where $r(u)$ is the range to target, then the effect of sway can be seen as a timing error [Soumekh, 1994; Callow, 2003], giving

$$d(t,u) \approx \overline{d}(t - \frac{2}{c}w(u)\cos\theta, u), \qquad (2.1)$$

where $\overline{d}(t,u)$ is the aberration free echo and θ is the angle to target. Wide-beam motion compensation techniques must account for the dependency of the timing error on the angle to target θ [Callow et al., 2004].

2.1.1 Narrow beam approximation

A simplification may be made by assuming narrow beam-widths (small θ), giving

$$d(t,u) \approx \overline{d}(t - \frac{2}{c}w(u), u). \qquad (2.2)$$
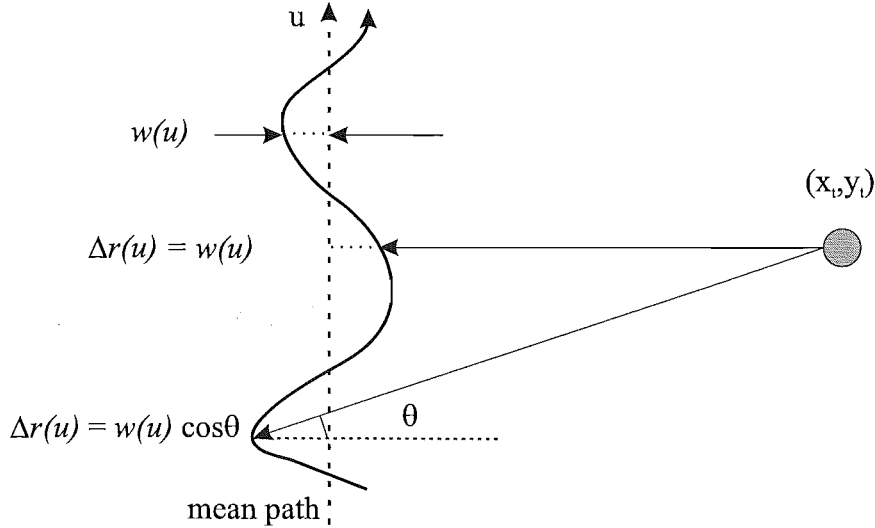
Figure 2.1 Effect of sway on path length. A sway of $w(u)$ affects the path length $r(u)$ by $\Delta r(u)$.

The comparatively narrow beam-widths used in SAR mean this assumption is used widely [Jakowatz et al., 1996, Sec 4.2]. SAS has wider beam-widths, however the narrow-beam approximation is still commonly used [Hawkins, 1996; Bellettini and Pinto, 2002]. For KiwiSAS, a beam-width of approximately 30 degrees introduces errors in the path-length of less than 3.5%. Narrow-beam motion compensation will be used throught this thesis.

2.2 SPOTLIGHT IMAGING

Consider spotlight data that has been range-compressed, polar formatted and range-aligned, forming the *phase history* signal $G(x, v)$.

2.2.1 Motion compensation

A spotlight SAR typically is a high-Q system with small beam-widths and long standoff-ranges. This allows a direct mapping between the pulse-number u and the along-track spatial-frequency v [Jakowatz et al., 1996, Sec 3.4]. This allows sway to be modelled as a phase error independent of range, giving

$$G(x, v) = \overline{G}(x, v) \exp\left[j\phi_e(v)\right], \qquad (2.3)$$

where $\phi_e(v) = \frac{4\pi f_0}{c} w(u = -v \frac{c r_0}{4\pi f_0})$, f_0 is the centre frequency and r_0 is the reference range to the centre of the patch. Although more sophisticated models are possible, this is the common starting point for most autofocus methods [Jakowatz and Wahl, 1993; Wahl et al., 1994a; Ye et al., 1999; Fienup, 2000]. To be independent of operating

parameters, the results shown in this thesis are for the estimation of the phase error $\phi_e(v)$ in radians, not the sway $w(u)$ in meters.

To correct a signal for an estimated phase error $\widehat{\phi}(v)$, a negative phase correction can be applied using

$$\widehat{G}(x,v) = G(x,v) \exp\left[-j\widehat{\phi}(v)\right].$$ (2.4)

2.2.2 Image reconstruction

The image estimate can be formed with an inverse Fourier transform of the signal in the azimuth direction, i.e.,

$$\widehat{g}(x,y) = \mathcal{F}_{v\to y}^{-1}\left\{\widehat{G}(x,v)\right\}.$$ (2.5)

Because the phase error is in the azimuth direction, the aberrated image is smeared only in the along-track direction.

2.3 STRIPMAP IMAGING

Due to the wider bandwidth (lower-Q) systems generally employed in stripmap SAS, a different motion compensation method is used. Image reconstruction is a more difficult problem for stripmap imagery compared to spotlight imagery.

2.3.1 Motion compensation

Consider a narrow-beam timing error given by (2.2) as

$$d(t,u) \approx \overline{d}(t - \frac{2}{c}w(u), u).$$ (2.6)

The received echo is down-sampled to baseband, giving

$$\widetilde{d}(t,u) \approx \widetilde{\overline{d}}(t - \frac{2}{c}w(u), u) \exp\left[-j\frac{4\pi f_0}{c}w(u)\right],$$ (2.7)

where f_0 is the centre frequency. The first term is the time-shifted complex envelope, and the second term is caused by the phase shift of the carrier. Taking the Fourier transform of (2.7) in the fast-time (t) variable gives

$$\widetilde{D}(f,u) \approx \widetilde{\overline{D}}(f,u) \exp\left[j\frac{4\pi(f+f_0)}{c}w(u)\right].$$ (2.8)

Thus sway introduces a phase change proportional to the total frequency, not just the centre frequency.

To correct data for a given estimate of sway $\widetilde{w}(u)$, first the data requires a Fourier transform in the fast-time (t) direction giving $\widetilde{D}(f,u)$. A phase correction can then be

applied using

$$\widehat{D}(f,u) \approx \widetilde{D}(f,u) \exp\left[-j\frac{4\pi(f+f_0)}{c}\widehat{w}(u)\right].$$ (2.9)

The corrected data is then obtained with an inverse Fourier transform,

$$\widehat{d}(t,u) = \mathcal{F}_{f\to t}^{-1}\left\{\widehat{D}(f,u)\right\}.$$ (2.10)

2.3.2 Narrowband approximation

For a high-Q system, a narrow-band approximation can be used giving the further simplification;

$$\widehat{D}(f,u) \approx \widetilde{D}(f,u) \exp\left[-j\frac{4\pi f_0}{c}\widehat{w}(u)\right].$$ (2.11)

As in the spotlight case, this has a constant phase error with range with the phase correction $\phi(u) = \frac{4\pi f_0}{c}\widehat{w}(u)$. A phase correction can then be made directly to the data with

$$\widehat{d}(t,u) \approx \widetilde{d}(t,u) \exp\left[-j\frac{4\pi f_0}{c}\widehat{w}(u)\right].$$ (2.12)

2.3.3 Image reconstruction

A number of different synthetic aperture stripmap image inversion, or image reconstruction, methods have been developed for the estimation of the image reflectivity $g(x,y)$ from the data $d(t,u)$. Originally, SAS systems used a time-delay and sum technique, known as time domain correlation or exact matched-filtering [Gough and Hawkins, 1997]. Although exact, this method has a large computational load. The recently developed algorithm, fast factorised back projection, is an efficient time-domain inversion algorithm [Ulander et al., 2001]. Spatial frequency domain methods such as the range-Doppler [Bamler, 1992; Carrera et al., 1995; Curlander and McDonough, 1996], chirp-scaling [Cumming et al., 1992; Runge and Bamler, 1992; Raney et al., 1994; Hawkins, 1996], and the wavenumber [Stolt, 1978; Cafforio et al., 1991; Soumekh, 1994; Callow, 2003] algorithms provide a more efficient inversion, but require assumptions about the collection geometry. The results shown in this thesis have used a wavenumber inversion algorithm.

Chapter 3

SYNTHETIC APERTURE IMAGE STATISTICS

Due to the coherent nature of synthetic aperture (SA) imagery, the measured echo from a rough surface is not deterministic. Interference between randomly positioned scatterers causes the measured echo to be a random value, where the mean is the desired backscatter coefficient. Any function of the measurements, such as sharpness, is also a random variable. It is worthwhile therefore to define and understand the statistics of SA images. The statistics of a SA image are developed in Section 3.1, based on the well known statistical properties of speckle. The statistics of sharpness are developed in Section 3.2. Both are compared to statistics of simulated and field SAS images from KiwiSAS.

3.1 STATISTICS OF SPECKLE

If the resolution size of a SA image is large with respect to the system wavelength, there are multiple independent scatterers in each resolution cell. When these scatterers are illuminated by a coherent source, the reflections constructively and deconstructively interfere. This interference causes the measured echo from each element to be random giving the image a granular appearance. This granular appearance is known as speckle. An example of a SAS image of a flat patch of seafloor is shown in Fig. 3.1.

A commonly assumed model for speckle is a multiplicative random noise process with circular Gaussian statistics. This is known as *fully developed speckle*. Conditions for which include: [Goodman, 1986]

- There are a large number of scatterers contributing to the signal in each resolution cell.

- The contributing scatterers are independent.

- No scatterer is so strong that it dominates the reflected signal.

- The phase of each scatterer is random, i.e., uniformly distributed over $[0, 2\pi]$.
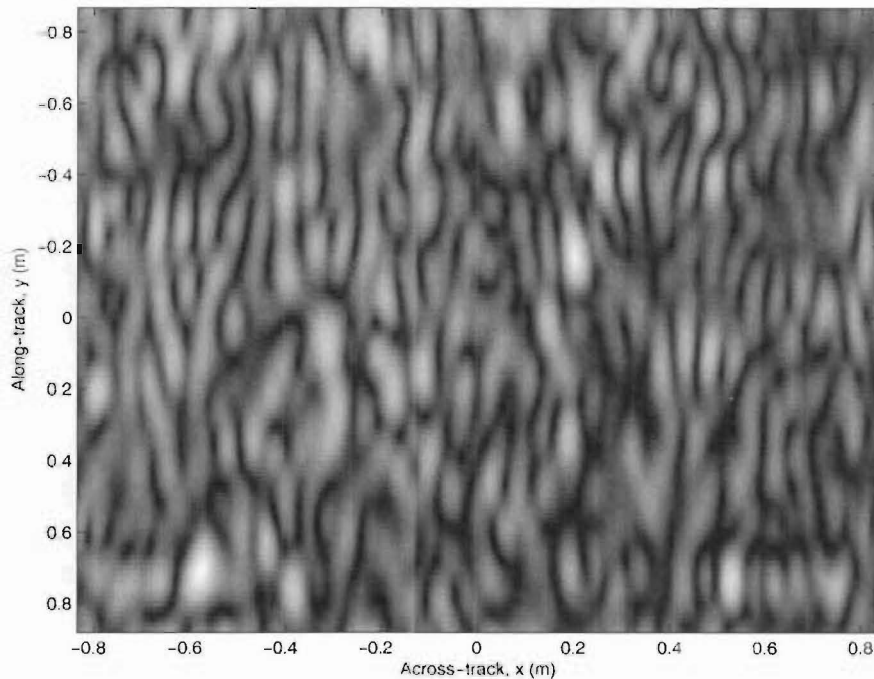
Figure 3.1   A SAS intensity image of a flat patch of seafloor, showing speckle. Across-track resolution is 0.05m, along-track resolution is 0.20m. Pixels have been interpolated to see fine structure.

When the resolution of a real aperture system becomes small compared with the size of a wavelength, the above assumptions are violated and the statistics depart from the classical behaviour. This is difficult to model for synthetic aperture sonar, so the resulting statistics are uncertain. As new generation, higher resolution SAS's come into use, it will become more apparent if this is an issue. Analysis of KiwiSAS image statistics (for example Fig. 3.2) shows the assumptions of fully developed speckle are accurate.

Speckle arises in many forms of imaging, mostly from observing a fairly coherent source reflected from a rough surface or propagating through a medium with random refractive index fluctuations. For example, laser, ultrasound, or radar imaging all show speckle. Similarly, a short exposure image of an astronomical object taken through a turbulent medium (such as earth's atmosphere) also has speckle-like structure. Since speckle has been studied for a number of years, the statistics of speckle are well understood. Good references on the subject include [Dainty, 1975; Goodman, 1976; 1986; George and Sinclair, 1976]. This section develops the statistics of speckle for SA imaging, focusing on SAS imagery. This is compared to measured statistics of field SAS images.

3.1.1 First-order statistics

Speckle can be modeled by a random walk in the complex plane where each step in the walk is the echo received from a single scatterer from within a resolution cell. These echos are coherently summed at the receiver and the resultant vector recorded. The value of the resultant field is a random variable and needs to be described by stochastic methods. From the assumptions for fully developed speckle, it follows using the central-limit-theorem that the probability density function of the real (X) and imaginary (Y) parts of the field are uncorrelated and are zero mean, σ_X^2 variance Gaussian random variables [Goodman, 1975];

$$\Pr(X, Y) = \frac{1}{2\pi\sigma_X^2} \exp\left(-\frac{X^2 + Y^2}{2\sigma_X^2}\right). \tag{3.1}$$

Thus the magnitude $M = \sqrt{X^2 + Y^2}$ follows the Rayleigh distribution;

$$\Pr(M) = \begin{cases} \frac{M}{\sigma_X^2} \exp\left(-\frac{M^2}{2\sigma_X^2}\right) & M \geq 0 \\ 0 & \text{otherwise,} \end{cases} \tag{3.2}$$

and the speckle intensity $I = X^2 + Y^2$ has a negative exponential distribution given by

$$\Pr(I) = \begin{cases} \frac{1}{\sigma_I} \exp\left(-\frac{I}{\sigma_I}\right) & I \geq 0 \\ 0 & \text{otherwise,} \end{cases} \tag{3.3}$$

where $\sigma_I = 2\sigma_X^2$. The n^{th} moment of intensity can be shown to be

$$\mathrm{E}\left[I^n\right] = \begin{cases} n!\,\mathrm{E}\left[I\right]^n & n \text{ positive integer} \\ \Gamma(n+1)\,\mathrm{E}\left[I\right]^n & n > -1. \end{cases} \tag{3.4}$$

Thus the mean and variance of the intensity are

$$\mathrm{E}\left[I\right] = \sigma_I, \tag{3.5}$$

$$\begin{aligned} \mathrm{Var}\left[I\right] &= \mathrm{E}\left[I^2\right] - \mathrm{E}\left[I\right]^2 \\ &= \sigma_I^2. \end{aligned} \tag{3.6}$$

Since X and Y are uncorrelated, the speckle phase θ is independent of the intensity and has a uniform distribution given by

$$\Pr(\theta) = \begin{cases} \frac{1}{2\pi} & -\pi \leq \theta < \pi \\ 0 & \text{otherwise.} \end{cases} \tag{3.7}$$

This distribution applies to an ensemble of the intensities of a single point over
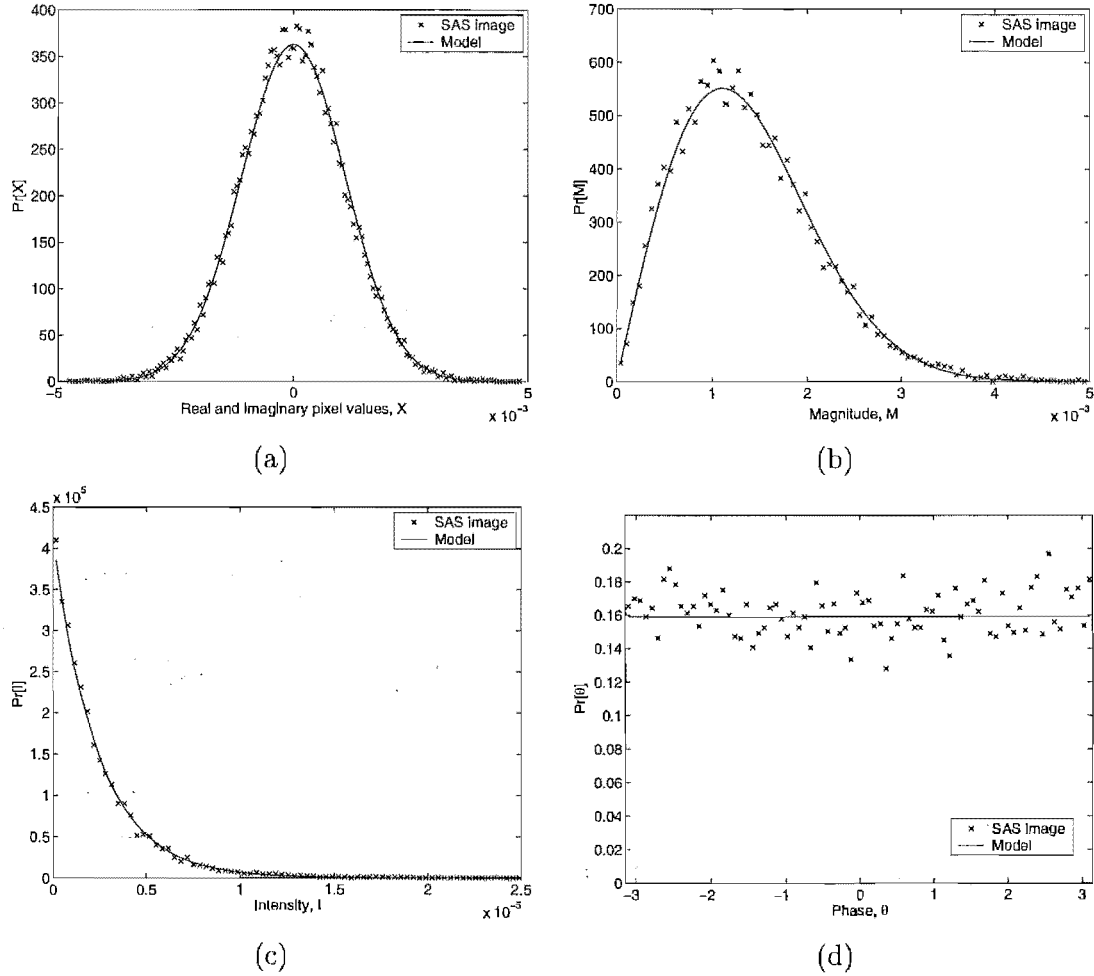
Figure 3.2 Measured probability distribution function of pixels over a patch of seafloor in a field SAS image. (a) Pixel real and imaginary components, compared to model (3.1). (b) Pixel magnitude, compared to model (3.2). (c) Pixel intensity, compared to model (3.3). (d) Pixel phase, compared to model (3.7).

different speckle realisations. It is tempting to apply it over a number of pixels in the same image. This can be done if the statistics of each point are equal, which in general is not true (the process is not ergodic). However, it can be approximately applied over a number of pixels of a single speckle realisation if each pixel has the same expected value and the patch contains many independent 'speckles'. For example, Fig. 3.2 shows the probability distribution of pixels on a small patch of a bland seafloor field SAS image. This shows the expected Gaussian distribution of the real part, Rayleigh distribution for magnitude, negative-exponential distribution for intensity, and uniform distribution for phase of fully developed speckle.

A common measure of the level of speckle is the *speckle contrast* defined as the ratio of the standard-deviation to the mean of the intensity of an area of image;

$$C \equiv \frac{\mathrm{Std}\,[I]}{\mathrm{E}\,[I]}. \tag{3.8}$$
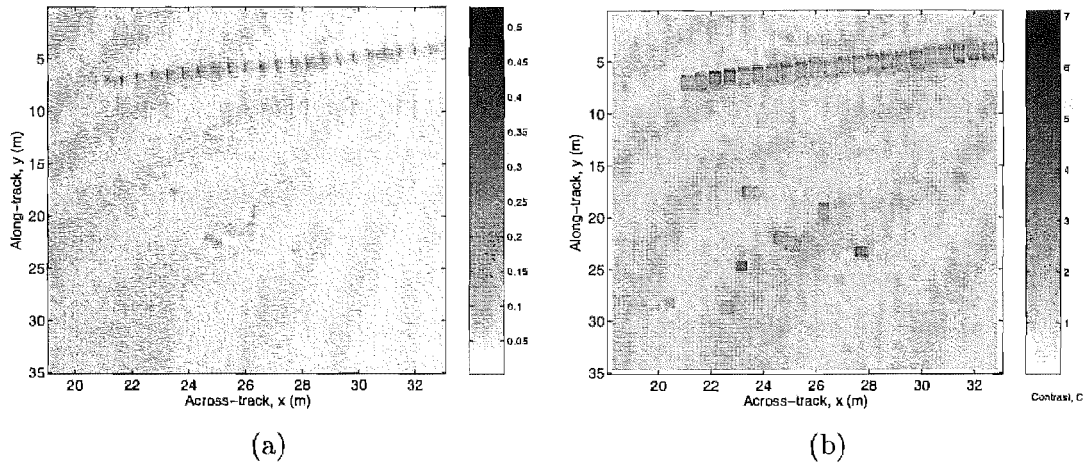
Figure 3.3 Variation in speckle contrast over field SAS image. (a) SAS root magnitude image. (b) Speckle contrast of SAS image measured over 16x8 pixel window.

From (3.5) and (3.6), the contrast of fully developed speckle is unity. Fig. 3.3 shows the variation in speckle contrast measured over small patches of a field SAS with strong targets. The contrast in areas with no targets is close to unity.

Speckle contrast can be reduced by a number of speckle reduction techniques, which also change the speckle distribution. In SAR, it is commonplace to deal with multi-look images. A number of independent looks of the same scene are summed in intensity to reduce the speckle level. Correspondingly, the intensity distribution is not the same as that of fully developed speckle. The intensity may be modeled as a gamma distribution with a controlling parameter of N_{eff}, the effective number of independent looks used to form the image [Henderson and Lewis, 1998]. Multi-look techniques are not currently commonly used in SAS, so pixel statistics will be considered from single-look images only. Multi-look techniques for speckle reduction in SAS are considered in more detail by Fortune et al. [2003b].

Some studies modelling the statistics of sonar returns from various surfaces can show a departure from the standard Rayleigh magnitude model, finding a better fit using log-normal [Gensane, 1989] and generalised K-distributions [Dunlop, 1997; Jakeman and Pusey, 1976]. However, these studies look at the distribution of pixels over a large region; not the distribution of a single pixel. Spatial structure on the seafloor, such as sediment ripples or large stones, is resolved, affecting the region's pixel statistics. For KiwiSAS data, if the statistics of a large patch are examined, there is a variation in mean value over the patch and the magnitude distribution differs from Rayleigh. If a smaller patch is used, the distribution matches a Rayleigh distribution well as shown in Fig. 3.2.

3.1.2 Second-order statistics

It is often assumed that each pixel of speckle is independent or is spatially *delta-correlated*. However, to have delta-correlated speckle requires infinite system bandwidth (temporal and spatial). SA image speckle results from a wide-bandwidth random process being filtered by the imaging system, so necessarily has spatial correlation. To describe speckle fully, consider the coarseness of its spatial structure or the *speckle size*. This can be best described by the autocorrelation function of the speckle or by its Fourier transform, power spectral density, or Weiner spectrum [Goodman, 1976]. The measured autocorrelation of a patch of speckle in a field SAS image is shown in Fig. 3.4. The speckle size is approximately the size of the main peak of the speckle autocorrelation.
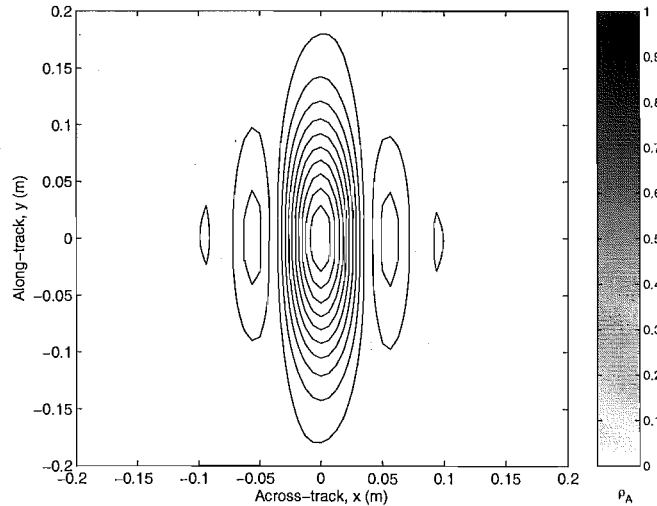
Figure 3.4 Measured 2D correlation coefficient $\rho_A(x, y)$ of a patch of speckle in SAS field image.

Speckle correlation

Consider the autocorrelation of the observed complex field $A(x, y)$, referred to as the *mutual intensity* in optics, given by

$$R_A(\Delta x, \Delta y) = E\left[A(x + \Delta x, y + \Delta y)A^*(x, y)\right]. \tag{3.9}$$

The speckle Weiner spectrum, or power spectral density, is the Fourier transform of the speckle autocorrelation, given by

$$S_A(f_x, f_y) = \int\int\limits_{-\infty}^{\infty} R_A(\Delta x, \Delta y) \exp\left(-j2\pi(f_x\Delta x + f_y\Delta y)\right) d\Delta x\, d\Delta y. \tag{3.10}$$

Consider the geometry shown in Fig. 3.5 for a monochromatic signal. The reflection
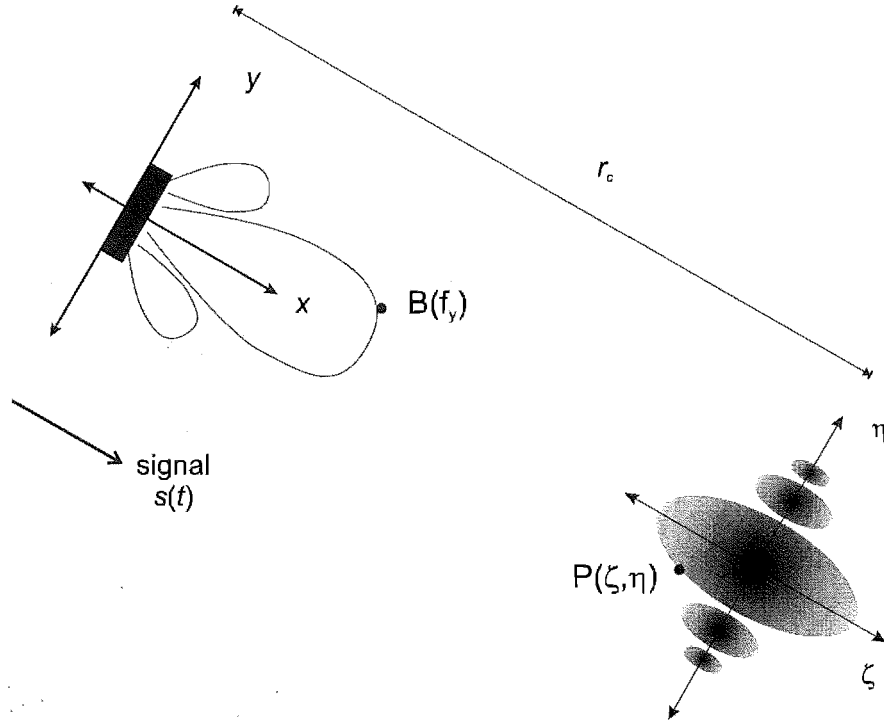
Figure 3.5 Geometry of speckle formation in side-looking imagery. $P(\zeta, \eta)$ is amplitude of scattering area observed by transducer at $A(x, y)$. Transducer has beam-pattern of $B(f_y)$.

from an illuminated area of a rough surface (scattering area) freely propagates over space, with $A(x, y)$ the observed field in a plane parallel to the scattering area. The autocorrelation of the observed field is related to the complex amplitude of the field incident on the scattering area $P(\zeta, \eta)$ by [Goodman, 1975],

$$R_A(\Delta x, \Delta y) = \frac{\kappa}{r_c^2 \lambda^2} \int\!\!\!\int_{-\infty}^{\infty} |P(\zeta, \eta)|^2 \exp\left[j\frac{2\pi}{r_c\lambda}(\zeta\Delta x + \eta\Delta y)\right] d\zeta\, d\eta, \qquad (3.11)$$

where κ is a proportionality constant, λ is the wavelength, and r_c is the separation of the scattering and observation planes. Taking a Fourier transform gives the speckle Weiner spectrum of

$$S_A(f_x, f_y) = \kappa |P(r_c\lambda f_x, r_c\lambda f_y)|^2. \qquad (3.12)$$

This shows $P(\zeta, \eta)$ can be considered as a filter band-limiting the spatial frequencies contributing to the speckle from the wide-bandwidth scatterers on the seafloor. Since $A(x, y)$ is a zero mean process, a normalised version of the covariance (the correlation

coefficient) is given by

$$\rho_{\mathrm{A}}(\Delta x, \Delta y) = \frac{\mathrm{R_A}(\Delta x, \Delta y)}{\mathrm{R_A}(0,0)} \tag{3.13}$$

$$= \frac{\int\limits_{-\infty}^{\infty}\!\!\!\int |P(\zeta, \eta)|^2 \exp\left[j\frac{2\pi}{r_c\lambda}(\zeta\Delta x + \eta\Delta y)\right] d\zeta\, d\eta}{\int\limits_{-\infty}^{\infty}\!\!\!\int |P(\zeta, \eta)|^2\, d\zeta\, d\eta}. \tag{3.14}$$

In SAR and optical wavelengths, intensity (rather than complex amplitude) are used. Thus speckle *intensity* correlation is used. Intensity correlation is related to the complex correlation by

$$\mathrm{R_I}(\Delta x, \Delta y) = \mu_I{}^2 \left[1 + |\rho_{\mathrm{A}}(\Delta x, \Delta y)|^2\right], \tag{3.15}$$

with correlation coefficient

$$\rho_{\mathrm{I}}(\Delta x, \Delta y) = |\rho_{\mathrm{A}}(\Delta x, \Delta y)|^2. \tag{3.16}$$

For simulation purposes, the correlation coefficient of the real (or imaginary) parts of the echo is useful. It can be shown that this is equal to the complex correlation coefficient, giving

$$\rho_{\mathrm{X}}(\Delta x, \Delta y) = \rho_{\mathrm{Y}}(\Delta x, \Delta y) = \rho_{\mathrm{A}}(\Delta x, \Delta y). \tag{3.17}$$

In SA images, the scattering area $P(\zeta, \eta)$ is the patch illuminated by the beam-pattern. Since the two dimensions are independent, the correlation coefficient can be separated;

$$\rho_{\mathrm{A}}(\Delta x, \Delta y) = \rho_{\mathrm{Ax}}(\Delta x)\, \rho_{\mathrm{Ay}}(\Delta y). \tag{3.18}$$

Speckle correlation in along-track

In the along-track direction, consider a rectangular aperture illuminating a flat surface. The surface is then imaged with a (possibly separate) rectangular aperture. For a uniformly illuminated rectangular aperture, the aperture illumination function is given by

$$Q(y) = \mathrm{rect}\left(\frac{y}{D}\right), \tag{3.19}$$

where D is the length of the aperture. From diffraction theory, the radiation pattern in the far-field is shaped by the Fourier transform of the aperture illumination function [Soumekh, 1994]. Thus a rectangular aperture has an amplitude beam pattern given by

$$B(f_y) = \mathrm{sinc}\,(f_y D), \tag{3.20}$$

where f_y is the wavenumber of the along-track direction. For one-way propagation, $f_y = \frac{k}{2\pi}\sin(\theta) = \frac{\sin(\theta)}{\lambda}$, where θ is angle off boresight. Combining the effect of a transmit aperture (length D_T) and receiver aperture (length D_R), the combined transfer function at a single frequency w is

$$H(w, x, y) \approx P(w) B_T(w, x, y) B_R(w, x, y) \frac{\exp\left(-j2k\sqrt{x^2 + y^2}\right)}{(x^2 + y^2)}. \tag{3.21}$$

This uses the 'stop-and-hop' assumption [Hawkins, 1996], where the two-way phase is the time taken for the transmitted signal to travel from the transmit aperture to the point (x, y) and back to the (stationary) receive aperture. The two-way wavenumber increases at twice the rate of the one-way system. Thus a moving active imaging system, such as spotlight or stripmap SA imagery, has twice the spatial bandwidth and twice the angular resolution of a passive system. The overall amplitude pattern scales as $f_y = \frac{2k}{2\pi}\sin(\theta)$ [Hawkins, 1996] giving

$$
\begin{aligned}
B(f_y) &= B_T\left(\frac{f_y}{2}\right) B_R\left(\frac{f_y}{2}\right) \\
&= \operatorname{sinc}\left(\frac{f_y D_T}{2}\right) \operatorname{sinc}\left(\frac{f_y D_R}{2}\right).
\end{aligned} \tag{3.22}
$$

For a scattering area $P(\zeta, \eta)$ centreed at $(r_c, 0)$, $\sin\theta = \eta/\left[r_c^2 + \eta^2\right]^{1/2}$. For $r_c \gg \eta$, $\sin\theta \approx \eta/r_c$ giving a scattering area amplitude function of

$$P(\eta) \propto \operatorname{sinc}\left(\frac{D_T \eta}{2r_c \lambda}\right) \operatorname{sinc}\left(\frac{D_R \eta}{2r_c \lambda}\right). \tag{3.23}$$

Substituting (3.23) into (3.12) gives the speckle Weiner spectrum in the along-track direction as

$$S_{\mathrm{Ay}}(f_y) \propto \left|\operatorname{sinc}\left(f_y D_T/2\right) \operatorname{sinc}\left(f_y D_R/2\right)\right|^2 \tag{3.24}$$

This shows the speckle spectrum, and hence speckle size, is independent of range or wavelength as expected. Taking an inverse Fourier transform gives the speckle autocorrelation as

$$\rho_{\mathrm{Ay}}(\Delta y) \propto \left[\operatorname{rect}\left(\frac{2\Delta y}{D_T}\right) \odot \operatorname{rect}\left(\frac{2\Delta y}{D_R}\right)\right] \star \left[\operatorname{rect}\left(\frac{2\Delta y}{D_T}\right) \odot \operatorname{rect}\left(\frac{2\Delta y}{D_R}\right)\right]. \tag{3.25}$$

The measured autocorrelation of a patch of speckle from a field SAS image is compared to the model (3.25) in Fig. 3.6(a), showing a reasonable match.

Speckle correlation in across-track

In SAS or SAR, an image is formed in across-track by echo-ranging. Distance is measured by the delay between a transmitted and received pulse. Delay τ is mapped to
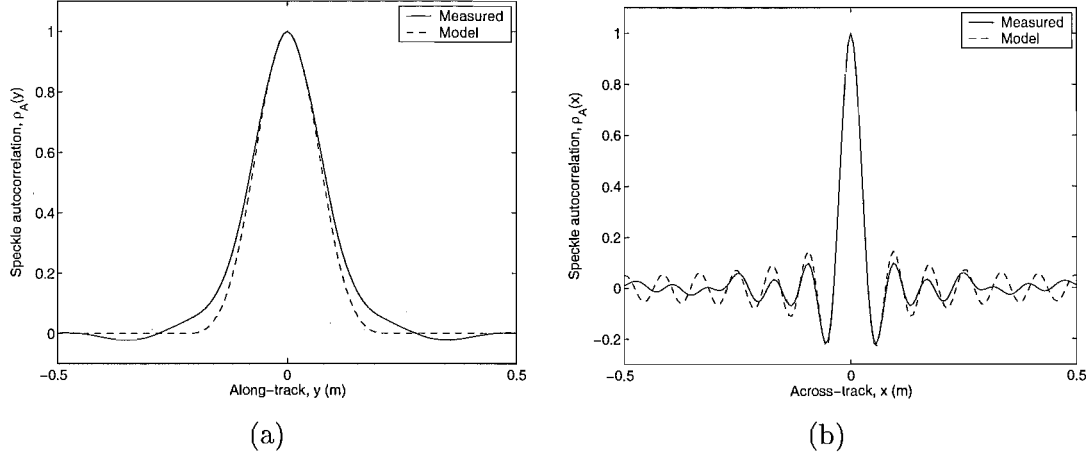
Figure 3.6 Measured autocorrelation of patch of speckle in SAS field image. (a) In along-track direction, compared to model given by (3.25). (b) In across-track direction, compared to model given by (3.33).

range x by

$$x = \frac{c\tau}{2}. \tag{3.26}$$

A standard technique to increase SNR is to transmit an extended, wide-bandwidth signal $s(t)$, and *pulse-compress* or *match filter* the received echo $e(t, y)$ by correlating with the transmitted signal. Thus the image obtained of a point at range x is found by

$$d(x, y) = \int_{-\infty}^{\infty} e(t, y) s^* \left(t - \frac{2x}{c} \right) dt \tag{3.27}$$

$$= e \left(\frac{2x}{c}, y \right) \star s \left(\frac{2x}{c} \right). \tag{3.28}$$

If a perfect point target is imaged, the received echo $e(t, y)$ is a delayed and scaled version of the transmitted signal $s(t)$, i.e. $e(t, y) = a \cdot s(t - \tau(y))$ and

$$d(t, y) = a \cdot s(t - \tau(y)) \star s(t)$$
$$= a \cdot ss(t - \tau(y)), \tag{3.29}$$

where $ss(t)$ is the autocorrelation of the signal; $ss(t) = s(t) \star s(t)$.

The factor causing the band-limiting of spatial frequencies in the speckle image is not the width of the beam-pattern, but the bandwidth of the transmitted chirp. If $S(f)$ is the spectrum of the transmitted signal $(s(t) \longleftrightarrow S(f))$, the illuminated scattering patch in across-track is given by

$$P(k_x) = \left| S \left(\frac{ck_x}{2} \right) \right|^2. \tag{3.30}$$

Transforming to the spatial co-ordinates of the patch with $k_x = \frac{\zeta}{\lambda z}$ gives

$$P(\zeta) = \left| S\left(\frac{c\zeta}{2\lambda z}\right) \right|^2 . \qquad (3.31)$$

Substituting into (3.12) gives the speckle spectrum

$$S_A(k_x) \propto \left| S\left(\frac{ck_x}{2}\right) \right|^4 . \qquad (3.32)$$

Taking an inverse Fourier transform gives the speckle autocorrelation in across-track

$$\rho_A(\Delta x) \propto ss\left(\frac{2\Delta x}{c}\right) \star ss\left(\frac{2\Delta x}{c}\right) . \qquad (3.33)$$

KiwiSAS transmits a linear FM chirp which has the baseband form

$$s_b(t) = \text{rect}\left(\frac{t}{T_c}\right) \exp\left(j\pi\frac{B_c}{T_c}t^2\right) , \qquad (3.34)$$

where B_c is the chirp bandwidth and T_c the chirp duration. The autocorrelation of the chirp is given by

$$ss_b(t) = T_c \,\text{tri}\left(\frac{t}{T_c}\right) \text{sinc}\left(\text{tri}\left(\frac{t}{T_c}\right) B_c t\right) . \qquad (3.35)$$

An approximate form of the Fourier transform of the chirp $s_b(t)$ can be obtained using the principle of stationary phase [Hawkins, 1996], giving

$$S_b(f) = \text{rect}\left(\frac{f}{B_c}\right) \sqrt{\frac{jT_c}{B_c}} \exp\left(-j\pi\frac{T_c f^2}{B_c}\right) \qquad (3.36)$$

The measured autocorrelation in across-track of a patch of speckle from a field SAS image is compared to the model using (3.33) and (3.35) in Fig. 3.6(b), showing a reasonable match.

Speckle correlation and system impulse response

It is often stated that the speckle correlation is equal to (or closely related to) the system impulse response [Henderson and Lewis, 1998; Vachon and Raney, 1989; Raney, 1983]. This statement requires some clarification. A useful (but not exact) way to consider speckle is as a delta-correlated, white-noise random process (scatterers) being filtered by a band-limited filter (SA imaging system). The average power-spectral-density of the filter output is an approximation of the squared magnitude of the filter response $S_A(\omega) \approx |H(\omega)|^2$. Hence the autocorrelation of the filter output (speckle) is not an approximation of the system impulse response, but is an approximation of the impulse response *autocorrelation*; $R_A(\tau) \approx h(\tau) \star h(\tau)$. This can be seen from (3.12) and (3.11)

where the illuminated patch $P(\zeta, \eta)$ is equivalent to the system transfer function $H(\omega)$. In the across-track direction, (3.33) shows the speckle autocorrelation is equal to the autocorrelation of $ss(2x/c)$, where $ss(2x/c)$ is the impulse response in the across-track direction. In the along-track direction, (3.25) shows the speckle autocorrelation is equal to the autocorrelation of the impulse response of $\text{rect}\left(\frac{2\Delta y}{D_T}\right) \odot \text{rect}\left(\frac{2\Delta y}{D_R}\right)$. Fig. 3.7 compares the measures speckle autocorrelation to the magnitude of a simulated point target and its autocorrelation. This comparison confirms that the impulse-response autocorrelation is a closer match to the speckle correlation than the impulse-response alone.
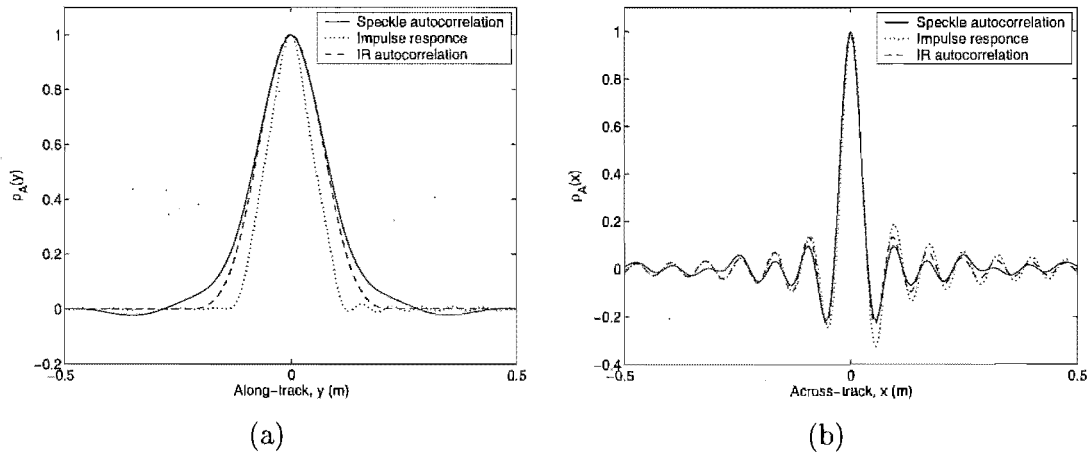
Figure 3.7   Measured autocorrelation of patch of speckle in simulated SAS image, compared to response of simulated point target and autocorrelation of impulse response. (a) In along-track direction. (b) In across-track direction.

While this similarity applies for a perfectly focused image, the speckle spectrum and system transfer-function diverge if the SA processor is defocused. The system impulse-response is sensitive to scene coherence, system coherence and system focus. The speckle autocorrelation is independent of scene coherence and system focus [Raney, 1983]. Thus it is possible to estimate the system impulse response (or its Fourier transform; the system transfer function) from the speckle autocorrelation (or speckle intensity spectrum) of a deliberately defocused speckle image [Vachon and Raney, 1989]. Defocus does not blur the speckle, but it does blur any correlated component of the image, reducing the image correlation function to the system psf. Fig. 3.8 shows the effect of system defocus on the speckle autocorrelation of a SAS image is minimal. In fact, there is no way that first or second-order measures of distributed Gaussian random scenes (fully developed speckle) can be used to infer scene coherence [Raney, 1983]. This indicates that measures of speckle alone (i.e. if no coherent target is present in an image) cannot be used to focus an image.
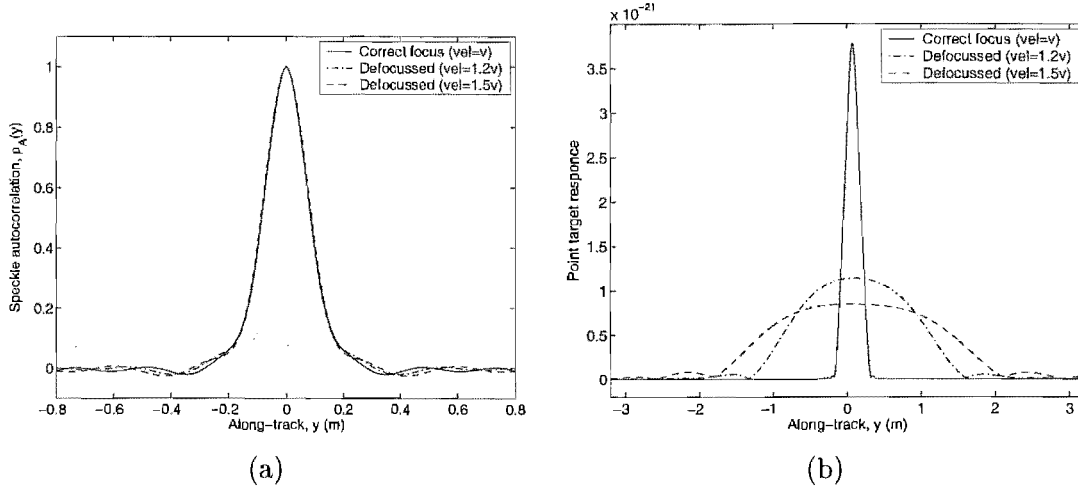
Figure 3.8 Effect of defocus in SAS image by varying the reconstruction velocity. (a) Autocorrelation of patch of speckle in along-track direction. (b) Response of simulated point target in along-track direction.

Speckle size

To indicate the average size of the observed speckle, a measure called the *speckle size* is often used. It can be defined as the statistical average of the distance between adjacent regions of maximum or minimum brightness, or more usefully, as the distance shifted before the speckle autocorrelation drops to close to an arbitrarily small value. This is sometimes called the *correlation length* and is a similar measure.

It is often stated that the speckle size is equal to the resolution of the system. Although not a precise statement (as it depends on what the definitions of resolution and speckle size are), it is a useful rule of thumb. However, because the speckle autocorrelation can be considered as the autocorrelation of the impulse response, the speckle correlation is normally wider. If both receiver and transmitter are of extent D in along-track, the impulse response is approximately triangular in shape, of width D. This has a half power (3dB) width of $D/2$, the normally quoted along-track resolution. The speckle autocorrelation goes to zero at $\Delta y = D$. At a shift of $D/2$, the correlation drops by $\frac{1}{4}$. The most comparable measure to the 3dB width is for the correlation to drop to a half, which occurs at $\Delta y = D/2.77$. Considering the inexact nature of the measure, system resolution is a good indication of speckle size. In the across-track direction, the impulse response of a rectangular shaped pulse is approximately a sinc. The autocorrelation of a sinc is unchanged, making the speckle size and system resolution almost identical, (see Fig. 3.7). The following expressions for speckle size in across-track (L_x) and along-track (L_y) will be used

$$L_x = \frac{c}{2B_c} \tag{3.37}$$

$$L_y = \frac{D}{2.77}. \tag{3.38}$$

Effect of range and aperture sampling

Synthetic aperture imagery is normally sampled. Thus the correlation between adjacent pixels depends on the sampling rate. If the sampling rate is low, the sample spacing can be larger than the speckle size. Thus each pixel can be accurately described as uncorrelated. Many authors describe this as delta-correlated speckle, although care needs to be taken as this is dependent on the sampling rate. However, if the system is sampled at a rate lower than the speckle size, the system will be undersampled. This is because the speckle size is approximately the size of the system resolution.

In SAS, pixel correlation in along-track drops to a half at $D/2.77$. However, $D/4$ sampling is required to sample along-track adequately [Hawkins, 1996]. Thus most systems will have significant correlation between adjacent pixels at normal operating speed. In across-track, the speckle size is approximately $\frac{c}{2B_c}$, where B_c is the chirp bandwidth. The sample spacing is given by $\frac{c}{2f_s}$ where f_s is the sampling frequency. For Nyquist sampling, it is necessary that $f_s > B_c$. Thus adjacent samples have some correlation in across-track. In KiwiSAS, $B_c = 20\,\text{kHz}$ and $f_s = 30\,\text{kHz}$, thus there is a small correlation between adjacent pixels. The more oversampled the system, the higher the correlation.

A useful measure of the total amount of correlation (which is dependent on the sampling rate) is the correlation area A_c. This is defined by

$$A_c \equiv \sum_{r=1-N_m}^{N_m-1} \sum_{s=1-N_n}^{N_n-1} |\rho_{\text{A}}(r,s)| . \tag{3.39}$$

It is apparent that the higher the sampling rate compared to the correlation length, the closer the spacing of r, s and the higher the correlation area. For delta-correlated speckle, the correlation area is unity. The correlation area can be calculated using (3.25), (3.33) and (3.18). For KiwiSAS, the correlation area is shown as a function of the ratio of the speckle size to the sample spacing in Fig. 3.9. The high level of quantisation in the correlation area as Δy is changed occurs as pixels move from outside to inside the correlation area. The large value of correlation area for $\Delta x = L_x$ is due to the large sidelobes in the correlation in the across-track direction.

3.1.3   Imaging with speckle

For fully developed speckle, it is standard to model speckle as a random process appearing as a multiplicative modulation of the scene (intensity) reflectivity [Lowenthal and Arsenault, 1970; Goodman, 1976; Lee, 1986; Marron and Morris, 1986], i.e.,

$$I(x,y) = V(x,y)U(x,y), \tag{3.40}$$

(a)   (b)
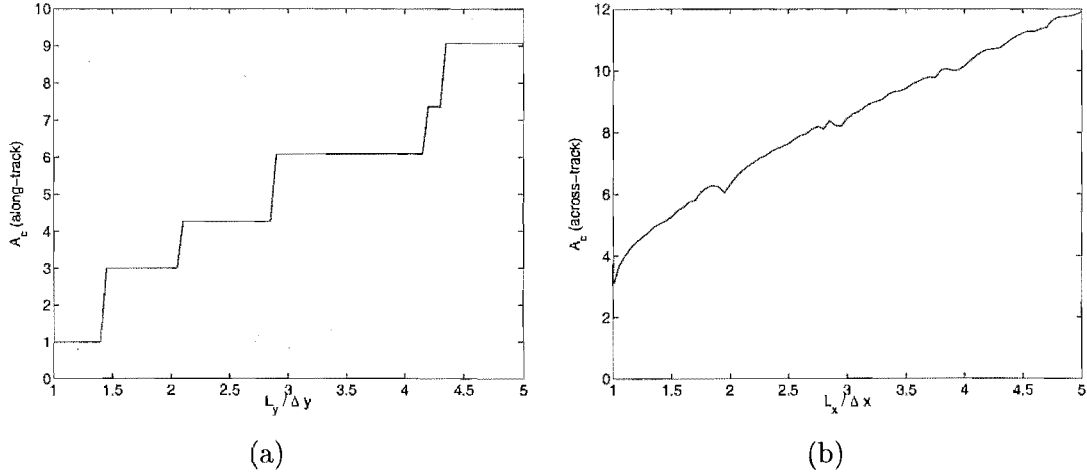
Figure 3.9 Correlation area as a function of the ratio of speckle-size to sample spacing for KiwiSAS system. (a) In along-track direction. (b) In across-track direction.

where $I(x,y)$ is the image intensity at point (x,y), $V(x,y)$ is the noncoherent image of the object and $U(x,y)$ is the speckle noise. In any measured image, there is a component of additive noise but this is usually much smaller than the speckle noise for reverberation-limited imagery and is ignored in this analysis.

This multiplicative model is valid only for fully developed speckle and when there is only a small change in contrast within each neighborhood being modeled. It has been shown [Tur et al., 1982] that if the object being imaged has spatial details which cannot be resolved by the coherent system, the model has significant errors in that region.

Using this model, the speckle noise statistics are assumed constant over the whole image. Any variation in the mean intensity of a point, caused by a target or fading with range, is contained in the noncoherent component $V(x,y)$. There is still an arbitrary scale factor between the two components in (3.40). Some authors remove this by defining $\mu_U = 1$. For the analysis in this chapter, the arbitrary scale factor is left undefined. If U is stationary, the mean and variance of image intensity I is given by

$$\mathrm{E}\left[I(x,y)\right] = \mu_U V(x,y) \qquad\qquad (3.41)$$

$$\mathrm{Var}\left[I(x,y)\right] = \sigma_U^2 V^2(x,y), \qquad\qquad (3.42)$$

where

$$\mu_U = \mathrm{E}\left[U(x,y)\right]$$

$$\sigma_U^2 = \mathrm{Var}\left[U(x,y)\right].$$

3.2   STATISTICS OF IMAGE SHARPNESS

A popular image sharpness measure consists of the mean of all the image pixels intensity raised to a power, i.e.,

$$S_\beta = \frac{1}{N_m N_n} \sum_{m=1}^{N_m} \sum_{n=1}^{N_n} I^\beta[m,n]. \tag{3.43}$$

To simplify notation, pixels will be referred to by a single index p, with $I[m,n] = I_p$ giving,

$$S_\beta = \frac{1}{N_p} \sum_p I_p^\beta \tag{3.44}$$

where $N_p = N_m N_n$ is the number of pixels summed. Using the multiplicative noise model from (3.40), this can be written as

$$S_\beta = \frac{1}{N_p} \sum_p V_p^\beta U_p^\beta. \tag{3.45}$$

Defining the random variable $Z_p \equiv U_p^\beta$ gives

$$S_\beta = \frac{1}{N_p} \sum_p V_p^\beta Z_p. \tag{3.46}$$

The distribution of the sharpness depends on the noncoherent image V and the distribution of the modified speckle intensity Z.

3.2.1   Distribution of Z

Consider a fully developed speckle image, where U has a probability distribution given by (3.3) with mean intensity σ_U. The probability distribution of Z is given by

$$\Pr(Z) = \begin{cases} \frac{Z^{(1/\beta - 1)}}{\beta \sigma_U} \exp\left(-\frac{Z^{1/\beta}}{\sigma_U}\right) & Z \geq 0 \\ 0 & \text{otherwise.} \end{cases} \tag{3.47}$$

From (3.4), the mean of Z is

$$\mathrm{E}[Z] = \mu_Z = \beta! \sigma_U^\beta \tag{3.48}$$

if β is an integer. Otherwise $\beta!$ is replaced with $\Gamma(\beta+1)$ where appropriate. To simplify the notation, factorial notation will continue to be used, although β is not restricted to being an integer. The result in (3.48) also follows directly from Reed's Gaussian moment theorem [Reed, 1962]. The variance is given by

$$\begin{aligned} \mathrm{Var}[Z] = \sigma_Z^2 &= \mathrm{E}[Z^2] - (\mathrm{E}[Z])^2 \\ &= \left[(2\beta)! - (\beta!)^2\right] \sigma_U^{2\beta}. \end{aligned} \tag{3.49}$$

It can then be shown that

$$\mathrm{E}\left[X_p^4 X_q^2\right] = \mathrm{E}\left[X_p^2 X_q^4\right] = \mathrm{E}\left[Y_p^4 Y_q^2\right] = \mathrm{E}\left[Y_p^2 Y_q^4\right] = 12\left|\rho_\mathrm{A}\right|^2 + 3 \qquad (3.50)$$

$$\mathrm{E}\left[X_p^4 X_q^4\right] = \mathrm{E}\left[Y_p^4 Y_q^4\right] = 24\left|\rho_\mathrm{A}\right|^4 + 72\left|\rho_\mathrm{A}\right|^2 + 9. \qquad (3.51)$$

The most commonly used sharpness measure is $\beta = 2$, giving $Z = U^2$ and

$$\mathrm{E}\left[Z_p Z_q\right] = \sigma_U^4/16\left(2\,\mathrm{E}\left[X_p^4 X_q^4\right] + 8\,\mathrm{E}\left[X_p^4 X_q^2\right]\mathrm{E}\left[X_p^2\right] + 2\,\mathrm{E}\left[X_p^4\right]\mathrm{E}\left[X_q^4\right] + 4\,\mathrm{E}\left[X_p^2 X_q^2\right]^2\right)$$
$$= 4\sigma_U^4\left(\left|\rho_\mathrm{A}\right|^4 + 4\left|\rho_\mathrm{A}\right|^2 + 1\right).$$

$$(3.52)$$

The correlation coefficient of Z is then

$$\rho_Z = \frac{\mathrm{E}\left[Z_p Z_q\right] - \mathrm{E}\left[Z_p\right]\mathrm{E}\left[Z_q\right]}{\sqrt{\mathrm{Var}\left[Z_p\right]\mathrm{Var}\left[Z_q\right]}}$$
$$= \frac{\left|\rho_\mathrm{A}\right|^4 + 4\left|\rho_\mathrm{A}\right|^2}{5}. \qquad (3.53)$$

3.2.2 Mean of sharpness measure

Since Z is spatially stationary, $\mathrm{E}\left[Z_p\right]$ and $\mathrm{Var}\left[Z_p\right]$ are constant for all p,

$$\mathrm{E}\left[S_\beta\right] = \mathrm{E}\left[Z\right]\frac{1}{N_p}\sum_p V_p^\beta$$
$$= \mathrm{E}\left[Z\right]\widehat{S_\beta} \qquad (3.54)$$

where $\widehat{S_\beta}$ is the sharpness of the noncoherent image of the object. For fully developed speckle, (3.48) gives

$$\mathrm{E}\left[S_\beta\right] = \beta!\sigma_U^\beta\widehat{S_\beta}. \qquad (3.55)$$

Fig. 3.10 shows the mean sharpness value for a number of simulated speckle patterns as σ_U is varied. The results closely match those predicted in (3.55). This matches results for laser speckle [Marron and Morris, 1986] and SAR [Paxman and Marron, 1988].

3.2.3 Variance of sharpness measure

The variance of the sharpness is given by

$$\mathrm{Var}\left[S_\beta\right] = \frac{\sigma_Z^2}{N_p^2}\sum_p\sum_q V_p^\beta V_q^\beta \,\rho_{Z_\mathrm{p},Z_\mathrm{q}}. \qquad (3.56)$$

Thus the sharpness variance depends on the speckle correlation coefficient as well as the autocorrelation of the image scene V. This general form for the sharpness variance
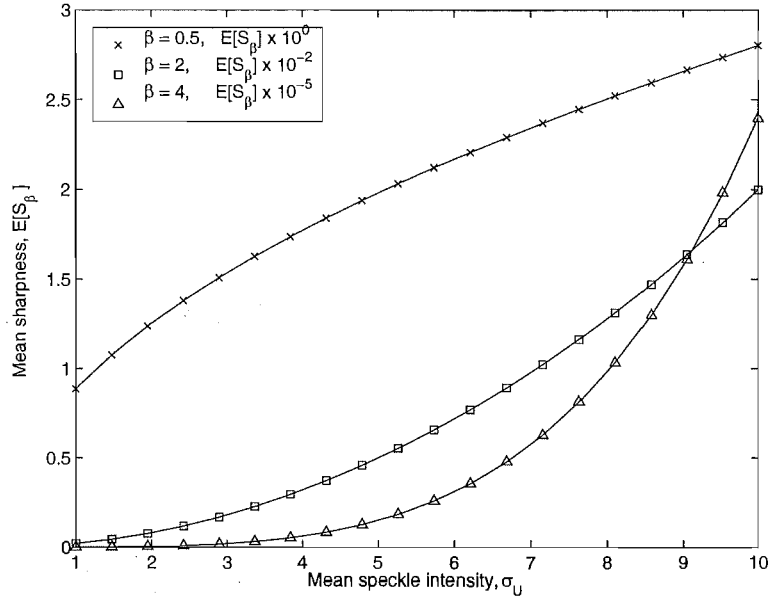
Figure 3.10 Mean sharpness of simulated speckle images as a function of the mean speckle intensity σ_U. Data is compared with the model (3.55) (solid lines).

can be calculated directly with difficulty. For a simpler expression, assumptions need to be made about the form of the speckle correlation, or of the object being imaged.

Uncorrelated speckle

Evaluation of the sharpness variance is straightforward if the speckle is delta-correlated, i.e. $\rho_U(\Delta x, \Delta y) = \rho_Z(\Delta x, \Delta y) = 0$ for $\Delta x, \Delta y \neq 0$. Equation (3.56) then becomes,

$$
\begin{aligned}
\text{Var}\,[S_\beta] &= \frac{\sigma_Z^2}{N_p^2} \sum_p V_p^{2\beta} \\
&= \frac{\sigma_Z^2}{N_p} \widehat{K_\beta},
\end{aligned}
\tag{3.57}
$$

defining

$$
\widehat{K_\beta} \equiv \frac{1}{N_p} \sum_{p=1}^{N_p} V_p^{2\beta}.
$$

Many authors make the assumption that speckle is delta-correlated, although in SAS to be adequately sampled requires the speckle of adjacent pixels to have some correlation (Section 3.1.2). The KiwiSAS system is normally only just adequately sampled, meaning adjacent pixels have small correlation. For fully developed, delta-correlated speckle, (3.49) gives

$$
\text{Var}\,[S_\beta] = \frac{1}{N_p} \left[(2\beta)! - (\beta!)^2 \right] \sigma_U^{2\beta} \widehat{K_\beta}.
\tag{3.58}
$$

Fig. 3.11 shows how $\text{Var}[S_2]$ varies with σ_U for a number of simulated uncorrelated speckle patterns. The data matches (3.58) closely.
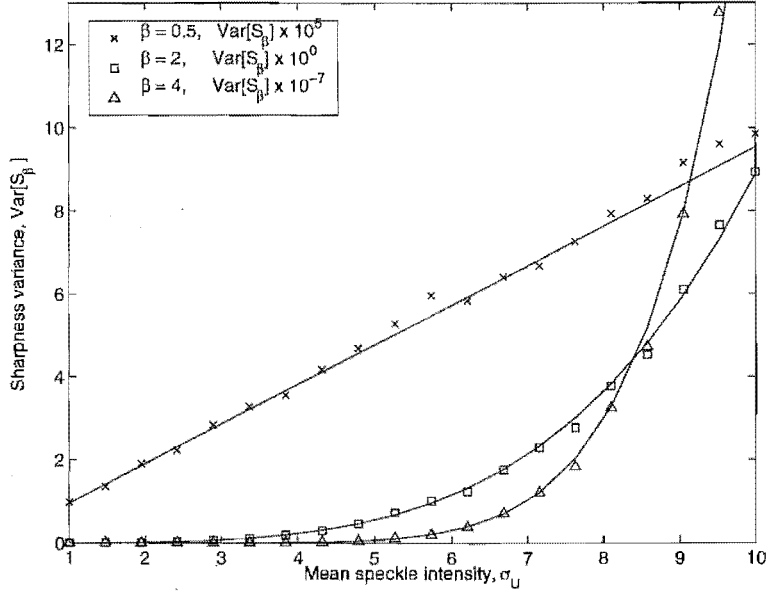
Figure 3.11 Variance of sharpness of simulated uncorrelated speckle images as a function of the mean speckle intensity σ_U. Data is compared with the model (3.58) (solid lines).

Slowly varying, correlated speckle

Rearranging (3.56) gives

$$\text{Var}[S_\beta] = \frac{\sigma_Z^2}{N_p^2} \sum_{m=1}^{N_m} \sum_{n=1}^{N_n} \sum_{r=1-N_m}^{N_m-1} \sum_{s=1-N_n}^{N_n-1} V^\beta[m,n] V^\beta[m+r,n+s] \, \rho_Z(r,s). \tag{3.59}$$

Consider that $V[m,n]$ varies slowly over the correlation area of the speckle, i.e., $V[m,n] \approx V[m+r,n+s]$ for r,s less than the speckle size. The sharpness variance becomes

$$\text{Var}[S_\beta] \approx \frac{\sigma_Z^2}{N_p^2} \sum_{m=1}^{N_m} \sum_{n=1}^{N_n} V^{2\beta}[m,n] \sum_{r=1-N_m}^{N_m-1} \sum_{s=1-N_n}^{N_n-1} \rho_Z(r,s)$$

$$\approx \frac{\sigma_Z^2}{N_p} \widehat{K_\beta} A_z, \tag{3.60}$$

where A_z is the correlation area of Z given by

$$A_z = \sum_{r=1-N_m}^{N_m-1} \sum_{s=1-N_n}^{N_n-1} \rho_Z(r,s). \tag{3.61}$$

The calculated correlation area of Z is shown in Fig. 3.12, as the sample spacing is changed. This result shows A_z is approximately linear to the correlation area A_c.

Fig. 3.13 shows the variance of the sharpness measured for a number of simulated speckle patterns as the correlation area is changed. The results match (3.60) closely. The quantity N_p/A_z may be considered the number of independent speckles in the image. From (3.60), the variance of sharpness is inversely proportional to the number of speckles in the image.
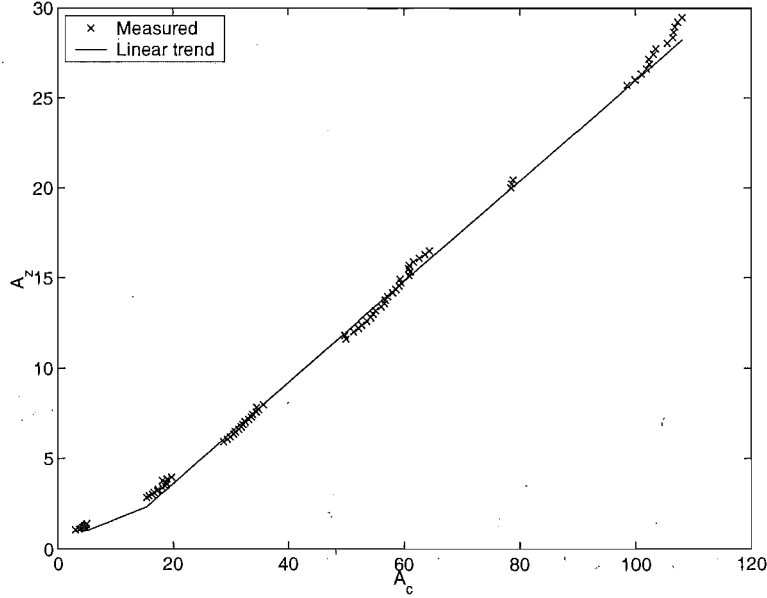
Figure 3.12 The correlation area A_z as a function of the speckle correlation area A_c, for $Z = U^2$. Linear model is $A_z = 0.28A_c - 2$, for $A_z > 1$.

Point target plus correlated speckle

Consider a scene comprising uniform speckle V_c and a point target V_t of target-to-speckle strength ratio α; that is,

$$V_c[m,n] = 1,$$

$$V_t[m,n] = \begin{cases} \alpha & \text{for } [m,n] = [Nt_x, Nt_y], \\ 0 & \text{otherwise,} \end{cases}$$

$$V[m,n] = V_c[m,n] + V_t[m,n].$$

From (3.60)

$$\text{Var}\,[S_\beta] = \frac{\sigma_Z^2}{N_p^2} \sum_m \sum_n \sum_r \sum_s \rho_Z(r,s)\,[V_c[m,n] + V_t[m,n]]^\beta$$

$$\times\,[V_c[m+r,n+s] + V_t[m+r,n+s]]^\beta. \quad (3.62)$$
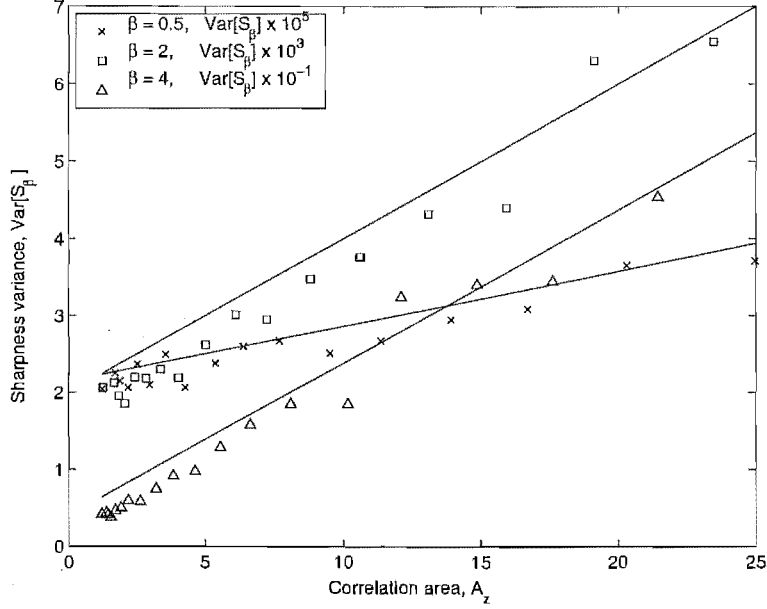
Figure 3.13 Variance of sharpness S_β as a function of speckle correlation area A_z. Data is compared with the model from (3.60) (solid lines).

If $\alpha \gg 1$, $V^\beta \approx V_c^\beta + V_t^\beta$ giving,

$$\text{Var}\,[S_\beta] \approx \frac{\sigma_Z^2}{N_p^2} \sum_m \sum_n \sum_r \sum_s \rho_Z(r,s) \left[V_c^\beta[m,n] + V_t^\beta[m,n] \right]$$
$$\times \left[V_c^\beta[m+r,n+s] + V_t^\beta[m+r,n+s] \right]. \quad (3.63)$$

As $V_c[m,n] = 1$,

$$\text{Var}\,[S_\beta] \approx \frac{\sigma_Z^2}{N_p^2} \sum_m \sum_n \sum_r \sum_s \rho_Z(r,s)$$
$$\times \left[1 + V_t^\beta[m,n] + V_t^\beta[m+r,n+s] + V_t^\beta[m,n]V_t^\beta[m+r,n+s] \right],$$
$$\approx \frac{\sigma_Z^2}{N_p^2} \left[\sum_m \sum_n \sum_r \sum_s \rho_Z(r,s) + \sum_m \sum_n V_t^\beta[m,n] \sum_r \sum_s \rho_Z(r,s) \right.$$
$$+ \sum_m \sum_n \sum_r \sum_s V_t^\beta[m+r,n+s]\, \rho_Z(r,s) \qquad (3.64)$$
$$\left. + \sum_m \sum_n \sum_r \sum_s V_t^\beta[m,n]V_t^\beta[m+r,n+s]\, \rho_Z(r,s) \right],$$
$$\approx \frac{\sigma_Z^2}{N_p^2} \left[\alpha^{2\beta} + (N_p + 2\alpha^\beta)A_z \right].$$

Fig. 3.14 shows the measured variance of the sharpness of a number of simulated speckle patterns with a point target of various α. The results match (3.64) reasonably closely. Increasing α increases the sharpness variance. As $\alpha \to 0$, (3.64) becomes (3.60).
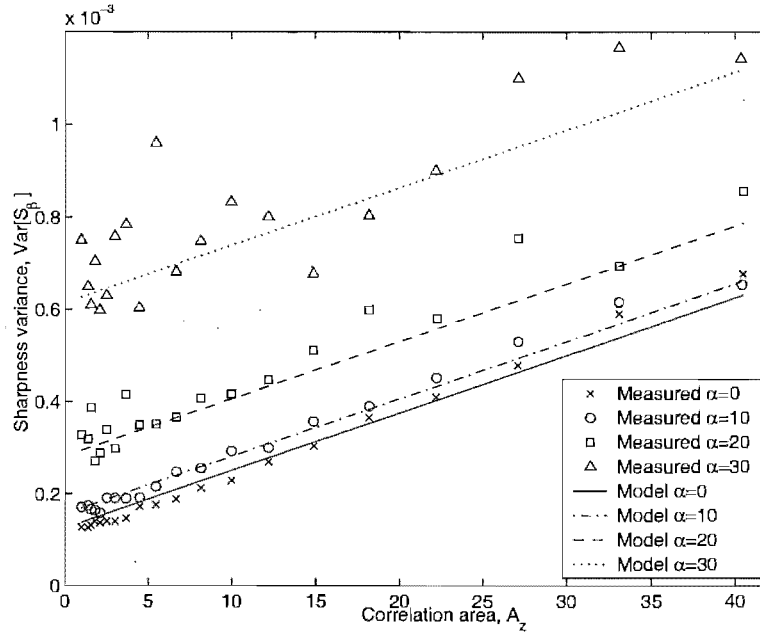
Figure 3.14 Variance of sharpness S_2 for point target strength α plus speckle as a function of speckle correlation area A_z. Speckle intensity σ_U is set so $\mathrm{E}[S_2] = 1$. Data is compared against the model from (3.64).

3.2.4   Sharpness distribution

To determine the probability distribution of the sharpness, assumptions about the correlation between pixels is required. The sharpness ditribution for uncorrelated speckle is compared to the sharpness distribution of correlated speckle.

Uncorrelated speckle

If the speckle is uncorrelated, $Z(m, n)$ is a set of independent and identically distributed (iid) random variables. Thus S_β is a sum of iid random variables, weighted by the noncoherent term $V^\beta(m, n)$. Liapounov's version of the central limit theorem, dealing with independent, heterogeneously distributed variables W_t, states that the summation $X = \frac{1}{N} \sum_{t=1}^{N} W_t$ will be normally distributed under the following conditions [White, 1984]:

1. The variance of W_t is non-zero:

   $\mathrm{Var}\,[W_t] \neq 0, \quad \forall\ t.$

2. The second moment of W_t is finite:

   $\mathrm{E}\left[\,|W_t - \mathrm{E}\,[W_t]\,|^{2+\delta}\right] < \Delta < \infty$ for some $\Delta > 0$ and $\forall\ t.$

3. The quantity $\bar{\sigma}_N^2 = \frac{1}{N} \sum_{t=1}^{N} \mathrm{Var}\,[W_t]$ is non-zero for all N:

   $\bar{\sigma}_N^2 > \delta' > 0, \quad \forall\ N.$

In this case, $W_t = Z(m,n)V^\beta(m,n)$. The sharpness measure will be normally distributed under the following corresponding limitations on the noncoherent image V_p:

1. There is speckle, i.e., $\text{Var}[Z] \neq 0$. In an area of no return where $V_p = 0$, then $W_t = 0$, so the sharpness distribution is not affected. This has been confirmed by experiments.

2. V_p is finite for all points of the image.

3. $\widehat{K} = \frac{1}{N_p}\sum_{i=1}^{N_p} V_p^{2\beta} > \delta'' > 0$.
   This ensures that the variance of C is not zero, which will hold for all practical images.

If the sharpness S is normally distributed, the distribution of S is given by

$$\Pr(S) = \frac{1}{\sqrt{2\pi\,\text{Var}[S]}}\exp\left[-\frac{(S-\text{E}[S])^2}{2\,\text{Var}[S]}\right]. \tag{3.65}$$

Fig. 3.15 (a) compares the results of the sharpness of a number of simulated uncorrelated speckle patterns against the normal distribution given by (3.65), showing a close match.

Correlated speckle

If the speckle is correlated, the sharpness is a sum of non-independent random variables so it is not necessarily normally distributed. Fig. 3.15 (b) shows the probability distribution of simulated highly correlated speckle. It does differ from the normal distribution model but is close enough for the assumption to be useful. In KiwiSAS, the speckle correlation dies to very close to zero within two pixels, which is typical, making the normal distribution model quite accurate.

3.2.5 Results

The derived statistics are for different speckle realisations of the same scene. This could not be easily tested experimentally with SAS data but can be easily verified with speckle simulations.

Simulated images

A 128×128 pixel complex speckle image U was simulated with circular Gaussian distribution and with each pixel independent. The speckle was then multiplied by the desired non-coherent image V. This was repeated for 10 000 different speckle realisations and the sharpness measured for a number of different values of speckle intensity, correlation area and image type V.
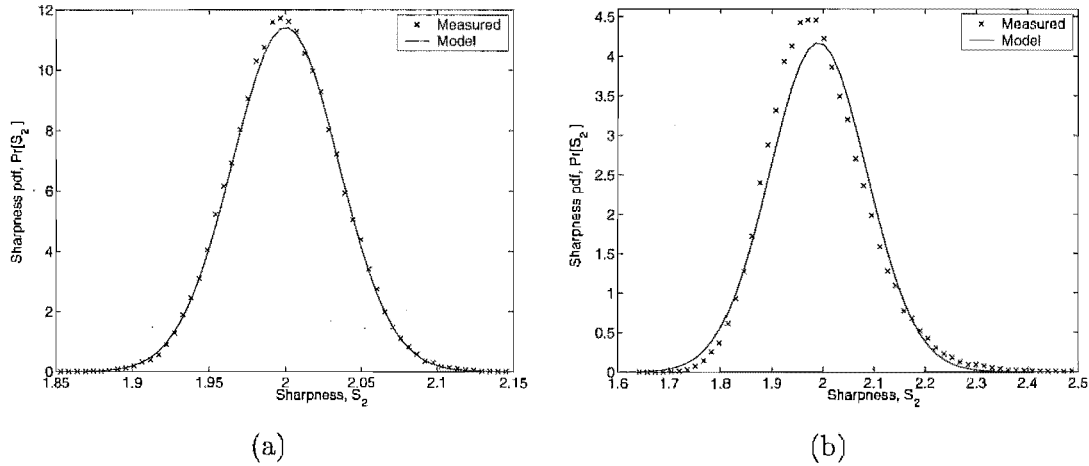
Figure 3.15 Probability distribution of sharpness (S_2) of simulated speckle patterns. Data is compared with the normal distribution model in (3.65). (a) Uncorrelated speckle. (b) Highly correlated speckle. (Correlation area $A_z = 63.5$).

The sharpness was first measured for a range of uncorrelated speckle patterns, varying the mean speckle intensity σ_U. The mean sharpness value matched that predicted by (3.55) well as that shown in Fig. 3.10. The variance of the same data is shown in Fig. 3.11. This fits the model of (3.57) well.

The mean speckle intensity was then held constant and the correlation area A_c varied by filtering an uncorrelated speckle image with a two dimensional filter. This is the simplest method for simulating correlated speckle [Raney and Wessels, 1988]. The mean of the sharpness measurement did not vary with correlation area. Fig. 3.13 shows how the variance of the sharpness varies with correlation area. This data fits the model of (3.60) well.

The image was then simulated to have a point target of various strengths α with background speckle. The correlation area A_c was varied. The resulting variance is shown in Fig. 3.14 and compared to the model in (3.64).

Fig. 3.15 shows the estimated probability distribution of the sharpness for uncorrelated and correlated speckle simulations. The uncorrelated speckle PDF matches the normal distribution model in (3.65) well. The PDF of a highly correlated speckle pattern is slightly skewed from the normal distribution model.

Synthetic aperture sonar images

To test the developed statistics on real data, a number of different speckle realisations of the same image is required. This is not strictly possible, but it can be approximated by taking a number of images of the same object. One way to achieve this is to take a long strip of SAS data of homogeneous, bland seafloor and divide it into small strips of data. The assumption is that each strip is a separate image of the same object.
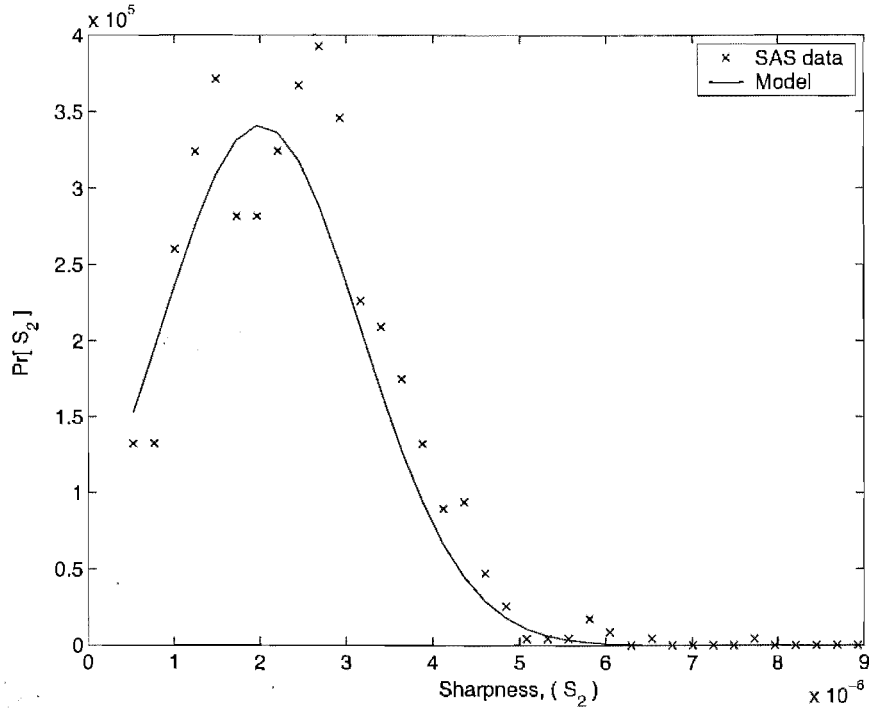
Figure 3.16 Probability distribution of sharpness of strips of SAS data compared with the model in (3.65). The model mean is calculated from (3.55) and the variance from (3.60). Measured correlation area from image is $Az = 1.5$.

The sharpness of each strip can then be calculated and the statistics of the sharpness compared with the model.

To estimate the noncoherent image $V(x,y)$, the intensity of each strip was averaged, then filtered to reduce the variation caused by speckle. The correlation area of the image was measured by averaging the autocorrelation of many patches. The mean and variance of the sharpness were calculated by use of (3.55) and (3.60), respectively. In Fig. 3.16 the probability distribution of the measured sharpness values of 1000 pings of SAS data is compared against that predicted by the model. This shows a reasonable match, considering the small amount of data used and the inexact nature of the comparison.

3.3 CONCLUSIONS

Speckle in a coherent image causes the image intensity and therefore sharpness measure to be a random variable. This chapter develops the probability distribution of the sharpness measure. The intensity of a point can be considered the product of the noncoherent image of the scene $V(x,y)$ and the speckle $U(x,y)$. The distribution of the sharpness measure is dependent on two major factors. The first is the nature of the imaged scene $V(x,y)$. The expected value of the sharpness is proportional to the sharpness of $V(x,y)$. The variance is proportional to the sharpness of $V^2(x,y)$. As a

target is brought into focus by a correct phase estimate, both the mean and variance of the sharpness will be at a peak. The second factor is the correlation of the speckle. The sharpness variance is inversely proportional to the number of independent speckles in the image N_p/A_z. Since many speckle-reducing filters increase the correlation length of speckle as the speckle variance is reduced [Fortune *et al.*, 2003b], speckle reduction does not reduce the variance of the sharpness. The sharpness is normally distributed for uncorrelated speckle and is close to normally distributed for correlated speckle.

Although the sharpness variance increases when there is a target in the image, the sharpness gradient is proportional to the target-to-speckle energy ratio. Thus the ability to perform sharpness optimisation is improved with an increased target-to-speckle ratio.

Chapter 4

ECHO CORRELATION

The most common methods of autofocus correlate adjacent echos to estimate the average phase difference between them. This includes the shear-average, phase-gradient autofocus (PGA), redundant phase centre (RPC) and displaced-phase centre antenna (DPCA) algorithms. This chapter shows a theoretical study of estimating the Fourier phase error of a spotlight image using echo-correlation. A new method, weighted phase difference estimation (WPDE), is introduced, showing superior performance over standard techniques for some images.

The chapter first reviews the well known problem of time-delay estimation (TDE). Various methods of complex cross-correlation are described. Fourier phase-difference estimation (PDE) is described, and shown to be a narrow-band equivalent of TDE. A model of the variance of the Fourier phase of an image is developed. This allows limits to the accuracy of PDE methods to be derived. Several assumptions are made in standard PDE methods that limit the performance on extended objects. The method of PDE is generalised and a new method, weighted phase difference estimation (WPDE) introduced. WPDE uses a measure of coherence to estimate the average signal to clutter ratio of each range bin. This is used to calculate the variance of the phase estimate from each range-bin using the model developed. The phase estimate is then weighted by the inverse variance to estimate the average Fourier phase for each ping. In addition, the variance of each phase estimate can be estimated.

Estimating the phase error of a stripmap image is more difficult, as a target does not have constant phase difference. Several ways to overcome this problem are shown. A new method, weighted phase curvature estimation (WPCE) follows from the development of phase difference estimation.

4.1   TIME DELAY ESTIMATION

Consider a real time signal, $e(t)$ monitored at two spatially separated sensors, with their measurements modelled as

$$p_1(t) = e(t) + n_1(t) \tag{4.1a}$$

$$p_2(t) = \alpha e(t - D) + n_2(t), \tag{4.1b}$$

with the noise $n_1(t)$ and $n_2(t)$ assumed stationary and uncorrelated with the signal. There are many applications where it is of interest to estimate the time delay D between the signals. This problem is also fundamental to many motion estimation methods such as DPCA [Bellettini and Pinto, 2002].

4.1.1   Cross correlation

A common method is to find the argument τ that maximises the cross correlation function of the two signals. A delay estimate $\widehat{D}$ can be found using

$$\widehat{D} = \arg\max_{\tau} \left\{ \widehat{R}_{p_1 p_2}(\tau) \right\}, \tag{4.2}$$

where

$$R_{p_1 p_2}(\tau) \equiv E\left[p_1(t) p_2{}^*(t + \tau) \right]. \tag{4.3}$$

For a finite observation time T, the cross correlation of ergodic signals can be estimated by[1]

$$\widehat{R}_{p_1 p_2}(\tau) = \frac{1}{T} \int_{\tau}^{T} p_1(t) p_2{}^*(t + \tau) \, dt \tag{4.4}$$

$$= \mathcal{F}_{f \to \tau}^{-1} \left\{ P_1(f) P_2{}^*(f) \right\}. \tag{4.5}$$

where $P_k(f)$ is the Fourier transform of $p_k(t)$. This estimator is biased, but has a lower mean-squared error than using the unbiased divisor $\frac{1}{T-\tau}$ instead of $\frac{1}{T}$, so is normally preferable [Jenkins and Watts, 1968]. The cross power spectral density $S_{p_1 p_2}(f) \equiv \mathcal{F}\{R_{p_1 p_2}\}$ is estimated by

$$\widehat{S}_{p_1 p_2}(f) = \frac{1}{T} P_1(f) P_2{}^*(f). \tag{4.6}$$

---

[1]In practice, the correlation obtained using (4.5) will give a different result due to circular convolution. This may be mitigated by employing guard bands at the edge of the time-series.

From the signal model (4.1), the signal spectra are related by

$$P_1(f) = E(f) + N_1(f) \tag{4.7a}$$

$$P_2(f) = \alpha E(f) \exp\left(-j2\pi f D\right) + N_2(f) \tag{4.7b}$$

and since the signal and noise are uncorrelated, the cross correlation is

$$R_{p_1 p_2}(\tau) = \alpha R_{ee}(\tau - D) + R_{n_1 n_2}(\tau) \tag{4.8}$$

and the cross PSD is

$$S_{p_1 p_2}(f) = \alpha \left|E(f)\right|^2 \exp\left(j2\pi f D\right) + S_{n_1 n_2}(f) \tag{4.9}$$

$$= \alpha S_{ee}(f) \exp\left(j2\pi f D\right) + S_{n_1 n_2}(f). \tag{4.10}$$

The signal autocorrelation $R_{ee}(\tau)$ will be maximum for $\tau = 0$. Thus for the noiseless case, $R_{p_1 p_2}(\tau)$ will be maximum for $\tau = D$. This is equivalent to estimating the average slope of the unwrapped phase of $S_{p_1 p_2}(f)$ over all f.

4.1.2 Generalised correlation

The accuracy of time delay estimation depends upon the number of independent time signals used to estimate the cross-correlation, the bandwidth and the signal to noise ratio (SNR) of the measured signals. If the signal and noise have different power spectra, each signal can be filtered to maximise the SNR and hence the accuracy of the delay estimate. This section shows the optimal filter to maximise accuracy of the time delay estimation and determines the accuracy of that estimation.

Consider passing the signal $p_k(t)$ through a filter with spectral response $H_k(f)$ to yield $q_k(t)$. The resultant cross correlation of q_1 and q_2 is then used to estimate D. A good choice of filter $H_k(f)$ will increase the SNR of q resulting in a more accurate estimate of D. The cross PSD of the filter outputs is

$$S_{q_1 q_2}(f) = H_1(f)H_2^*(f) S_{p_1 p_2}(f), \tag{4.11}$$

and the generalised correlation is the inverse Fourier transform,

$$R_{q_1 q_2}(\tau) = \mathcal{F}^{-1}\left\{W(f) S_{p_1 p_2}(f)\right\}, \tag{4.12}$$

where $W(f) = H_1(f)H_2^*(f)$ is a general *frequency weighting* function. The estimate of the delay $\widehat{D}$ maximises the generalised cross correlation $\widehat{R}_{q_1 q_2}$;

$$\widehat{D} = \arg\max_{\tau}\left\{\widehat{R}_{q_1 q_2}(\tau)\right\}, \tag{4.13}$$

where

$$\widehat{R}_{q_1 q_2}(\tau) = \mathcal{F}^{-1}\left\{ W(f)\widehat{S}_{p_1 p_2}(f) \right\}. \tag{4.14}$$

Knapp and Carter [1976] determined a maximum likelihood (ML) estimator, that minimised $\text{Var}\left[\widehat{D}\right]$. This used a frequency weighting of

$$W(f) = \frac{|\gamma_{12}(f)|^2}{|S_{p_1 p_2}(f)|\left[1 - |\gamma_{12}(f)|^2\right]}, \tag{4.15}$$

where the coherence at each frequency is defined as

$$\gamma_{12}(f) \equiv \frac{S_{p_1 p_2}(f)}{\sqrt{S_{p_1}(f)\, S_{p_2}(f)}} = \frac{E\left[P_1(f)P_2{}^*(f)\right]}{\sqrt{E\left[|P_1(f)|^2\right] E\left[|P_2(f)|^2\right]}}. \tag{4.16}$$

The coherence is unity if $n_1 = n_2 = 0$ and goes to zero as the noise level increases. The frequency dependent weighting $W(f)$ increases the weighting of frequencies with higher coherence or lower noise.

The Cramér-Rao lower bound (CRLB) of the variance of the estimated delay $\widehat{D}$ is given by [Knapp and Carter, 1976]

$$\text{Var}\left[\widehat{D}\right] \geq \left[2T \int_0^\infty (2\pi f)^2 \frac{|\gamma_{12}(f)|^2}{\left[1 - |\gamma_{12}(f)|^2\right]} \, df\right]^{-1}. \tag{4.17}$$

Knapp and Carter [1976] showed that an estimator using the maximum likelihood frequency weighting (4.15) achieved this lower bound. An ordinary cross-correlation with $W(f) = 1$ will not achieve the CRLB unless the signal and noise power spectra match each other. For $W(f) = 1$, the delay variance is

$$\text{Var}\left[\widehat{D}\right] = \frac{\int_{-\infty}^\infty (2\pi f)^2 \left[S_{p_1 p_1}(f)\, S_{p_2 p_2}(f) - |S_{p_1 p_2}(f)|^2\right] \, df}{T\left[\int_{-\infty}^\infty (2\pi f)^2 |S_{p_1 p_2}(f)| \, df\right]^2}. \tag{4.18}$$

Carter et al. [1973] showed for $S_{n_1 n_1} = S_{n_2 n_2} = S_{nn}$ the coherence is related to the power spectra by

$$|\gamma_{12}(f)|^2 = \frac{S_{ee}{}^2(f)}{\left[S_{ee}(f) + S_{nn}(f)\right]^2}. \tag{4.19}$$

Thus the ratio of the signal and noise power spectra at frequency f is given by

$$\chi(f) \equiv \frac{S_{ee}(f)}{S_{nn}(f)} = \frac{|\gamma_{12}(f)|}{1 - |\gamma_{12}(f)|} \tag{4.20}$$

and thus

$$\frac{|\gamma_{12}(f)|^2}{1 - |\gamma_{12}(f)|^2} = \frac{1}{2}\left(\frac{1}{\chi(f)} + \frac{1}{2\chi^2(f)}\right)^{-1}. \tag{4.21}$$

The CRLB (4.17) can be determined from the power spectra using

$$\text{Var}\left[\widehat{D}\right] \geq \left[T \int_0^\infty (2\pi f)^2 \left(\frac{1}{\chi(f)} + \frac{1}{2\chi^2(f)}\right)^{-1} df\right]^{-1}. \tag{4.22}$$

If the signal and noise have an equal power spectra ratio, and are band-limited so that $\chi(f) = \text{SNR}$ for $f_1 < f < f_2$, then the CRLB becomes

$$\text{Var}\left[\widehat{D}\right] \geq \left[(2\pi)^2 T \left(\frac{1}{\text{SNR}} + \frac{1}{2\text{SNR}^2}\right)^{-1} \int_0^\infty f^2 \, df\right]^{-1} \tag{4.23}$$

$$\geq \frac{1}{(2\pi)^2} \cdot \frac{1}{T} \cdot \frac{3}{f_2^3 - f_1^3} \left[\frac{1}{\text{SNR}} + \frac{1}{2\text{SNR}^2}\right]. \tag{4.24}$$

If the power spectra have bandwidth $B = f_2 - f_1$, centred at f_0, then the CRLB becomes

$$\text{Var}\left[\widehat{D}\right] \geq \frac{1}{(2\pi f_0)^2} \cdot \frac{1}{TB} \cdot \frac{1}{1 + \frac{B^2}{12f_0^2}} \cdot \left[\frac{1}{\text{SNR}} + \frac{1}{2\text{SNR}^2}\right]. \tag{4.25}$$

This is similar to that developed by Quazi [Quazi, 1981] for passive sonar, who used the approximation

$$\frac{1}{\text{SNR}} + \frac{1}{2\text{SNR}^2} \approx \begin{cases} \frac{1}{\text{SNR}} & \text{for high SNR,} \\ \frac{1}{2\text{SNR}^2} & \text{for low SNR.} \end{cases}$$

4.1.3 Complex correlation

Consider $p_k(t)$ as a band-limited signal at centre frequency f_0, bandwidth B. It is common to deal with a complex baseband version of the signal $\widetilde{p}_k(t)$. A correlation of the complex signals may be performed in time or frequency using

$$\widehat{R}_{\widetilde{p}_1 \widetilde{p}_2}(\tau) = \frac{1}{T} \int_\tau^T \widetilde{p}_1(t) \widetilde{p}_2^*(t + \tau) \, dt \tag{4.26}$$

$$= \mathcal{F}_{f \to \tau}^{-1} \left\{ \widetilde{P}_1(f) \widetilde{P}_2^*(f) \right\} \tag{4.27}$$

where $\widetilde{p}_k(t) \longleftrightarrow \widetilde{P}_k(f)$. If the noise and signal are uncorrelated and cover the same bandwidth, the cross PSD is

$$S_{\widetilde{p}_1 \widetilde{p}_2}(f) = \frac{\alpha}{T} |E(f)|^2 \exp\left(j2\pi(f + f_0)D\right) + S_{n_1 n_2}(f) \tag{4.28}$$

and the cross correlation is the inverse Fourier transform;

$$R_{\widetilde{p}_1 \widetilde{p}_2}(\tau) = \alpha \, R_{ee}(\tau - D) \exp\left(j2\pi f_0 D\right) + R_{n_1 n_2}(\tau). \tag{4.29}$$

Thus a time delay causes a delay in the peak of the baseband correlation and a phase shift proportional to the delay.

Magnitude of complex correlation

For a coarse estimate of the time delay, one can measure the shift of the peak correlation

$$\widehat{D_c} = \arg\max_{\tau} \left\{ \left| \widehat{R}_{\widetilde{p}_1 \widetilde{p}_2}(\tau) \right| \right\}. \tag{4.30}$$

The sampling of τ can be increased via interpolation. By ignoring the correlation phase, information using the carrier is lost. The CRLB can be derived by setting the centre frequency f_0 to zero in (4.25), giving

$$\mathrm{Var}\left[\widehat{D_c}\right] \geq \frac{1}{4\pi^2} \cdot \frac{1}{TB} \cdot \frac{12}{B^2} \cdot \left[\frac{1}{\mathrm{SNR}^2} + \frac{1}{2\mathrm{SNR}^2} \right]. \tag{4.31}$$

Phase of complex correlation

The phase of the complex correlation gives a more accurate estimate of delay, but is ambiguous due to possible 2π phase wrapping. The ambiguity can be resolved by choosing the result closest to the *coarse* delay estimated from the correlation envelope. A *fine* delay can then be estimated using

$$\widehat{D_f} = \frac{\angle \left\{ \widehat{R}_{\widetilde{p}_1 \widetilde{p}_2}(\widehat{D_c}) \right\}}{2\pi f_0} + \frac{k}{f_0}, \tag{4.32}$$

where

$$k = \arg\min_{k} \left\{ \left| \widehat{D_c} - \widehat{D_f} \right| \right\}. \tag{4.33}$$

This will resolve any ambiguity if the error on $\widehat{D_c}$ is less than $1/(2f_0)$. This approach is sometimes referred to as a quasi-narrow-band approach [Shippey et al., 1998] as the phase only is used for the final measurement. This same technique is used in DPCA for SAS motion estimation [Bellettini and Pinto, 2002].

The CRLB for an unambiguous narrow-band correlation is derived by setting the bandwidth B to zero in (4.25) (this does not apply to the time-bandwidth product term), giving

$$\mathrm{Var}\left[\widehat{D}_{\mathrm{nb}}\right] \geq \frac{1}{(2\pi f_0)^2} \cdot \frac{1}{TB} \cdot \left[\frac{1}{\mathrm{SNR}^2} + \frac{1}{2\mathrm{SNR}^2} \right]. \tag{4.34}$$

This is the same as the CRLB developed by Bellettini and Pinto for the timing estimate

for DPCA [Bellettini and Pinto, 2002], assuming a statistically homogeneous reverber-
ation consisting of a Gaussian random process with a flat power spectral density. For a
narrow-band (high Q) system, the loss in accuracy over a full real correlation is minor.

4.1.4 Narrow-band (phase only) correlation

A time delay results in both a time-shift and phase shift of a full complex correlation
(4.29). Thus the phase shift at zero lag gives a direct estimate of the time delay
without the need for a full complex cross correlation at all lags. A narrow-band time
delay estimate can be formed via

$$\widehat{D}_{\text{nb}} = \frac{1}{2\pi f_0} \angle \left\{ \int \widetilde{p_1}(t)\widetilde{p_2}^*(t)\,\mathrm{d}t \right\} \tag{4.35}$$

$$= \frac{1}{2\pi f_0} \angle \left\{ \int \widetilde{P_1}(f)\widetilde{P_2}^*(f)\,\mathrm{d}f \right\}. \tag{4.36}$$

Narrow-band correlation has major computational savings over a full correlation. How-
ever, it assumes the autocorrelation of the signal R_{ee} is real only, which it is if the base-
band signals are symmetric about zero frequency. If not, the estimate will be biased
by

$$\widehat{D}_{\text{bias}} = \frac{1}{2\pi f_0} \angle \left\{ R_{\text{ee}}(0) \right\}. \tag{4.37}$$

In practice, it is difficult to ensure the signal spectra are symmetric. However, this will
only add a constant phase error, which means it is still useful to estimate a differential
phase error.

4.1.5 Noncoherent correlation

Instead of correlating the complex signals, it is possible to correlate the magnitude of
the signals using

$$\widehat{R}_{|p_1||p_2|}(\tau) = \frac{1}{T} \int_\tau^T |\widetilde{p_1}(t)|\,|\widetilde{p_2}(t+\tau)|\,\mathrm{d}t \tag{4.38}$$

and form a time delay estimate using

$$\widehat{D}_{\text{env}} = \arg \max_\tau \left\{ \widehat{R}_{|p_1||p_2|}(\tau) \right\}. \tag{4.39}$$

The information contained in the phase of the signal is lost, with a corresponding
reduction in accuracy. Although the result is less accurate than other methods, it is
unambiguous and useful for estimating the coarse time shift in a signal.

4.2 SPOTLIGHT FOURIER PHASE ERROR ESTIMATION

A similar problem exists in estimating the Fourier phase error of an spotlight image aberrated by a slant-range platform motion of $w(v)$. The baseband, range-compressed signal history is sampled so that $G_x[v] = G_x(v\Delta v)$ where v is the azimuth index and Δv the azimuth spacing.[2] The phase error $\phi_e(v) = \frac{4\pi f_0}{c} w(v)$ is independent of range, thus the aberrated signal history at a single range-bin x can be modelled by

$$G_x[v] = \overline{G_x}[v] \exp\left(j\phi_e[v]\right), \tag{4.40}$$

where $\overline{G_x}[v]$ is the unaberrated (ideal) range-compressed signal history. The image is formed by inverse Fourier transforming in the azimuth direction;

$$g_x[y] = \mathcal{F}_{v\to y}^{-1}\left\{G_x[v]\right\}. \tag{4.41}$$

It is useful to look at the components making up the phase of the signal $G_x[v]$. A common approach considers only the strongest target in the range-bin separately from other scatterers [3] [Eichel and Jakowatz, 1989; Jakowatz and Wahl, 1993; Ye *et al.*, 1999]. Suppose the strongest scatterer in range-bin x is at along-track position y_0, initial phase ϵ_x and echo magnitude a_x. The signal at this range-bin can be modelled as a single point target plus clutter, giving [4]

$$G_x[v] = a_x \exp\left[j(\epsilon_x + \omega_x v)\right] \exp\left(j\phi_e[v]\right) + N_x[v], \tag{4.42}$$

where $\omega_x = 2\pi y_0 \Delta v$. The signal to clutter ratio (SCR), defined as

$$\mathrm{SCR} \equiv \frac{a_x^2}{\mathrm{E}\left[|N_x|^2\right]}, \tag{4.43}$$

determines the variance of the phase of the signal in this range-bin. This determines the ultimate limit in accuracy of any estimate of phase error $\phi_e[v]$. Fourier phase-error estimation techniques estimate the phase-error ϕ_e, which has the property of being constant with range. The phase of the noise $N_x[v]$ is a random-variable, varying with range and azimuth. In addition, the strongest scatterer gives a linear $\omega_x v$ and constant ϵ_x phase offset to each range-bin. Various techniques exist for estimating ϕ_e, by removing the bias $\epsilon_x + \omega_x v$ caused by the target and averaging over range. Shear average is the simplest, estimating the phase difference between adjacent range-bins, thus removing the effect of ϵ_x. However there remains an offset caused by ω_x, constant with azimuth, varying with range, that reduces the accuracy of the estimate. Phase

---

[2] All signals are complex-baseband versions $\widetilde{G_x}[v]$, the tilde will be omitted for clarity.

[3] This model will be generalised in Section 4.3.

[4] The Fourier phase error $\phi_e[v]$ affects the clutter $N_x[v]$ also, but as the phase of the clutter is considered random, this can be ignored.

gradient autofocus (PGA) removes the bias of ω_x by shifting the strongest scatterer to the centre of the image, In addition, PGA increases the SNR via spatial windowing.

It is possible to directly estimate phase, removing the need to integrate phase differences. However, the phase offset ϵ_x needs to be directly estimated at each range-bin, and care needs to be taken not to introduce errors through phase-wrapping. This method is discussed in Section 4.2.3. Sharpness maximisation, discussed in Chapter 6, is not affected by ϵ_x or ω_x and is only affected by noise.

4.2.1 Phase difference estimation (shear average)

It is difficult to directly estimate absolute Fourier phase as each range-bin is offset by a varying linear phase, making direct summation over range impossible. It is possible to estimate the *phase difference* between adjacent azimuth bins using the shear average

$$C_x[v] \equiv G_x[v+d]\, G_x^*[v] \tag{4.44}$$

$$= \overline{G_x[v+d]}\, \overline{G_x^*[v]} \exp\left[j\Delta\phi[v]\right], \tag{4.45}$$

where the phase difference $\Delta\phi[v] \equiv \phi_e[v+d] - \phi_e[v]$. The phase can then be estimated by integrating the estimated phase differences;

$$\widehat{\phi}[v+d] = \widehat{\phi}[v] + \Delta\widehat{\phi}[v] \tag{4.46}$$

Generally phase difference estimation is performed on adjacent azimuth bins ($d = 1$), as the method requires coherence between the two bins which is maximum for adjacent bins. This technique is used in a number of fields, known as Knox-Thompson speckle interferometry [Knox and Thompson, 1974] in astronomical imaging, spatial correlation autofocus for sparse radio arrays [Attia and Steinberg, 1989], or shear-average in spotlight SAR [Fienup, 1989]. There is no reason shear-average cannot be extended to further separated bins $d > 1$. This is discussed in Section 4.3.5.

It is interesting to note the similarity between the shear average product (4.44) and interferometry. Both estimate the phase of the product $\theta = \angle\{z_1 z_2^*\}$. For interferometry, z_1 and z_2 are the signals from different receivers or passes. For phase difference estimation, z_1 and z_2 are adjacent along-track echos or pings. Thus techniques of interferometry can be of use. Phase difference estimation has the additional property that the parameter of interest, the Fourier phase error ϕ_e, is constant with range. This allows averaging of the interferogram over range samples.

Modelling the signal as a single point target plus clutter (4.42), the shear average

(4.45) becomes

$$C_x[v] = (a_x \exp\left[j(\epsilon_x + \omega_x(v+d) + \phi_e[v+d])\right] + N_x[v+d])$$

$$(a_x \exp\left[-j(\epsilon_x + \omega_x v + \phi_e[v])\right] + N_x{}^*[v]) \tag{4.47}$$

$$= |a_x|^2 \exp\left[j(\Delta\phi[v] + \omega_x d)\right] + S'_{NN}[v], \tag{4.48}$$

assuming the noise is uncorrelated with the target and between azimuth bins. Comparing the product of adjacent azimuth bins (4.48) and the cross power spectra for time-delay estimation (4.28) it is apparent that estimating the phase difference is the same problem as time delay estimation, for $\Delta\phi[v] = 2\pi(f + f_0)D$ with a phase offset $\omega_x d$. Thus if $f_0 \gg B$, Fourier phase-difference estimation is the *narrow-band* equivalent of the time delay estimation problem with $\Delta\phi = 2\pi f_0 D$. The equivalent of frequency f in time delay estimation, is the range x. The equivalent of the signal spectra $P_1(f)$ is then a single ping, $G_x[v_1]$. A discrete version of the narrow-band correlation estimation (4.36) is given by

$$\Delta\widehat{\phi}[v] = \angle\left\{\sum_x G_x[v+d]\, G_x{}^*[v]\right\}, \tag{4.49}$$

which is the shear-average estimator. The shear-average estimator has been shown to give the maximum likelihood solution under the assumptions that the image consists of a single, central (zero phase) point target per range-bin, and the noise power is constant with range [Jakowatz and Wahl, 1993].

Results from time-delay estimation can be used in phase difference estimation problems. The variance can be related using

$$\text{Var}\left[\Delta\widehat{\phi}\right] = (2\pi f_0)^2 \text{Var}\left[\widehat{D}\right]. \tag{4.50}$$

The equivalent of the time-bandwidth product TB is the number of independent range samples[5], given by

$$N'_x = N_x \frac{B_c}{f_s}, \tag{4.51}$$

where N_x is the number of range samples, B_c is the pulse bandwidth and f_s the sampling frequency. The phase difference estimate equivalent of the narrow-band CRLB (4.34) for a constant SCR becomes

$$\text{Var}\left[\Delta\widehat{\phi}\right] \geq \frac{1}{N'_x}\left[\frac{1}{\text{SCR}} + \frac{1}{2\text{SCR}^2}\right]. \tag{4.52}$$

In time-delay estimation, a normal correlation is optimum if the signal and noise have an equal ratio of power spectra with regard to *frequency*. A generalised correlation Section 4.1.2 is optimum if the spectra are not equal. For the sheared-product phase

---

[5]The pulse duration $T_c = \frac{N_x}{f_s}$.

difference estimator to be optimum, the signal and clutter power of the *signal* needs
to be constant with *range*. This is not a reasonable assumption. A generalised model
not requiring this assumption, the equivalent of generalised correlation for time-delay
estimation, is developed in Section 4.3.

4.2.2 PGA algorithm

Each range-bin has a different linear offset $\omega_x v$, due to the strongest scatterer (4.42)(4.48).
The PGA algorithm removes this offset by circularly shifting the strongest scatterer in
each range-bin to the centre of the image [Eichel *et al.*, 1989]. This step improves the
accuracy of the phase estimation [Wahl *et al.*, 1994a] [Jakowatz *et al.*, 1996]. This is
performed in the image domain to produce $g_x^{(cs)}[y]$.

PGA performs *spatial* windowing in along track to increase the SNR. The windowed
image is found by

$$g_x^{(w)}[y] = w[y] \, g_x^{(cs)}[y] \tag{4.53}$$

This assumes the spatial extent of the blurred target is within the window. Thus the
window includes all the signal and excludes clutter, increasing the SNR of the image.
This results in a more accurate phase estimation [Jakowatz *et al.*, 1996]. However, it
also reduces the order of the estimated phase error. A discussion on the best method
for windowing can be found in [Warner *et al.*, 2000] [Callow, 2003].

Following windowing, the phase-gradient is estimated in the same way as shear-
average discussed in the previous section. Thus PGA can be considered as prefiltering
the signal history $G_x[v]$ to obtain

$$G_x^{(w)}[v] = W[v] \odot G_x[v] \exp\left(2\pi \Delta v y_0 v\right), \tag{4.54}$$

where y_0 is the along-track position of the strongest scatterer in range-bin x and
$W[v] \longleftrightarrow w[y]$.

Originally, PGA used a different phase estimation kernel, the differentiation kernel
[Eichel and Jakowatz, 1989]. The performance of this kernel was demonstrated to be
inferior to the shear-average method (4.49) [Jakowatz and Wahl, 1993]. The terms
phase-gradient and phase-difference are often used interchangeably.

The windowing step is most effective for point-like targets. For areas of shadow or
textured clutter, such as trees in a SAR image or sand-ripples in a SAS image, spatial
windowing will not be effective since the signal is not spatially isolated around the
brightest point. For shadows, windowing will decrease the SCR. This can be seen when
a shadow is centre-shifted in Fig. 4.9(c). After centre shifting, the centre of the image
contains only speckle.

4.2.3 Direct phase estimation

A different phase estimation method was suggested by Ye et al. [Ye *et al.*, 1999]. For a number of phase estimates ϕ_x, phase gradient estimation finds the phase of a sum of the estimates using

$$\Delta\widehat{\phi} = \angle\left\{\sum_x a_x \exp\left(j\Delta\phi_x\right)\right\},$$

whereas direct phase estimation directly sums the phase estimates using

$$\widehat{\phi} = \frac{\sum_x b_x \phi_x}{\sum_x b_x}.$$

For this method to work, several challenges need to be overcome. Firstly, as phase values are being summed, phase wrapping becomes a problem. Any constant phase offset ϵ_x will upset the phase addition process. The correct weights b_x must also be determined. The method suggested by Ye et al. [Ye *et al.*, 1999], termed weighted least-squares (WLS) phase estimation, used a local phase unwrapping (LPU) technique to prevent errors from phase wrapping and directly estimate the phase offsets ϵ_x.[6] The best (maximum likelihood) weightings for the phase estimates are the inverse variance of the phase estimates σ_x^{-2}. The method to estimate σ_x^{-2} involved modelling the signal as a single prominent point plus clutter, then estimating the SNR from the relationship between the mean magnitude and mean intensity of the signal. An estimate of the variance is then formed from the estimated SNR. An improved method for estimating σ_x^{-2} involving a measure of the signal coherence is discussed later, in Section 4.3.1. One important difference is that phase-gradient techniques estimate the phase-gradient $\Delta\phi[v] \equiv \phi[v+d] - \phi[v]$. If the error in $\widehat{\phi}[v+d]$ and $\widehat{\phi}[v]$ are independent, the variance of the phase-gradient is twice the variance of the direct-phase estimate, i.e.

$$\mathrm{Var}\left[\Delta\widehat{\phi}\right] = 2\,\mathrm{Var}\left[\widehat{\phi}\right]. \qquad (4.55)$$

4.3 GENERALISED PHASE DIFFERENCE ESTIMATION

The model used in the previous section, assumed the signal comprised of a single point target plus clutter with constant statistics with range. The performance of the shear-average estimator will suffer if this does not apply. This section generalises this model and develops the maximum likelihood estimator weighting for each range-bin, by measuring the average SCR of each range-bin. This is analogous to the generalised correlation method for time delay estimation [Knapp and Carter, 1976]. Whereas generalised time-delay estimation applies a varying frequency weighting, generalised phase difference estimation applies a range-weighting to the correlation between pings.

---

[6]Linear phase offsets were first removed using a circular shifting procedure as in PGA.

One can consider that a target is the portion of the signal that remains constant ping to ping, while clutter is the portion that does not. The more coherent the signal, the lower the phase variance of the signal and the more accurate the phase difference estimate. It is possible to use a measure of the coherence between azimuth bins as a measure of the average SCR and thus the accuracy of the phase estimate. Coherence is also widely used as a measure of the accuracy of phase difference estimates in time delay estimation [Knapp and Carter, 1976] and interferometry [Rodriguez and Martin, 1992; Just and Bamler, 1994].

4.3.1 Estimating phase variance

The phase difference between adjacent azimuth bins $\theta_x[v]$ can be measured with

$$\theta_x[v] = \angle \left\{ G_x[v+1] G_x{}^*[v] \right\}. \tag{4.56}$$

To weight the phase difference estimate from each range-bin appropriately, it is important to model the variance of this phase difference measurement. This depends on the coherence of the adjacent azimuth bins, which in turn depends on the signal to clutter ratio. This section develops a model of the phase difference variance for two different signal models.

For estimating the time-delay between two signals x_1 and x_2, Knapp and Carter [1976] determined the maximum-likelihood weighting for minimising the estimate variance. This involved the complex coherence function, or normalised cross-spectral density defined as

$$\gamma_{x_1 x_2}(f) \equiv \frac{S_{x_1 x_2}(f)}{\sqrt{S_{x_1}(f) \, S_{x_2}(f)}} \tag{4.57}$$

Various methods exist to estimate the coherence function [Carter et al., 1973]. Some form of ensemble average is required. The average coherence of a single range bin can be estimated by averaging over azimuth bins using

$$\widehat{\gamma_x} = \frac{\sum_v G_x[v+1] G_x{}^*[v]}{\sqrt{\sum_v |G_x[v+1]|^2 \sum_v |G_x[v]|^2}}. \tag{4.58}$$

More sophisticated measures of coherence are available. For a discussion on estimating coherence in SAR images, see [Touzi et al., 1999].

Constant point target

Consider a perfect point target, with a constant signal α_x and phase $\omega_x v$, added to white circular Gaussian speckle noise with mean intensity μ_U. The signal may be modelled as

$$\overline{G_x}[v] = \alpha_x \exp(-j\omega_x v) + N_x[v], \tag{4.59}$$

with $N_x[v]$ a white circular Gaussian random variable, with $\mathrm{E}\left[|N_x|^2\right] = \mu_U$. The signal and noise are uncorrelated so that

$$\mathrm{E}\left[|G_x|^2\right] = |\alpha_x|^2 + \mu_U. \tag{4.60}$$

The signal to clutter ratio (SCR) of the signal at range-bin x is given by

$$\mathrm{SCR}_x = \frac{\alpha_x^2}{\mathrm{E}\left[|N_x|^2\right]} = \frac{\alpha^2}{\mu_U}. \tag{4.61}$$

The image is the inverse Fourier transform of the signal, giving

$$\overline{g_x}[y] = N_y \alpha_x \delta[y - y_0] + n_x[y], \tag{4.62}$$

where $y_0 = \omega_x/(2\pi\Delta v)$, N_y is the number of azimuth samples. The image clutter n_x is white circular Gaussian speckle noise with $\mathrm{E}\left[|n_x|^2\right] = N_y \mu_U$ and $y_0 = \omega_x/(2\pi\Delta v)$. The ratio of the intensity of the point target to the average intensity of the clutter is $N_y \alpha_x^2/\mu_U$, so is dependent on the SCR and number of pings on the target.

With a phase error $\phi_e[v]$, the aberrated signal-history is

$$G_x[v] = \alpha_x \exp(-j\omega_x v) \exp(j\phi_e[v]) + N'_x[v], \tag{4.63}$$

where $N'_x = N_x \exp(j\phi_e)$ has the same statistical properties as N_x. The signal coherence is given by

$$\gamma_x \equiv \frac{\mathrm{E}\left[G_x[v+1]G_x{}^*[v]\right]}{\sqrt{\mathrm{E}\left[|G_x[v+1]|^2\right]\mathrm{E}\left[|G_x[v]|^2\right]}} \tag{4.64}$$

$$= \frac{|\alpha_x|^2 \exp\left[j(\Delta\phi_e - \omega_x)\right]}{|\alpha_x|^2 + \mu_U}$$

$$= \frac{\exp\left[j(\Delta\phi_e - \omega_x)\right]}{1 + \mathrm{SCR}_x^{-1}} \tag{4.65}$$

Thus, $\angle\{\gamma_x\} = \Delta\phi_e - \omega_x$ and $|\gamma_x| = \frac{1}{1+\mathrm{SCR}_x^{-1}}$. The latter is a well known result for TDE [Carter *et al.*, 1973] and additive (thermal) noise in interferometric SAR [Rodriguez and Martin, 1992; Just and Bamler, 1994]. Thus the SCR can be derived from the coherence using

$$\mathrm{SCR}_x = \frac{|\gamma_x|}{1 - |\gamma_x|}. \tag{4.66}$$

Fig. 4.1 compares the known SCR for a point target to that measured using the coherence (4.66), showing a good match.

From (4.63), the signal history is a constant background (the signal), plus a speckle pattern (the noise). The distribution of the resultant has been developed by Goodman
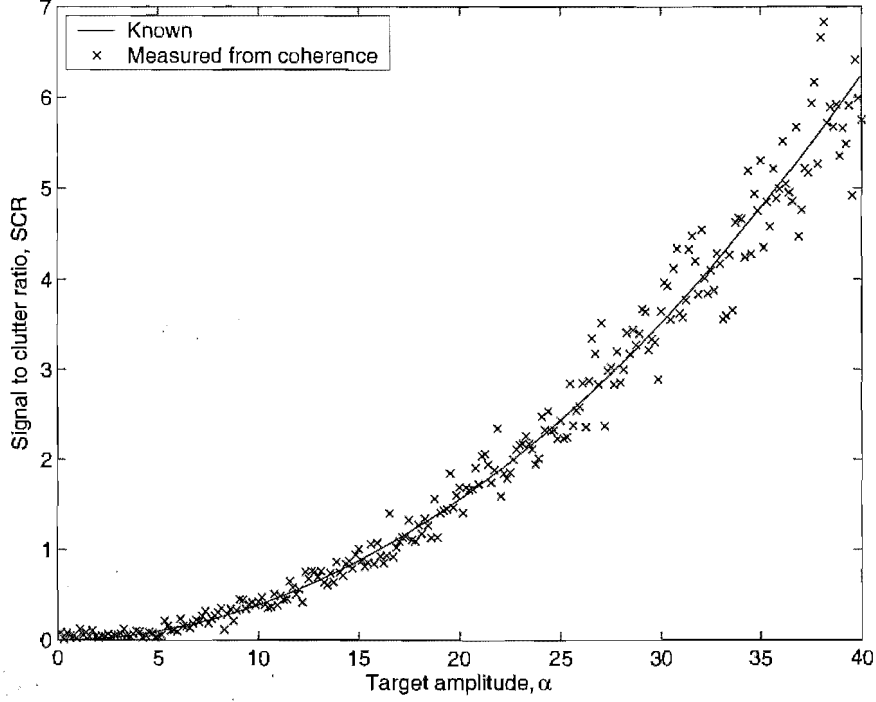
Figure 4.1 The signal to clutter ratio of a point target amplitude α in clutter with mean intensity $\mu = 1$, for $Ny = 256$. The known $SCR = \frac{\alpha^2}{N_y \mu_U}$ is compared to that measured using the coherence $\frac{|\gamma_x|}{1-|\gamma_x|}$.

[Goodman, 1975]. The probability density function of the intensity $I_G \equiv |G_x|^2$ is

$$\Pr[I_G] = \begin{cases} \frac{1}{\mu_U} \exp\left(-\frac{I_G + \alpha_x^2}{\mu_U}\right) I_0 \left(2\frac{\sqrt{I_G \alpha_x^2}}{\mu_U}\right), & I_G \geq 0 \\ 0 & \text{otherwise,} \end{cases} \qquad (4.67)$$

where $I_0()$ is a modified Bessel function of the first kind, zero order. The probability distribution of the phase $\xi_x \equiv \angle\{G_x\}$, for $\mathrm{E}[\xi_x] = 0$ is given by

$$\Pr[\xi_x] = \frac{e^{-\mathrm{SCR}_c}}{2\pi} + \sqrt{\frac{\mathrm{SCR}_x}{\pi}} \cos \xi_x \exp\left(-\mathrm{SCR}_x \sin^2 \xi_x\right) \Omega\left(\sqrt{2\mathrm{SCR}_x} \cos \xi_x\right)$$

$$\text{for } -\pi \leq \xi_x \leq \pi, \qquad (4.68)$$

where

$$\Omega(b) \equiv \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{b} \exp\left(-\frac{y^2}{2}\right) dy. \qquad (4.69)$$

When the target is weak ($\mathrm{SCR} \to 0$), the phase becomes uniformly distributed ($\Pr[\xi_x] \to \frac{1}{2\pi}$). When the target is strong ($\mathrm{SCR} \gg 1$), the variance of the phase reduces and $\Pr[\xi_x]$ is approximately Gaussian [Goodman, 1975].

The variance of the phase ξ_x can be approximated by [7]

$$\hat{\sigma}_{\xi_x}^2 \approx \frac{1}{2}\left(\frac{1}{\text{SCR}_x} + \frac{1}{2\text{SCR}_x^2}\right). \tag{4.70}$$

Fig. 4.2 compares the above model to some simulated data. There is a good match for SCR > 1. Below this the phase wraps, whichreduces the measured variance. The model assumes unwrapped phase and $\hat{\sigma}_\xi^2 \to \infty$ as SCR $\to 0$. Substituting the coherence

(a)   (b)
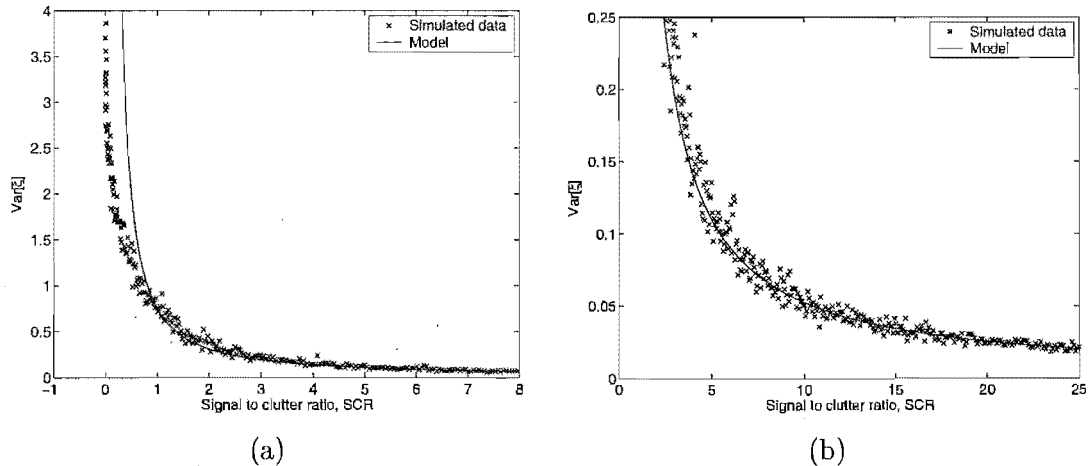
Figure 4.2   Variance of phase ξ of signal history of point target with additive, white, circular Gaussian noise. Simulated data is compared to model (4.70), $\hat{\sigma}_\xi^2 = \frac{1}{2}\left(\frac{1}{\text{SNR}} + \frac{1}{2\text{SNR}^2}\right)$. Simulated data is wrapped $[-\pi, \pi]$, so the max variance, for uniformly distributed phase, is $\pi^2/3 = 3.29$. (a) Small SNR. (b) Large SNR.

measure of the SCR (4.66) into the phase variance estimate (4.70), gives an estimate of the phase variance based on the signal coherence,

$$\text{Var}\,[\xi_x] \approx \frac{1 - |\widehat{\gamma_x}|^2}{4\,|\widehat{\gamma_x}|^2}. \tag{4.71}$$

The variance of the phase-difference between azimuth bins, $\theta_x[v] = \xi_x[v+1] - \xi_x[v]$ is given by $\text{Var}\,[\theta_x[v]] = \text{Var}\,[\xi_x[v+1]] + \text{Var}\,[\xi_x[v+1]]$ if ξ is independent between azimuth bins. Thus

$$\text{Var}\,[\theta_x] = 2\,\text{Var}\,[\xi_x] \tag{4.72}$$

$$\approx \frac{1 - |\widehat{\gamma_x}|^2}{2\,|\widehat{\gamma_x}|^2}. \tag{4.73}$$

This matches results in TDE [Knapp and Carter, 1976] and the limiting case of multi-look SAR interferometry [Rodriguez and Martin, 1992]. Fig. 4.3 compares the variance of the measured phase difference of a simulated point target with additive white circular

---

[7]Goodman [Goodman, 1975] used the approximation $\text{Var}\,[\xi] \approx \frac{1}{2\text{SNR}}$ but (4.70) gives a better approximation at lower SNR.

Gaussian noise, with the variance model of (4.73). There is a good match for SCR > 1. Below this, phase wrapping reduces the measured variance.
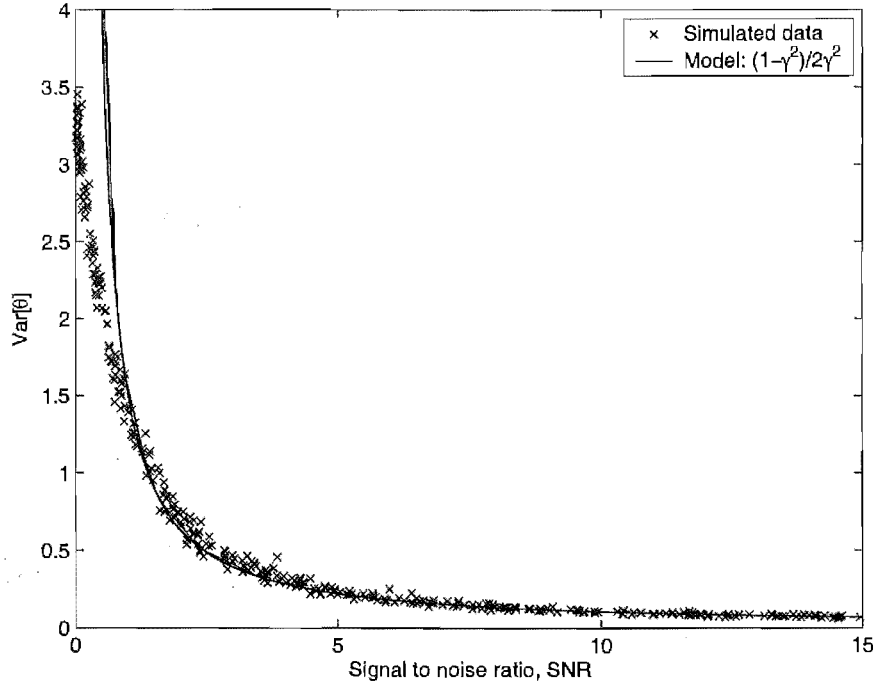
Figure 4.3 Variance of phase difference θ, between adjacent azimuth bins for perfect point target with additive white circular Gaussian noise. Simulated data is compared to model of (4.73) $\sigma_\theta^2 = \frac{1-\gamma^2}{2\gamma^2}$.

Speckle target

Not all targets are well modelled as a constant point source. The measured reflectivity of a rough surface in coherent imagery (SAR or SAS) is a random variable due to speckle (see Chapter 3). An extended rough target can be modelled as the mean scene reflectivity $e_x[y]$, multiplied by a complex, stationary, white, circular Gaussian random variable $n_x[y]$ (the speckle) filtered by $h_x[y]$ representing the imaging process;

$$g_x[y] = \Big(n_x[y]e_x[y]\Big) \odot h_x[y]. \tag{4.74}$$

A model of the signal history is the Fourier transform of the scene, giving

$$G_x[v] = \Big(N_x[v] \odot E_x[v]\Big)H_x[v], \tag{4.75}$$

with $E_x[v] \longleftrightarrow e_x[y]$, $N_x[v] \longleftrightarrow n_x[y]$ and $H_x[v] \longleftrightarrow h_x[y]$. The signal clutter $N_x[v]$ is also a complex, stationary, white, circular Gaussian random variable. At an azimuth position separated by d, the signal is then

$$G_x[v+d] = \Big(N_x[v] \odot E_x[v+d]\Big)H_x[v+d]. \tag{4.76}$$

Generally, the imaging system $H_x[v]$ can be considered broadband (slowly varying) and for a small separation $(d = 1)$, $H_x[v+1] \approx H_x[v]$. The statistics of the sheared product have been developed for this system model, as it is of interest in interferometry [Just and Bamler, 1994; Lee et al., 1994]. Of particular interest is the phase of the cross-product

$$\theta_x[v] = \angle \{G_x[v + d]G_x{}^*[v]\}. \tag{4.77}$$

The probability density function $f_\theta(\theta)$ of the phase difference for a single look (single range-bin) is given by [Just and Bamler, 1994; Lee et al., 1994]

$$f_\theta(\theta) = \frac{1 - |\gamma|^2}{2\pi \left[1 - |\gamma|^2 \cos^2(\theta - \theta_0)\right]} \left\{ 1 + \frac{|\gamma| \cos(\theta - \theta_0) \cos^{-1}[-|\gamma| \cos(\theta - \theta_0)]}{\left[1 - |\gamma|^2 \cos^2(\theta - \theta_0)\right]^{0.5}} \right\}, \tag{4.78}$$

where $\theta_0 = \mathbf{E}[\theta]$. Thus the statistics of the phase difference between two pings depends only on the coherence between them. Fig. 4.4(a) shows the pdf $f_\theta(\theta)$ for several values of coherence γ. If the signals have no correlation, $\gamma = 0$ and the phase pdf is uniformly distributed $[-\pi, \pi]$. For full correlation, $\gamma = 1$ and the pdf approaches a delta function at θ_0. The variance can be calculated from the pdf using

$$\mathrm{Var}\,[\theta] = \int_{-\pi}^{\pi} \theta^2 \, f_\theta(\theta + \theta_0) \, \mathrm{d}\theta. \tag{4.79}$$

Fig. 4.4(b) shows the variance of the phase difference of a single range-bin versus coherence. This is compared to the model for a constant target (4.73). These do not match well, showing the constant-target model does not give a good approximation of phase-variance for a single range-bin of a speckle target. Fig. 4.5 shows a simulated speckle target. The phase variance estimate obtained by integrating the pdf (4.79) is a close match to the measured phase variance. The constant-target model (4.73) is a poor match.

The variance of a single-look phase difference is normally large. A reduced variance estimate can be formed formed from using multiple looks. For phase difference estimation, each range-bin can be considered a separate look, with independent noise. If each look had the same coherence, the multi-look phase can be obtained by

$$\widehat{\phi}[v] = \angle \left\{ \frac{1}{N'_x} \sum_x G_x[v + 1]G_x{}^*[v] \right\}. \tag{4.80}$$

This is the shear-average phase difference estimator, described earlier (4.49). The multi-look phase-difference distribution has been given by Lee et al. [Lee et al., 1994]
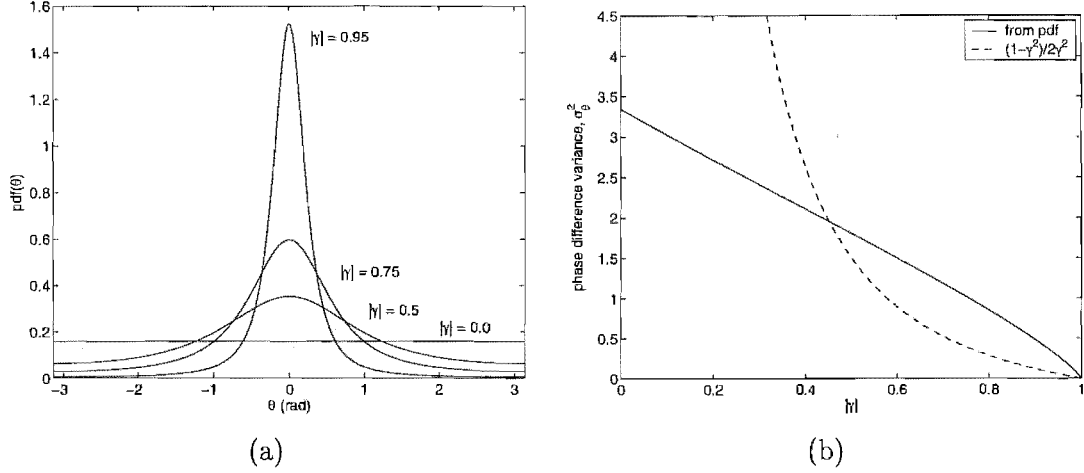
Figure 4.4 Statistics of phase difference θ between a single look of two signals, with $E[\theta] = 0$ and coherence γ. For $|\gamma| = 0$, phase is uniformly distributed $[-\pi, \pi]$ and $\sigma_\theta^2 = \pi^2/3$. For $|\gamma| = 1$, phase pdf converges toward a delta function and $\sigma_\theta^2 = 0$. (a) Probability density function of θ from (4.78). (b) Variance of theta versus coherence γ from integrating the pdf (4.79), compared to model $\sigma_\theta^2 = (1 - \gamma^2)/2\gamma^2$ from (4.73).

for N_x' independent looks as

$$f_{\widehat{\phi}}(\widehat{\phi}, N_x') = \frac{\Gamma(N_x' + \frac{1}{2})(1 - |\gamma|^2)^{N_x'} |\gamma| \cos(\widehat{\phi} - \phi_0)}{2\sqrt{\pi}\Gamma(N_x') \left(1 - |\gamma|^2 \cos^2(\widehat{\phi} - \phi_0)\right)^{N_x' + \frac{1}{2}}}$$

$$+ \frac{(1 - |\gamma|^2)^{N_x'}}{2\pi} F_1\left(N_x', 1; \frac{1}{2}; |\gamma|^2 \cos^2(\widehat{\phi} - \phi_0)\right), \quad (4.81)$$

where $F_1\left(N_x', 1; \frac{1}{2}; \beta\right)$ is a Gauss hypergeometric function. Note for $N_x' = 1$, (4.81) becomes (4.78). Fig. 4.6 shows the phase-difference variance with varied number of looks and coherence. Rodriguez and Martin [Rodriguez and Martin, 1992] described the CRLB of the multi-look phase-difference variance as

$$\text{Var}\left[\widehat{\phi}\right] = \frac{1 - |\gamma|^2}{2N_x' |\gamma|^2}, \quad (4.82)$$

which was stated as a good approximation of variance for $N_x' > 4$. Fig. 4.6 compares this approximation to that obtained from integrating the pdf in (4.81). This shows it to be a fair approximation for a large number of looks N_x' and high coherence γ. It is a poor approximation for $N_x' < 8$ or $|\gamma| < 0.2$

The previous derivation assumes the coherence γ is constant over all looks. This is a reasonable assumption for multi-look interferometry, where each look is of the same scene, thus has the same SNR. For phase-difference estimation, phase estimates are averaged over range-bins. The scene, thus the coherence and phase variance, varies with each range-bin.
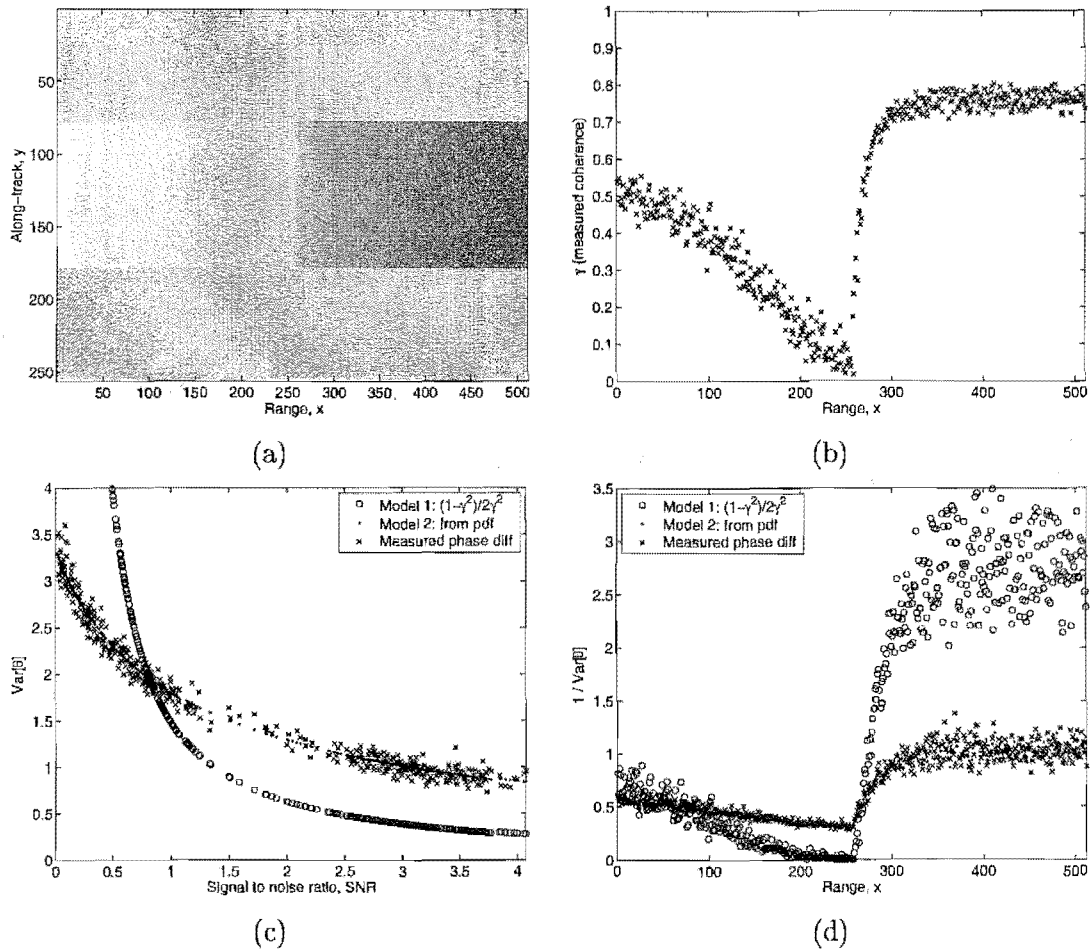
Figure 4.5   Simulation of speckle block target. (a) Simulated image intensity. Background has unity average intensity. Average target intensity varies from zero at $x = 1$, unity at $x = 256$ and 30 at $x = 512$. (b) Image coherence versus range. (c) Variance of phase difference θ versus SNR. Simulated data is compared to model (4.73) and from integrating the pdf (4.79). (c) Inverse variance of phase difference θ versus range. Simulated data is compared to model (4.73) and from integrating the pdf (4.79).
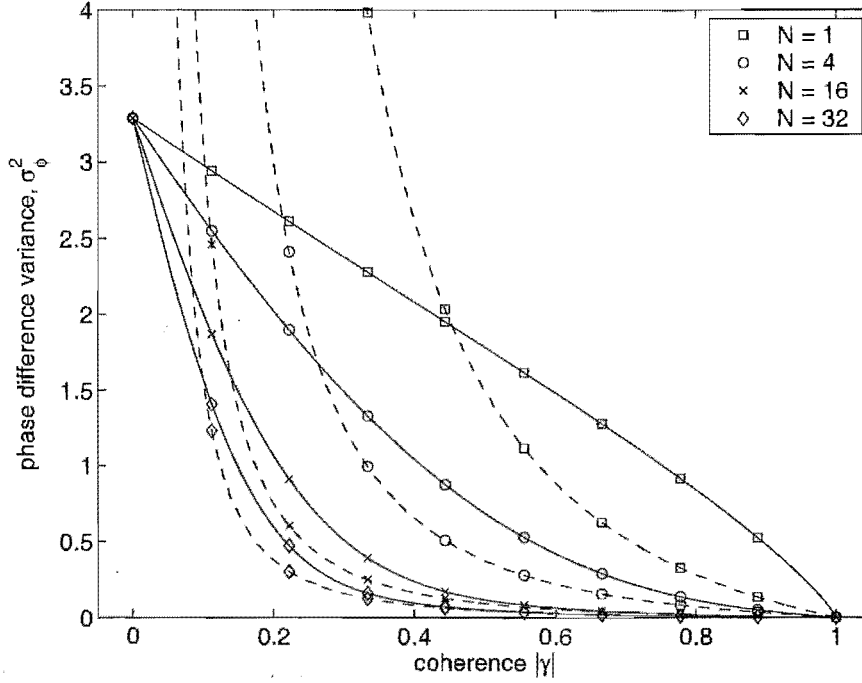
Figure 4.6 Variance of phase difference versus coherence γ with N looks. Solid line is model from integrating pdf in (4.81). Dashed line is approximate model in (4.82) $\mathrm{Var}\left[\widehat{\phi}\right] = \frac{1-|\gamma|^2}{2N_x'|\gamma|^2}$.

4.3.2 Maximum likelihood estimation

Consider N independent estimates of a parameter Y_n, each with a different variance σ_n^2. The maximum-likelihood estimate is given by [Beck and Arnold, 1977]

$$\widehat{Y}_{\mathrm{ML}} = \frac{\sum_n Y_n \sigma_n^{-2}}{\sum_n \sigma_n^{-2}}, \tag{4.83}$$

which has a variance of

$$\mathrm{Var}\left[\widehat{Y}_{\mathrm{ML}}\right] = \left[\sum_n \sigma_n^{-2}\right]^{-1}. \tag{4.84}$$

Thus the direct maximum-likelihood estimation of phase, or weighted least-squares estimate [Ye *et al.*, 1999] is given as

$$\widehat{\phi}_{WLS}[v] = \frac{\sum_x \widehat{\psi}_x[v]\sigma_\psi^{-2}[x]}{\sum_x \sigma_\psi^{-2}[x]} \tag{4.85}$$

where $\sigma_\psi^2[x]$ is the variance of the phase estimate $\widehat{\psi}_x[v]$ at range-bin x.

The *weighted phase difference* approach averages the phase difference θ between

azimuth bins, using

$$\Delta\widehat{\phi}_{WPD}[v] = \angle \left\{ \sum_x \exp\left(j\theta_x[v]\right) \widehat{\sigma}_\theta^{-2}[x] \left[\sum_x \widehat{\sigma}_\theta^{-2}[x] \right]^{-1} \right\} \qquad (4.86)$$

$$= \angle \left\{ \sum_x \widehat{\sigma}_\theta^{-2}[x] \frac{G_x[v+1]G_x{}^*[v]}{|G_x[v+1]G_x{}^*[v]|} \right\}. \qquad (4.87)$$

Using the variance estimate based on coherence (4.73) in the maximum-likelihood phase-gradient estimator (4.87) gives

$$\Delta\widehat{\phi}_{WPD}[v] = \angle \left\{ \sum_x \frac{2\,|\widehat{\gamma}_x|^2}{1 - |\widehat{\gamma}_x|^2} \frac{G_x[v+1]G_x{}^*[v]}{|G_x[v+1]G_x{}^*[v]|} \right\}, \qquad (4.88)$$

which is the same as time-delay generalised correlation (4.12) with the maximum-likelihood frequency weighting (4.15) from Knapp and Carter [Knapp and Carter, 1976]. Alternatively, the better phase variance approximation (4.79) for coherent images, can be used in the weighted phase gradient estimator (4.87) for the maximum likelihood phase estimator of coherent images.

This can be compared with the shear-average method, which weights each phase estimate by $|G_x[v+1]G_x{}^*[v]|$. Fig. 4.7 compares the mean weighting given to each range bin by shear-average to the estimated inverse variance by the two different models for the scene in Fig. 4.5(a). This is compared to the inverse variance of the measured phase difference. Shear-average weights phase estimates by signal energy, not signal coherence. Thus those range bins with higher coherence and lower energy, such as shadows, get a lower weighting. If speckle is brighter, for example at closer ranges, it will get a higher weighting in shear-average even though the phase estimate from it is completely random. Shear-average is not normalised by signal energy. The phase-variance model of (4.79) matches the measured phase variance closely, thus weighting each range-bin appropriately.

One advantage of shear-average is that each individual phase estimate in azimuth and range is weighted by $|G_x[v+1]G_x{}^*[v]|$, whereas for the weighted approach (4.87), each azimuth sample is weighted by the same average weighting for that range-bin. An improved approach would be to set the *average* weight of the range-bin to $\widehat{\sigma}_\theta^{-2}[x]$, but allow the weighting to vary in azimuth by using

$$\Delta\widehat{\phi}_{WPD}[v] = \angle \left\{ \sum_x \widehat{\sigma}_\theta^{-2}[x] \frac{G_x[v+1]G_x{}^*[v]}{\sum_v |G_x[v+1]G_x{}^*[v]|} \right\}. \qquad (4.89)$$

Table 4.1 shows that using this method (methods 2 and 4), leads to an improved phase estimation in all test images, compared to giving all estimates in the range bin the same weighting (methods 1 and 3).
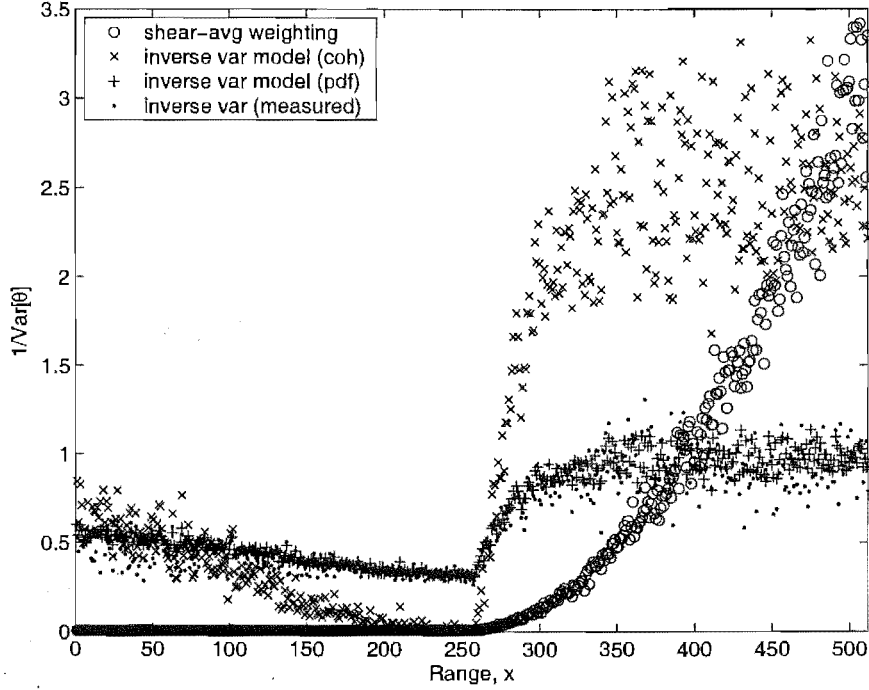
Figure 4.7 Mean weighting of shear average, $|G_x[v + 1]G_x{}^*[v]|$ compared to inverse variance and inverse variance models for each range-bin. Image is that shown in Fig. 4.5(a). Inverse variance of measured phase difference is compared to coherence model (4.73) and from integrating the pdf (4.79).

The total variance of the phase gradient estimate depends on the sum of the inverse variance at each range bin by

$$\text{Var}\left[\Delta\widehat{\phi}\right] = \left[\sum_x \sigma_\theta^{-2}\right]^{-1}. \tag{4.90}$$

using the approximation in (4.73), this becomes

$$\text{Var}\left[\Delta\widehat{\phi}\right] = \left[\sum_x \frac{2\,|\widehat{\gamma_x}|^2}{1 - |\widehat{\gamma_x}|^2}\right]^{-1}. \tag{4.91}$$

This compares directly to the CRLB of the variance of a time-delay estimate (4.17). If the SCR is equal for each range bin, substituting (4.66) into (4.91) gives

$$\text{Var}\left[\Delta\widehat{\phi}\right] = \frac{1}{N_x'}\left[\frac{1}{\text{SCR}} + \frac{1}{2\text{SCR}^2}\right]. \tag{4.92}$$

which meets the CRLB (4.52).

Adjustment for phase wrapping

Weighting each phase estimate by its inverse variance as in (4.89) would be ideal if the estimated phase variance was for unwrapped phase. Then as $\gamma \to 0$, the phase becomes

uniformly distributed, $\sigma_\theta^2 \to \infty$ and the weighting becomes zero. This is desirable, as there is no useful information in the phase if it is uniformly distributed. However, the phase variance model used in (4.79) is for wrapped phase. Using this model, as $\gamma \to 0$, the phase becomes uniformly distributed $[-\pi, \pi]$, $\sigma_\theta^2 \to \pi^2/3$ and the lowest weighting is $3/\pi^2 = 0.304$. Those range bins with no useful phase information still contribute to the overall phase estimate. A preferable weighting to use is thus

$$W[x] = \left| \widehat{\sigma}_\theta^{-2} - 0.304 \right| \tag{4.93}$$

to reduce the contribution from range-bins with low SCR. Table 4.1 shows that using this adjusted weighting (method 5) lead to an improved phase estimation in all test images, compared to without the adjustment (method 4).

4.3.3   Angular dependent scattering

Previously, the coherence was measured by averaging in azimuth. This is sufficient if the signal has no angular dependence, i.e. the magnitude of the return from a target is equal for all azimuth angles. However, some complicated targets exhibit angular dependent scattering (non-Lambertian scattering) [Callow, 2003]. The coherence estimate, hence the phase variance estimate, will then not be accurate for all azimuth positions.

A measure of the *average* coherence at azimuth position v can be made using

$$\widehat{\gamma}_v = \frac{\sum_x G_x[v+1] G_x{}^*[v]}{\sqrt{\sum_x |G_x[v+1]|^2 \sum_x |G_x[v]|^2}}. \tag{4.94}$$

An estimate of the N_x' look phase variance can be made using the approximation (4.82), valid for $N_x' > 8$,

$$\mathrm{Var}\left[\widehat{\phi}[v]\right] = \frac{1 - |\gamma_v|^2}{2N_x' |\gamma_v|^2}, \tag{4.95}$$

A scene containing a target with angular dependent scattering was simulated. The variation of the signal with azimuth is shown in Fig. 4.8(a). Clearly, the signal to clutter ratio, and thus the phase variance with vary with azimuth also. The phase variance at each azimuth position estimated using (4.95) matches the measured phase variance well. Using a single estimate for all azimuth position does not give a good estimate of the phase variance.

4.3.4   Direct centre shifting

Modelling the target as a single bright scatterer plus clutter (4.42), each range bin has a linear phase offset of $\omega_x v$, giving the phase gradient a constant offset of $\omega_x = 2\pi y_0 \Delta v$. The PGA algorithm removes this offset by shifting the brightest point y_0 to the centre of the image. However, for an image which does not consist of a single bright scatterer,
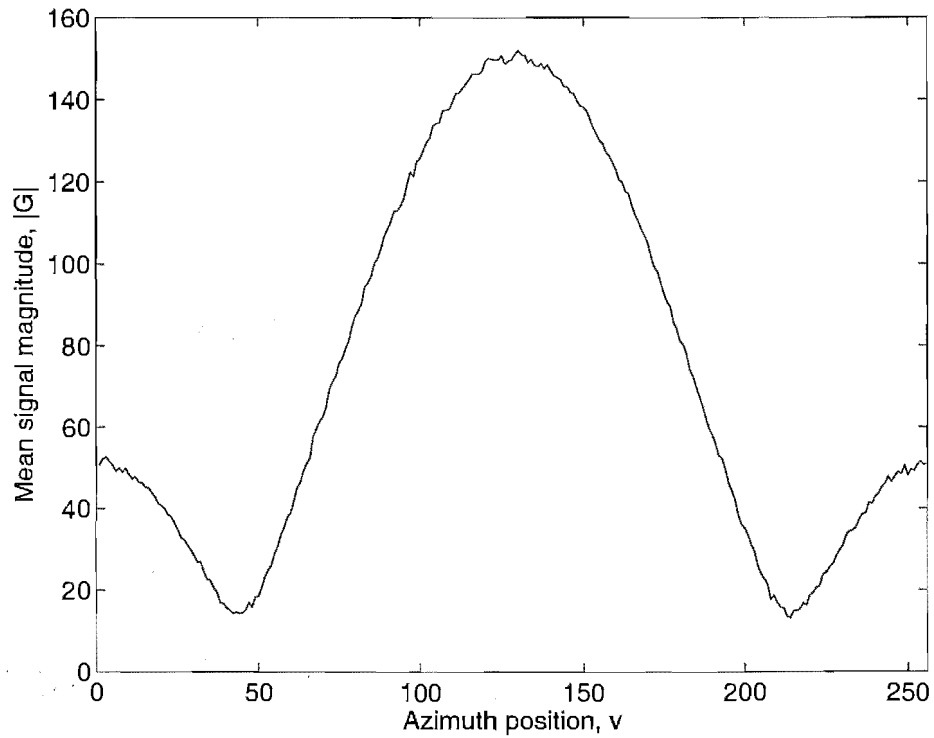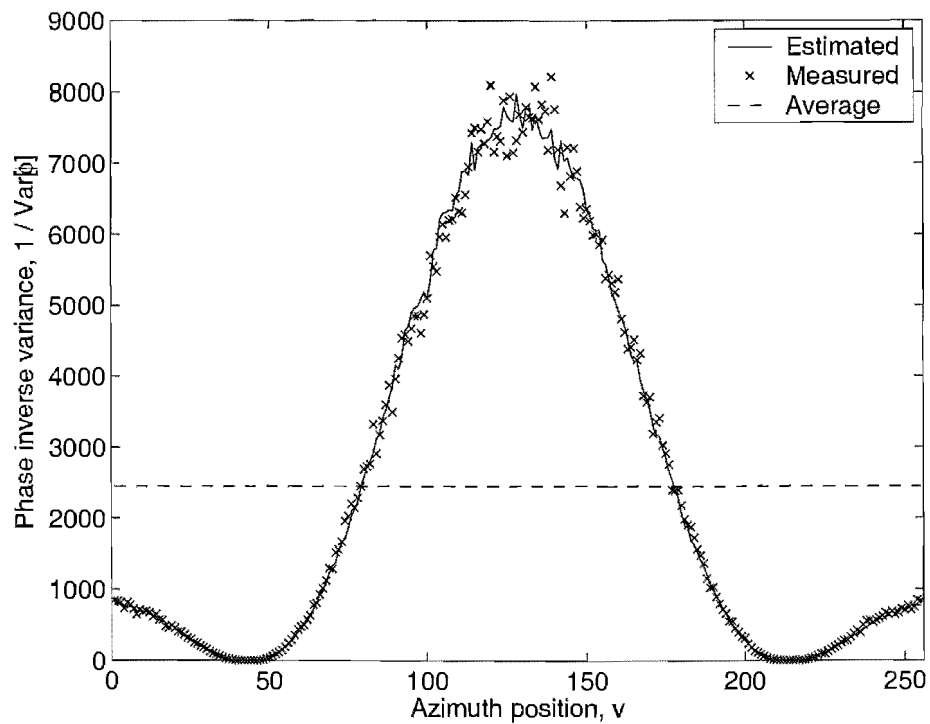
Figure 4.8 Simulation of scene with angular dependent scattering. (a) The mean signal magnitude at each azimuth position. (b) The inverse variance of the phase estimate. Phase inverse variance estimated using (4.95). Phase inverse variance at each azimuth position averaged over 80 different image simulations, with 20 different average signal to clutter ratios. The average value is calculated using (4.92).

this approach may leave a residual phase offset. This will cause an error when the phase-gradient is averaged over all range-bins. A superior method is to directly estimate and remove the mean phase-gradient in the signal, or Fourier domain. As the measured phase wraps, this is most conveniently performed by removing the mean measured phase gradient from each range-bin. Once the phase gradient $\theta_x[v] = \angle\{G_x[v+1]G_x{}^*[v]\}$ is calculated, the mean phase gradient can be calculated using

$$\overline{\theta_x} = \angle\left\{\sum_x \exp(j\theta_x[v])\right\} \tag{4.96}$$

and the mean removed using

$$\widetilde{\theta_x}[v] = \angle\left\{\exp\left[j(\theta_x[v] - \overline{\theta_x})\right]\right\}. \tag{4.97}$$

The phase is added using phasors rather than directly to remove phase-wrapping problems. This is equivalent to removing a linear Fourier phase trend using

$$\widetilde{G_x}[v] = G_x[v]\exp(-j\overline{\theta_x}v) \tag{4.98}$$

or shifting the image position using

$$\widetilde{g_x}[y] = g_x[y - y_0] \tag{4.99}$$

where $y_0 = \frac{\overline{\theta_x}}{2\pi\Delta v}$. Fig. 4.9 shows the result of the different schemes for removing linear phase trends. Fig. 4.9(b) shows the image with the brightest point shifted to the centre. This does not work well for the asymmetric speckled block and shadow, leaving a residual phase error which will bias the Fourier phase estimation techniques.

4.3.5   Higher order estimation

Phase gradient algorithms estimate the difference in phase between adjacent azimuth bins. It is possible to estimate the phase difference for larger separations. The redundant information that is available when more than two azimuth bins are used results in a more accurate estimation if the error at each separation is independent. Jakowatz and Wahl [1993] developed the eigenvector phase estimation method which used more than one separation to estimate the phase. If M azimuth bins were used to estimate the phase at a single azimuth position, then the CRLB of the error in the estimated phase is

$$\text{Var}\left[\widehat{\phi} - \phi_e\right] \geq \frac{1 + M\,\text{SNR}}{M\,N_x'\,\text{SNR}^2} = \frac{1}{N_x'}\left[\frac{1}{\text{SNR}} + \frac{1}{M\,\text{SNR}^2}\right], \tag{4.100}$$

which for $M = 2$ becomes (4.52). This is analogous to a multi-channel time-delay estimation with M hydrophones or receivers [Quazi, 1981]. The approach by Jakowatz

(a)

(b)

(c)

(d)
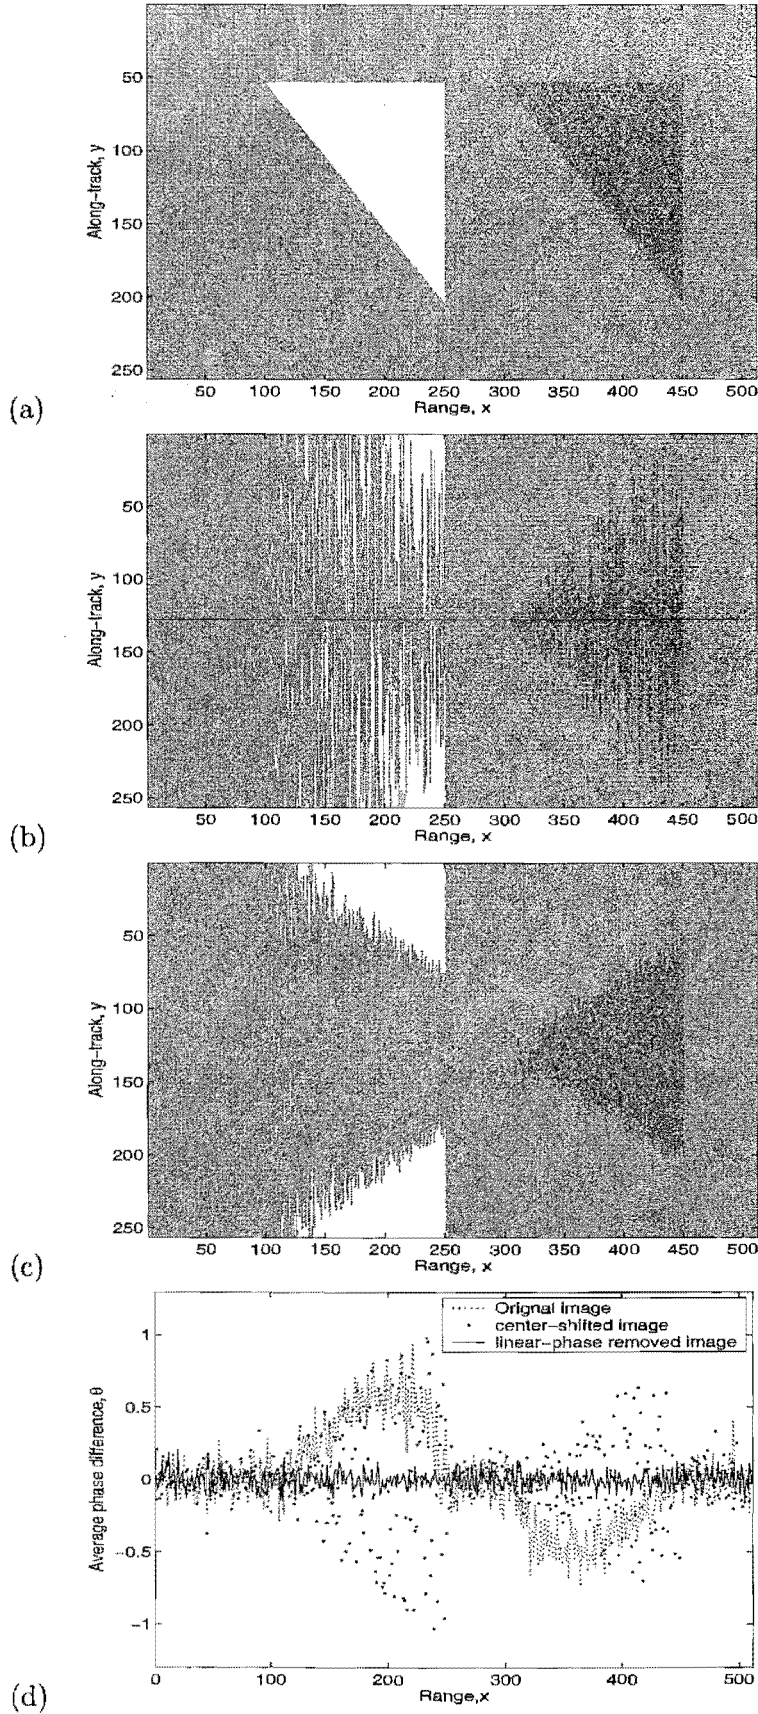
Figure 4.9 Centre shifting of spotlight images. (a) Original image. (b) Image with brightest point shifted to centre. (c) Image with mean phase difference removed. (d) Mean phase difference versus range.

and Wahl [1993] modelled the image at each range bin as a single point target with additive white noise. This section develops a high-order method with a more general image model, using ping-ping coherence to weight the measures appropriately.

The phase gradient at separation d, $\Delta\phi_d[v] = \phi_e[v+d] - \phi_e[v]$ can be estimated using

$$\Delta\widehat{\phi_d}[v] = \angle\left\{\sum_x \widehat{\sigma_d}^{-2}[x] \frac{G_x[v+d]G_x{}^*[v]}{|G_x[v+d]G_x{}^*[v]|}\right\}. \tag{4.101}$$

The variance of this estimate can be approximated by

$$\widehat{\sigma_d^2} = \left[2\sum_x \frac{|\widehat{\gamma_d}[x]|^2}{1-|\widehat{\gamma_d}[x]|^2}\right]^{-1}. \tag{4.102}$$

The estimates can be combined using

$$\widehat{\phi}[v] = \frac{\sum_d \left(\widehat{\phi}[v-d] + \Delta\widehat{\phi_d}[v-d]\right)\widehat{\sigma_d}^{-2}}{\sum_d \widehat{\sigma_d}^{-2}}. \tag{4.103}$$

However, directly summing phase estimates in (4.103) can cause problems if the phase estimates wrap. A more robust method is to sum phasors using

$$\widehat{\phi}[v] = \angle\left\{\sum_d \widehat{\sigma_d}^{-2}\exp\left(j\widehat{\phi}[v-d]\right)\exp\left(\Delta\widehat{\phi_d}[v-d]\right)\right\}. \tag{4.104}$$

This is a simpler approach than estimating the full multichannel cross-correlation coefficient [Benesty et al., 2004].

Fig. 4.10 plots the CRLB of the phase variance for several M from (4.100). Increasing M makes a small improvement to the variance, but it quickly becomes close to the asymptote. As $M \to \infty$, the lower bound reaches the asymptote $\sigma_\phi^2 \geq \frac{1}{N_x'.\text{SNR}}$. This plot shows only a small theoretical gain in using a higher-order estimator. In practice, using a higher-order estimator did not make a noticeable improvement in the accuracy of the phase estimation. However more work may show that a significant increase in performance is achievable using a higher-order estimator. Fig. 4.10 shows that the most gain is achievable at a low SNR.

4.4 RESULTS

Several scenes were simulated, with the Fourier phase error estimated by various means of echo-correlation. The scenes used are shown and described in Fig. 4.11. Any linear phase offset was first removed by using (4.97), then the phase gradient estimated using six different weightings. The variance of the error in the resulting phase gradient estimates over all azimuth bins was measured. The results are shown in Table 4.1.
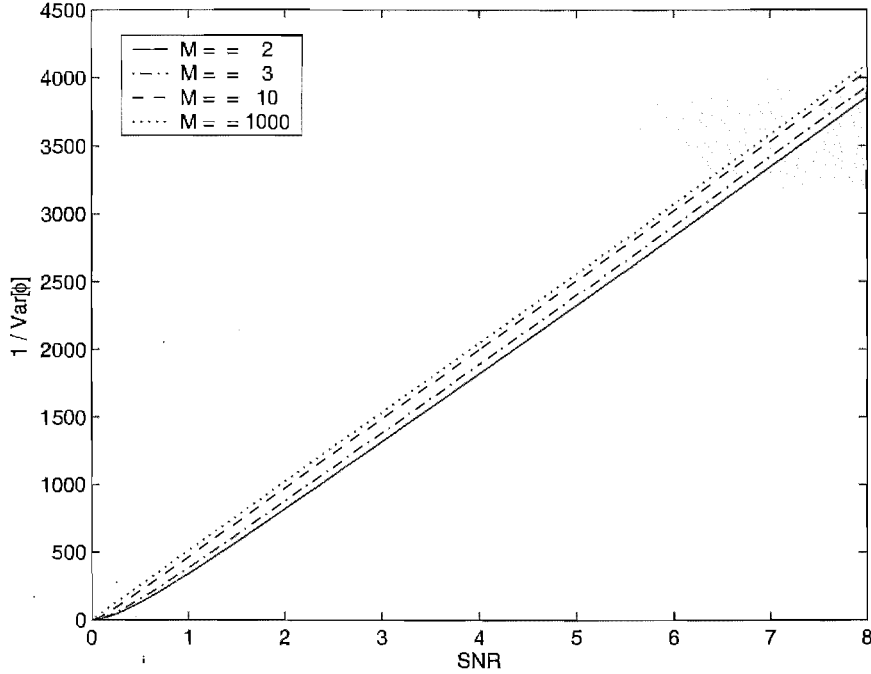
Figure 4.10 The CRLB (4.100) of the phase estimate variance versus SNR for $N'_x = 512$ for estimator of order M. Increasing M makes only minimal improvement to variance.

The results clearly show the importance of varying the azimuth weighting using (4.89) rather than using the same average weighting over a whole range-bin (4.87). This can be seen by comparing the results from method 2 to method 1 and method 4 to method 3. This is especially important for images that have varying coherence with azimuth (images C, D and F). For these images, the estimated phase-gradient variance was well below the measured levels. This is due to coherence in range reducing the number of independent samples averaged. (Noise is correlated in range).

When comparing the different range-weighting methods, the results show that using the estimated inverse variance from the simple coherence measure (method 2) outperforms shear-average for all images. However, using the (more accurate) variance model from the pdf (method 4) performs worse. The reason for this appears to be due to the model estimating wrapped phase variance, thus over-weighting ranges with low coherence. Using the adjusted pdf measure (method 5) improves performance of the estimator on all tested images. The adjusted pdf measure (method 5) performs similarly well to the simple coherence measure (method 2).

4.4.1 CRLB

In order to compare the results of echo-correlation to the derived CRLB (4.52), images of varying SCR were simulated and the variance of the phase-difference estimate was measured using both shear average and weighted phase difference estimation (WPDE).
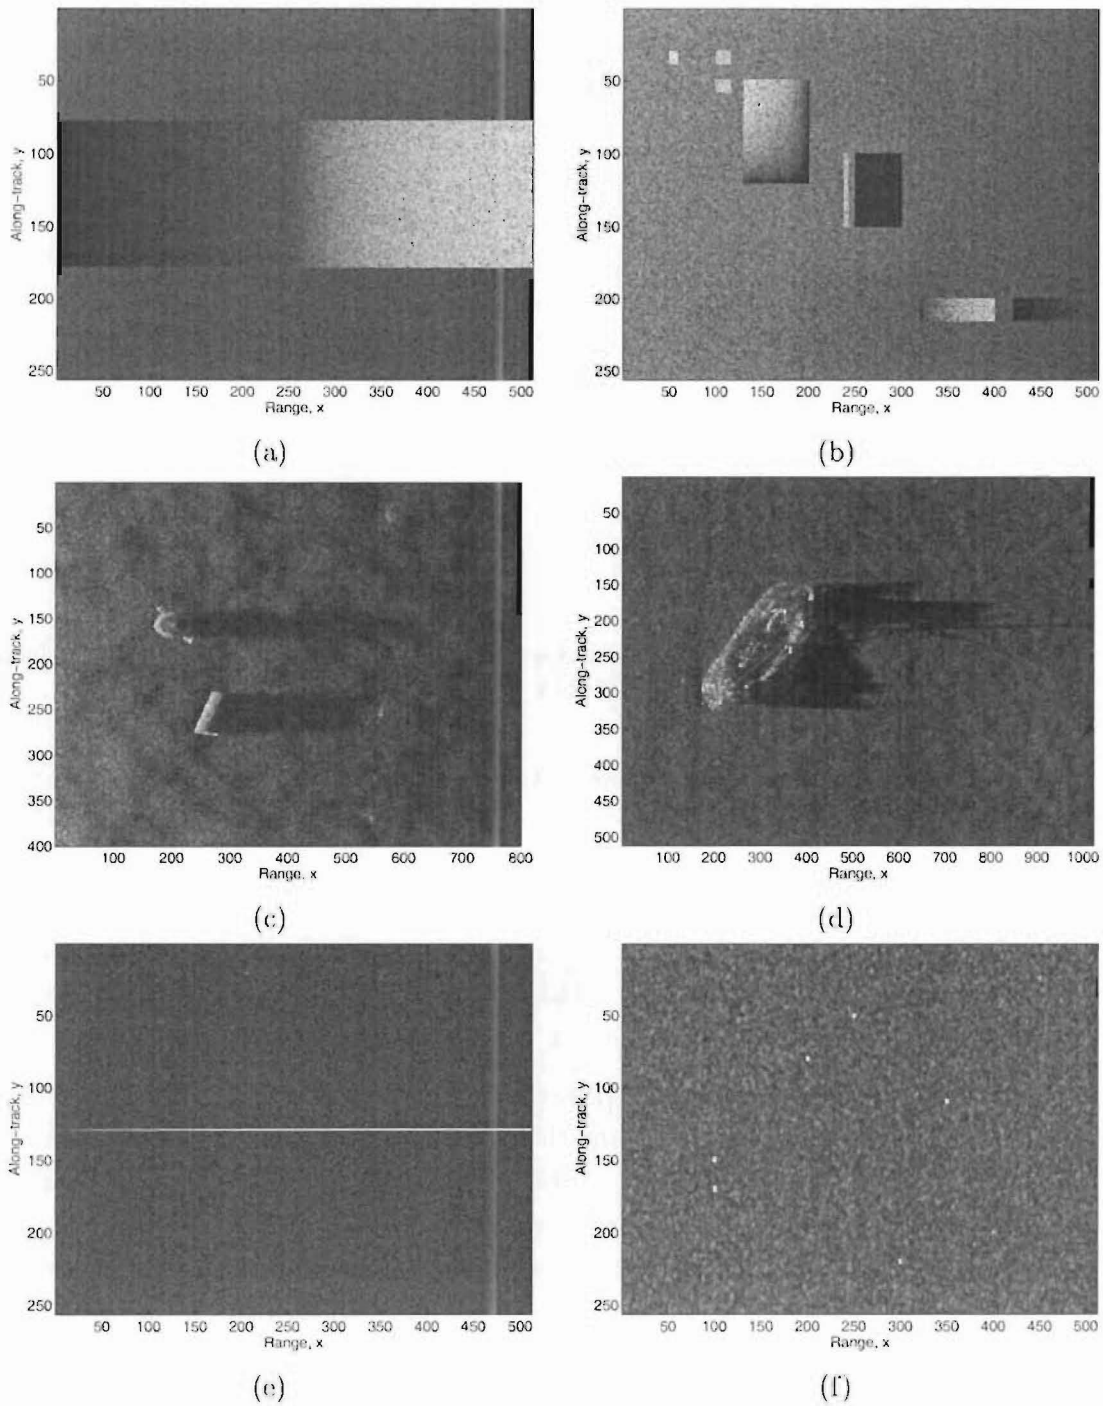
Figure 4.11 Scenes used to test echo-correlation methods. (a) A: Rectangular shadow and block target. (b) B: Combination of block targets and shadows. Background intensity decreases with range. (c) C: Simulated sonar image of mine-like objects with shadows. Seafloor has ripples and decreasing intensity with range. Scene is band-limited, white additive noise added. (d) D: Real sonar image of shipwreck with seafloor ripples and large shadow. (e) E: Point target in every range bin, increasing intensity. (f) F: Band-limited scene of several point targets and clutter.

Method	Scene					
	A	B	C	D	E	F
SA	0.00217	0.01118	0.01755	0.00830	0.000218	0.28715
1	0.00266	0.01078	0.01692	0.02259	0.000234	0.23361
2	0.00130	0.00576	0.01040	0.00673	0.000218	0.10391
3	0.00288	0.01193	0.01626	0.02780	0.000240	0.21866
4	0.00154	0.00870	0.01262	0.02044	0.000227	0.38981
5	0.00131	0.00563	0.00934	0.00677	0.000212	0.10963
$\widehat{\sigma_\phi^2}$	0.00143	0.00494	0.00311	0.00116	0.000245	0.02766

Table 4.1 Mean squared error of phase difference estimates of various simulated images. The scenes are shown and described in Fig. 4.11. Method SA: Shear average, (4.49). Method 1: WPDE, simple coherence measure (4.73), no azimuth averaging (4.88). Method 2: WPDE, simple coherence measure (4.73) with azimuth averaging (4.89). Method 3: WPDE, pdf measure (4.79), no azimuth averaging (4.87). Method 4: WPDE, pdf measure (4.79) with azimuth averaging (4.89). Method 5: WPDE, adjusted pdf measure (4.93) with azimuth averaging (4.89). $\widehat{\sigma_\phi^2}$: The phase gradient variance calculated from the coherence (4.91). (Assumes delta correlated clutter).

The varying azimuth method (4.89) was used with the simple inverse variance measure (method 2 above). The experiment was performed for both a point target, and a rough block target. The noise power was either constant with range, or varied with range. The inverse variance is shown in Fig. 4.12. For the images with constant SCR Fig. 4.12(a,c), both the shear-average and WPDE results have similar performance and cluster around the predicted CRLB. For the images in which the SCR varies with range Fig. 4.12(b,d), WPDE clearly outperforms shear-average with a lower phase-variance, especially at low SCR. WPDE results match the CRLB well, with shear-average failing to meet the CRLB. The performance gain was more significant for the point-target image.

4.5 STRIPMAP IMAGES

The effect of a slant-range platform sway of $w(u)$, assuming a narrow beam-width and narrow bandwidth, can be modelled as a 1D Fourier phase error in *along-track* using (2.12)

$$d(t, u) \approx \overline{d(t, u)} \exp\left[j2k_0 w(u)\right], \tag{4.105}$$

where $k_0 = 2\pi f_0/c$. The equivalent of shear-average can be performed by estimating the average phase difference between pings using

$$\Delta\widehat{\phi}(u) = \angle\left\{\sum_t d(t, u + \Delta u)d^*(t, u)\right\}. \tag{4.106}$$
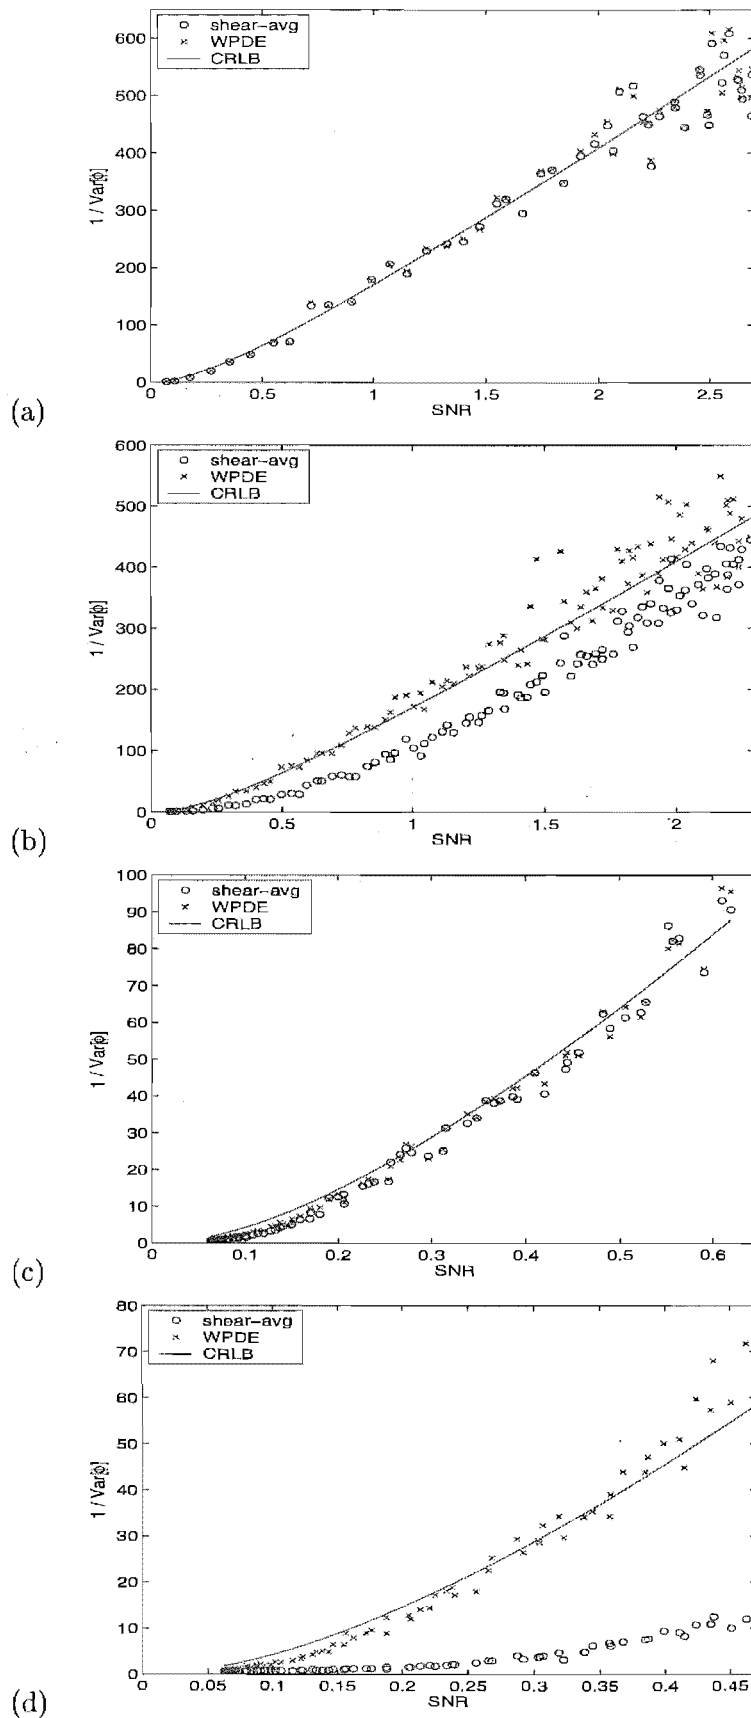
Figure 4.12 Inverse variance of the phase difference estimate using shear average and WPDE compared to the CRLB (4.92). WPDE uses (4.89) and the simple variance estimate (4.73). Image is 256 × 256 pixels. (a) Block target 101 pixels wide, noise constant with range. (b) Block target 101 pixels wide, noise power decreasing with range. (c) Constant point target in every range-bin, noise constant with range. (d) Constant point target in every range-bin, noise power decreasing with range.

A phase estimate can be obtain by integrating the phase-difference estimate

$$\widehat{\phi}(u + \Delta u) = \widehat{\phi}(u) + \Delta\widehat{\phi}(u), \tag{4.107}$$

and the sway estimated by scaling the phase estimate

$$\widehat{w}(u) = \frac{\widehat{\phi}(u)}{2k_0} = \frac{\widehat{\phi}(u)}{4\pi f_0/c}. \tag{4.108}$$

Echo-correlation behaves differently in stripmap and spotlight imaging. In spotlight, speckle is often uncorrelated ping-ping, whereas in stripmap, some along-track correlation occurs if the aperture is sampled sufficiently (see Section 3.1.2). Thus in stripmap, seafloor clutter can be used to estimate the Fourier phase error. In spotlight, a point target has linear Fourier phase, thus a constant Fourier phase difference. This offset can be removed by centre shifting, or the direct method discussed in Section 4.3.4. In stripmap, the phase of the pulse-compressed signal from a point target is hyperbolic, or approximately quadratic. Thus the phase difference is approximately linear. When estimating the shear average, this linear bias will cause errors in the region of a strong scattering target. Fig. 4.13(c) shows the phase bias caused by a strong reflector in a simulated scene.

This phase bias can be dealt with in one of four ways. The phase could be ignored by using a noncoherent shear-average; The weighting of the phase of strong-scattering targets causing bias could be reduced using amplitude weighting; The phase bias of strong targets can be calculated and compensated for, using the stripmap equivalent of PGA; Or the linear bias of targets can be removed by differentiating the phase gradient again, then estimating the phase curvature.

4.5.1 Noncoherent shear-average

Noncoherent shear-average works by correlating the *magnitude* of adjacent pings [Callow et al., 2001b; Callow, 2003]. The advantage of this is three fold. As it does not use the phase information, phase bias by strong targets does not affect the results. Secondly, large path-deviation can be estimated without phase-wrapping errors. Thirdly, the correlation length is extended. This allows correlation of pulses spaced at more than $D/2$ apart, where D is the extent of the receiver aperture. The disadvantage is that ignoring phase reduces the accuracy, especially in high Q systems, as shown in Section 4.1.5. The results in Table 4.2 show that non-coherent shear-average has a larger mean-squared error than techniques that use the phase information, as long as phase wrapping is successful for the other techniques.

(a)

(b)

(c)

(d)
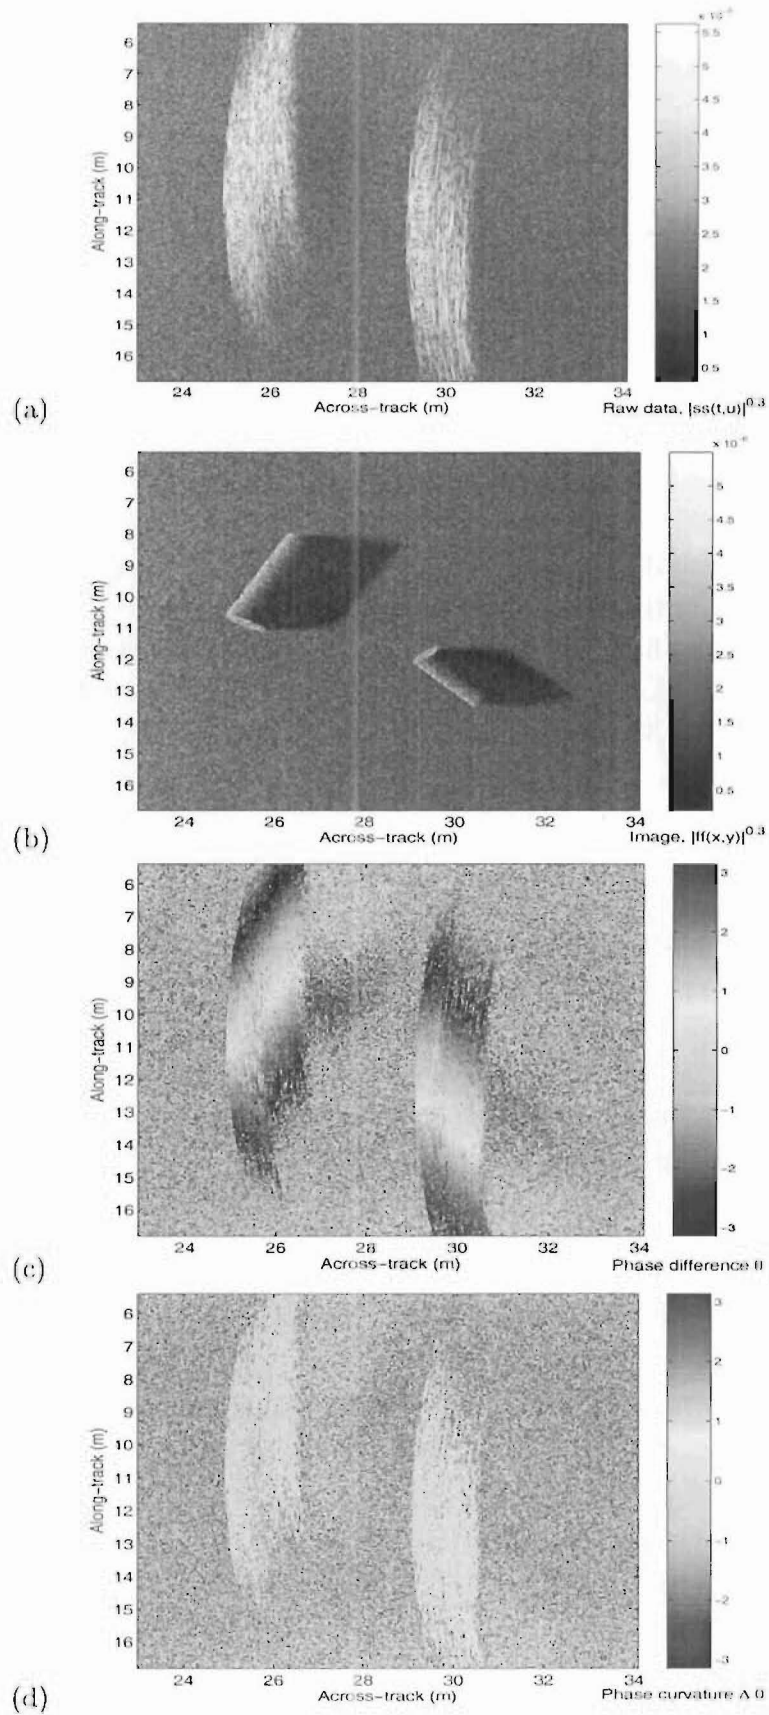
Figure 4.13 Simulated SAS scene of two cylinders on flat seafloor. (a) Magnitude of pulse compressed data $d(t, u)$. (b) Magnitude of reconstructed image $g(x, y)$. (c) Phase difference $\theta(u, t)$ between adjacent pings. There is an obvious linear bias across the return from strong targets and a smaller bias across the shadows. Phase difference of speckle is random. (d) Phase curvature $\Delta\theta(u, t)$ between adjacent pings. Linear bias across targets is removed.

4.5.2 Amplitude weighting

A strong scatterer biases the Fourier phase estimate. A method to reduce the bias is to reduce the weighting of strong-scattering targets. The advantage of this method is its speed and simplicity. The disadvantage is that it reduces the SNR, by effectively blocking out the strongest signal. For a discussion on amplitude weighting in stripmap shear-average, see [Callow *et al.*, 2001a; Callow, 2003]. An amplitude weighting may be applied using

$$\Delta\widehat{\phi}(u) = \angle\left\{\sum_t \beta[u,t]d(t, u + \Delta u)\, d^*(t, u)\right\}. \tag{4.109}$$

If the phase error is large, $\Delta\widehat{\phi}(u)$ will wrap. The estimated phase-gradient needs to be unwrapped before integrating to a phase estimate, or scaling to estimate the sway.

Table 4.2 shows results comparing three different amplitude weightings. These include using no weighting

$$\beta_N[u,t] = 1, \tag{4.110}$$

weighting all estimates the same

$$\beta_C[u,t] = \frac{1}{|d(t, u + \Delta u)\, d^*(t, u)|}, \tag{4.111}$$

and a weighting developed by Callow [Callow, 2003]

$$\beta_H[u,t] = \frac{1}{\alpha + |d(t, u + \Delta u)\, d^*(t, u)|^2}, \tag{4.112}$$

where alpha is selected so the peak weighting is above the noise floor. A value of $\alpha = \max\left\{|d(t, u + \Delta u)\, d^*(t, u)|^2\right\}/1000$ was found to work well and is used in the results shown. The results show the weighting β_H giving the best results, with β_C close in performance and β_N performing poorly. As β_H is essentially an ad-hoc weighting, there is further potential to develop an optimal weighting for stripmap shear-average.

4.5.3 Stripmap phase gradient autofocus

In spotlight imagery, the phase bias of strong targets can be removed by circular shifting. In stripmap, this step is not so simple. It is possible to estimate the position of strong targets and thus the phase bias caused by them. This is the essence of the equivalent form of PGA for stripmap imagery, stripmap phase gradient autofocus or SPGA. For details on the operation and performance of SPGA, see [Callow *et al.*, 2003; Callow, 2003].

4.5.4   Phase curvature

The phase of a point target is hyperbolic, or approximately quadratic. The phase gradient is then approximately linear. Differentiating again, the phase curvature $\Delta_\phi^2(u) = \Delta\phi[u + \Delta u] - \Delta\phi[u]$ is approximately constant. Thus strong targets do not bias the phase curvature. Fig. 4.13(c) shows a linear bias over the phase difference of a strongly reflecting target. In the phase curvature Fig. 4.13(d), the bias is removed.

The advantage of this method is that bias is eliminated, so the full SNR can be used. The phase curvature is significantly smaller than the phase-gradient, so is unlikely to wrap. This means larger sways can be estimated without phase-unwrapping. The disadvantage is that the phase curvature is noisier than the phase gradient. Thus the accuracy of the phase-curvature estimation can suffer, especially at low SCR. This can be seen as a higher phase gradient variance in Table 4.2. Also, a quadratic phase error cannot be estimated with the phase-curvature. Due to the double integration, phase curvature is poor at estimating low-order phase errors. This makes it poor for use as motion estimation technique. However, as high-order phase errors degrade the image more, it is a useful autofocus technique.

The phase curvature can be estimated using the two step process of

$$C(u, t) = d(t, u + \Delta u)\, d^*(t, u) \tag{4.113}$$

$$\Delta^2\phi(u) = \angle\left\{\sum_t W(t)C(t, u + \Delta u)\, C^*(t, u)\right\}. \tag{4.114}$$

The phase difference is the integral of the phase curvature,

$$\Delta\widehat{\phi}(u + \Delta u) = \Delta\widehat{\phi}(u) + \Delta^2\phi(u) \tag{4.115}$$

The results shown in Table 4.2 apply two different weightings,

$$W_N(t) = 1 \tag{4.116}$$

and

$$W_{ML}(t) = \frac{|\widehat{\gamma_C}(t)|^2}{1 - |\widehat{\gamma_C}(t)|^2} \frac{1}{\sum_t |C(t, u + \Delta u)\, C^*(t, u)|}, \tag{4.117}$$

where

$$|\widehat{\gamma_C}(t)|^2 = \frac{\sum_u |C(t, u + \Delta u)\, C^*(t, u)|^2}{\sum_u |C(t, u + \Delta u)|^2 \sum_u |C(t, u)|^2}. \tag{4.118}$$

This coherence weighting (W_{ML}) is the same as that developed for weighted phase-gradient estimation in Section 4.3. Table 4.2 shows a significant performance increase in using the coherence weighting, compared to using no weighting. In stripmap images, the mean signal coherence varies significantly more with u than with spotlight images (see Section 4.3.3).

Fig. 4.14 shows an example of using weighted phase curvature estimation (WPCE) to autofocus a field SAS image. Fig. 4.14(b) shows the phase curvature, with the strong targets showing lower phase variance and no linear bias. The coherence plots show large peaks when targets are present, signifying a higher SCR. The reconstructed image corrected with the phase estimate obtained using WPCE shows less blurring than the image reconstructed with no phase correction.

In stripmap images, a target return is localised to a limited number of pings. This leads to a large variation in signal coherence in along-track, as shown in Fig. 4.14(d). Rather than form a weighting by averaging coherence over all pings as in Fig. 4.14(c), a better approach may be to measure coherence over a small localised patch [Touzi *et al.*, 1999].

The echo of a strong target will not only experience a phase shift along the aperture, but may also shift between range-bins. This can be seen as a hyperbolic curve of a strongly reflecting target in the range-compressed data Fig. 4.13(a). It has a stronger effect (tighter curve) at short ranges than long ranges. This is known as footprint-shift or range migration in interferometry. It has the effect of decreasing the correlation between pings, which increases the phase estimate variance [Just and Bamler, 1994; Bamler and Hartl, 1998]. The range-migration ping-ping is minimal at a reasonable standoff range. It may be a problem if a higher order phase-curvature estimator is used (greater separation Δu). If range migration is a problem, it can be reduced by sub-banding. This reduces the range resolution, making any range-migration a smaller proportion of a range-cell, increasing the coherence between looks [Barclay *et al.*, 2005].

Phase curvature autofocus (PCA)

The phase curvature estimation technique described may be considered a simplification of the phase curvature autofocus (PCA) algorithm proposed by Wahl et al. [1994b]. WPCE directly measures the phase curvature of the phase corrupted data $d(t, u)$ making the assumption that the target phase history is quadratic. PCA operates on the reconstructed image $g(x, y)$. This allows the additional step of windowing point-like targets to increase the SCR. The image is then convolved with an appropriate along-track chirp for each range-bin. This spreads the target response out and *unfolds* the effect of a phase error [Hawkins, 1996, pages 152-153]. The hyperbolic phase variation of each target is removed, leaving a linear phase trend [Pat, 2000]. The phase curvature may then be estimated by averaging in range. This process has been simplified in WPCE by performing it directly in the data domain where the phase errors occur. WPCE is a faster technique to compute as no image reconstruction or convolution is required. Due to the windowing step and the more accurate hyperbolic phase matching, PCA is likely to be more accurate. However, the accuracy of the two methods have not been compared.

(a)

(b)

(c)

(d)

(e)

(f)
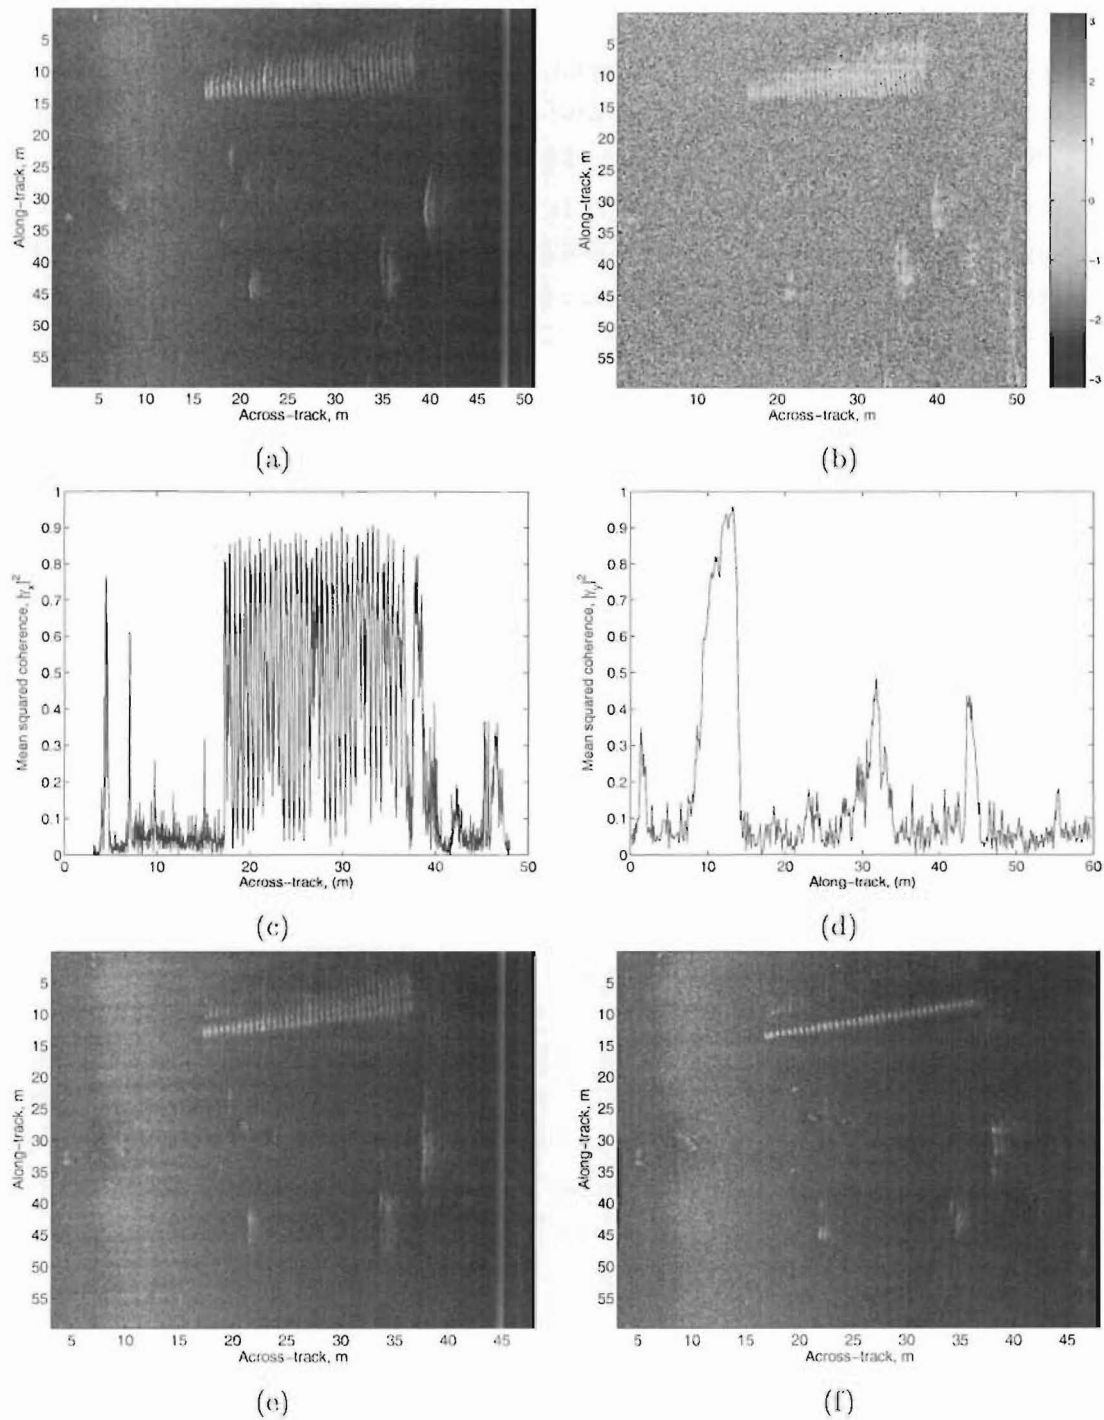
Figure 4.14 Example of using weighted phase curvature estimation (WPCE) to autofocus a field SAS image. Coherence weighting (W_{ML}) is used. (a) Magnitude of pulse-compressed data $d(t, u)$. (b) Phase curvature between adjacent pings $\Delta\theta(t, u)$. (c) Mean squared coherence at each range-bin $|\gamma_C(t)|^2$ (4.118). (d) Mean squared coherence at each ping $|\gamma_C(u)|^2$. (e) Original reconstructed image $g(x,y)$. (f) Reconstructed image corrected using WPCE $\hat{g}(x,y)$.

Wideband phase curvature estimation

Phase curvature estimation assumes a narrow bandwidth, as the phase error is assumed constant at all frequencies. It will lose accuracy for low-Q systems. A wide bandwidth version of PCA was described by Hawkins [Hawkins, 1996, page 155]. Hayes et al. [2002] described a wide-band, wide beam generalisation of PCA, estimating a 2-D phase error. A similar approach may be possible for WPCE. As a wide-band timing error causes a frequency dependent phase shift, the phase curvature could be estimated in the $D(f, u)$ domain.

4.5.5 Results

Single Scene

The methods of shear average, noncoherent shear average and phase curvature estimation were compared on a simulated SAS image. The image was given a known path error, the path-difference was estimated by each method, and the mean-squared error between the known and estimated path-difference recorded. This was repeated for 100 different random phase errors. The mean results are shown in Table 4.2.

The amplitude-weighted shear-average performed the best. This experiment performed phase unwrapping of the phase difference. If this was not performed, the performance of amplitude weighted shear-average was significantly poorer. The Callow weighting β_H was the best performing weighting, with a constant weighting being close in performance. Both were significantly better than using no weighting, which had similar performance to phase-curvature estimation using the coherence weighting W_{ML}. Using the coherence weighting led to a significant performance advantage over using no weighting. Phase-curvature estimation performed better than non-coherent shear-average, as long as the coherence weighting was used.

Method	Weighting	Error
Amplitude	None (β_N)	0.0552
	Constant (β_C)	0.0061
	Callow (β_H)	0.0038
Curvature	None (W_N)	0.4355
	Coherence (W_{ML})	0.0764
Noncoherent	None	0.2323

Table 4.2 Mean-squared error of path difference estimate using amplitude-weighted stripmap shear-average, phase-curvature estimation and noncoherent shear-average. Results are averaged over 100 trials with random known path corrupting a simulated SAS scene.

Varying SCR

A scene consisting of two cylinders on a flat seafloor was simulated with the targets and seafloor simulated separately. The targets could be added to the scene with a varying power, thus simulating the same scene with a varying signal to clutter ratio (SCR). The path difference was estimated using several different methods for 20 different instances of a random path for each scene. The mean-squared error of the path difference was measured for each path. The mean value over all paths was then recorded for a number of different scene signal to clutter ratios.
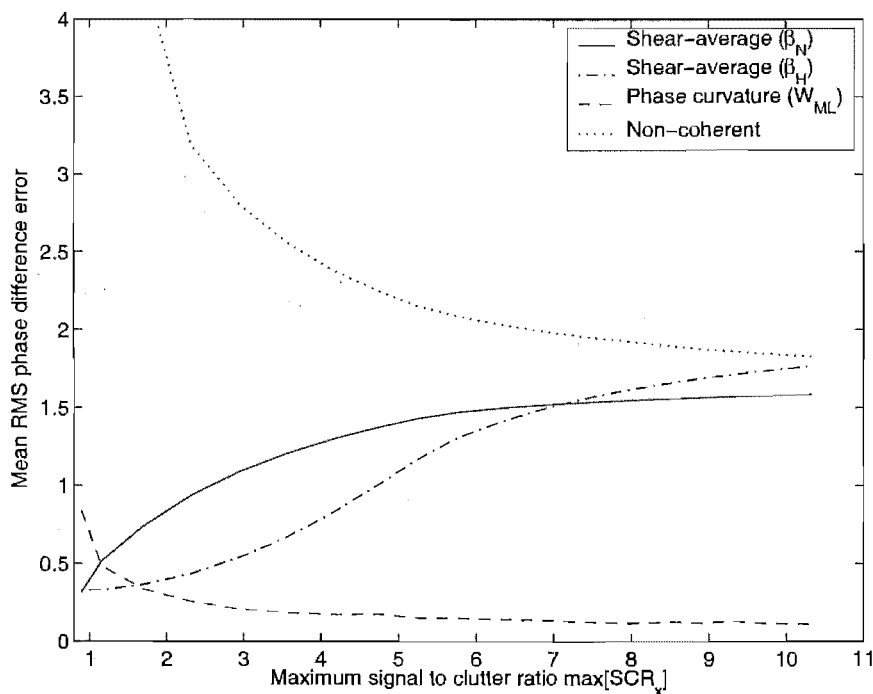
Figure 4.15   RMS path difference error of various path estimation methods at different signal to clutter ratios. Scene is shown in Fig. 4.13 with the targets added to the background with varying ratios. Errors are averaged over 20 random paths. SCR is the maximum over all range-bins.

The results, shown in Fig. 4.15 show the relative performance of the various phase estimation techniques changes with signal to clutter ratio. As the signal-to clutter ratio increases, the performance of the non-coherent shear average and weighted phase curvature estimation techniques improve. However, the amplitude-weighted shear-average method performs worse. This is predictable, as a strong target biases the shear-average method, so a brighter target will bias it more. The amplitude weighting technique suppresses the return from a strong target, so an improving SCR will not improve the performance. In contrast, the phase curvature and non-coherent shear-average techniques use the phase estimate from the strong target as well as the background. As the target gets stronger, the ping-ping coherence will improve and the phase variance decreases. Weighted phase curvature estimation outperforms all other techniques tried by a significant margin at a high signal to clutter ratio. At a low SCR, amplitude-

weighted shear-average performs better than weighted phase curvature estimation. It remains to be seen whether one can accurately predict the cross-over point, so that the SCR can be measured and used to decide which technique will be more effective.

4.6 CONCLUSIONS

A weighted phase difference estimation (WPDE) is shown to be an improvement over the standard technique for estimating the average phase difference between echos (shear-average). WPDE is a narrow-band implementation of the generalised correlation method of time-delay estimation. The phase difference estimate at each range-bin is weighted by an estimate of the inverse variance of that estimate, using a measure of the mean signal coherence. This gives an improved weighting, compared to the signal energy weighting used in shear average.

A model of the image phase variance is developed. The phase variance of a single range-bin has a different model for an extended rough target, compared to a constant point target. When the phase estimate from a number of range-bins are combined, the point target model is also accurate for an extended target.

A number of methods of weighting the phase difference estimates are compared. It is important to normalise the weighting with the mean signal energy $\sum_v |G_x[v + 1] \, G_x^*[v]|$, so that the weighting varies with azimuth position. The inverse variance weighting should be for unwrapped phase, so that the weighting goes to zero for clutter only. If a wrapped phase model is used, a simple adjustment to zero the weighting gives improved results. The point target and extended target phase variance models give similar results over a range of images. Thus the simpler point-target model is preferable.

If a target with angular dependent scattering is imaged, the signal to clutter ratio will vary with azimuth. Thus the phase estimate variance will vary with azimuth also. The phase estimate variance at each azimuth position may be estimated by measuring the average signal coherence between adjacent echos.

A direct centre shifting method is proposed. The mean phase difference is directly estimated and removed from each range-bin. This results in reduced phase bias compared to shifting the single brightest point to the centre of the image.

WPDE is shown to meet the CRLB for delta correlated scenes. If the SCR varies with range, shear-average performs below the CRLB. For scenes with correlated speckle, the performance of WPDE is below the CRLB.

In a stripmap image, the phase difference estimate is biased since a strong target has a linear phase difference across the aperture. A weighted phase curvature estimator (WPCE) removes this bias at the cost of a noisier phase estimate. WPCE gives a better phase estimate than non-coherent shear average and amplitude-weighted shear average for images with high SCR. For a low SCR, amplitude weighted shear-average

performs better. WPCE is similar in concept to phase curvature autofocus (PCA), which is performed on the reconstructed image. Compared to a phase gradient algorithm, WPCE is noisier and is poorer at estimating low-order phase errors due to the double integration. WPCE also assumes the data is narrow-band. WPCE is fast to calculate however, so it may be a useful fast (real-time) stripmap autofocus technique.

Chapter 5

STATISTICAL AUTOFOCUS

This chapter describes a general framework for the estimation of blurring parameters (autofocus) of a stripmap or spotlight synthetic aperture sonar (SAS) or synthetic aperture radar (SAR) image. Motion of the imaging platform from a straight path can cause blurring of the image. Statistical autofocus (SAF) describes this motion by a set of parameters. These parameters are altered to minimise a cost function. The cost function combines a measure of the image quality, using image sharpness or contrast, and a measure of the likelihood of the estimated motion. Together these form a regularised sharpness maximisation problem. Section 5.3 looks at different methods of measuring image sharpness and the effect different sharpness measures have on the path estimation. A method of measuring the likelihood of a path estimate based on statistics of the platform motion is developed in Section 5.4. Section 5.5 shows the optimal method of parameterising the platform motion, while Section 5.6 looks at different optimisation methods to minimise the measured cost function.

The technique of sharpness maximisation is comparable to adaptive optics schemes used for real-time correction of atmospherically degraded telescope images through image sharpening [Muller and Buffington, 1974]. However, the phase variations are introduced to the image in post-processing through an optimisation algorithm, rather than a real-time feedback system. An extra difficulty of stripmap autofocus is the blurring is point-spread-variant, meaning the optimising elements are highly coupled. The method developed is different from other stripmap contrast optimisation schemes [Sutton et al., 2000], in that it treats the path as a whole, rather than breaking it into strips, so it avoids tight restrictions on scatterer positions. Rather than represent the path by simple low-order polynomial terms [Berizzi et al., 1996], the algorithm can autofocus an arbitrary path up to the desired precision, by including enough path parameters in the optimisation. Due to the point-spread-variant nature of the blurring, no analytic gradient of a stripmap image has been determined. This rules out using a highly efficient conjugate-gradient optimisation algorithm [Gough and Lane, 1998; Fienup, 2000].

5.1 HISTORY OF SHARPNESS MAXIMISATION

Muller and Buffington [1974] proposed the use of image sharpening techniques to correct phase distortion of astronomical images. This included a proof that certain sharpness measures are at a maximum when the phase error is zero. Hamaker et al. [1977] showed a different, simple proof of the sharpness metric introduced by Muller and Buffington. They also indicated the limited validity of other suggested metrics. Paxman and Marron [1988] showed the technique of maximising image-sharpness could be applied to speckled coherent imagery such as synthetic aperture radar (SAR). Sharpness maximisation was first developed for spotlight SAR by estimating a single motion parameter (acceleration) for a small image [Finley and Wood, 1985; Blacknell et al., 1992]. They compared results of using contrast optimisation and registration of multi-look images (map-drift), finding they gave similar results. Sharpness maximisation was also applied to Inverse Synthetic Aperture Radar (ISAR) [Berizzi and Corsini, 1996]. It was extended to higher order motions [Berizzi et al., 1996] and using entropy as a sharpness measure [Xi et al., 1999]. Xi et al. [1999] also developed a novel optimisation technique, the stage by stage approaching (SSA) algorithm. SSA has been compared to PGA on simulated and real SAR data [Morrison and Munson, 2002; Morrison, 2002]. Results show SSA works well but is computationally more demanding.

For a spotlight system, a closed-form expression can be obtained for the gradient of the sharpness metric with respect to phase-error parameters. This allows the use of a highly efficient conjugate-gradient search algorithm for the minimisation procedure [Gough and Lane, 1998], [Fienup, 2000]. Along with this, an arbitrary path, or much higher-order Fourier phase errors were estimated. Fienup and Miller [2003] explained how different sharpness measured worked and compared their performance on a variety of SAR images.

Autofocus of SAS images differs from SAR primarily because to date, SAS produces strip-map images. Limited attempts have been made to use sharpness maximisation on stripmap SAS images [Sutton et al., 2000]. The image was divided into strips and a single parameter (platform acceleration) estimated for each strip. The technique has been extended to arbitrary order phase error [Fortune et al., 2001a; 2001b; 2002] and a measurement of path-probability used to regularise the result. This technique was named statistical autofocus (SAF) and is developed further in this chapter.

5.2 BAYESIAN FRAMEWORK

Consider a general off-axis motion represented by $\mathbf{w}$ forming the measured aberrated data $\mathbf{d}$. This thesis will normally use $\mathbf{w}$ to represent a one-dimensional sway or phase error, but this can be extended to include other motions. There are an infinite number of possible paths that formed the data. Approaching the autofocus problem from a

statistical view, the problem is this. Given the measured data $\mathbf{d}$, what is the most likely estimate of the motion parameters $\hat{\mathbf{w}}$? The maximum likelihood estimate of motion is given by

$$\hat{\mathbf{w}}_{\mathrm{ML}} = \max_{\hat{\mathbf{w}}} \left\{ \Pr\left[\hat{\mathbf{w}}|\mathbf{d}\right] \right\}. \tag{5.1}$$

If some prior information on the distribution of $\mathbf{w}$ is available, this can be included in the estimation of $\hat{\mathbf{w}}$ using a maximum a posteriori (MAP) estimator. This is obtained from (5.1) using Bayes' theorem and taking logs, giving

$$\hat{\mathbf{w}}_{\mathrm{MAP}} = \max_{\hat{\mathbf{w}}} \left\{ \log\left(\Pr\left[\mathbf{d}|\hat{\mathbf{w}}\right]\right) + \log\left(\Pr\left[\hat{\mathbf{w}}\right]\right) \right\}. \tag{5.2}$$

This is a MAP (maximum a posteriori) estimator, with $\Pr[\hat{\mathbf{w}}]$ incorporates prior information on the path probability. The path probability can be calculated from known statistics of the platform motion as shown in Section 5.4.

The log-likelihood of the data, given a motion estimate, also needs to be determined. The measured data $\mathbf{d}$ and motion estimate $\hat{\mathbf{w}}$ are first combined into an image estimate $g(x,y)$. This comprises of a step of *motion compensation* and then *image reconstruction*. The processing involved is different for spotlight and stripmap imagery and is discussed further in Chapter 2. Since uncompensated platform motion blurs this image, the sharpness of the reconstructed image can be used to estimate the image's log-likelihood. Methods to measure image sharpness are shown in Section 5.3.

5.2.1 Statistical autofocus algorithm

Statistical autofocus is summarised in Fig. 5.1. The measured data $\mathbf{d}$ is corrected with the current path estimate $\hat{\mathbf{w}}$, and an image estimate $g(x,y)$ is formed. A sharpness metric is measured from the image. This is combined with the cost of the path estimate, to obtain a cost measure using (5.3). This measure is fed into an optimisation routine programmed to minimise this cost measure by perturbing a set of motion parameters $\mathbf{b}$, representing the path estimate $\hat{\mathbf{w}}$. Methods to parameterise the platform motion are discussed in Section 5.5.

The cost function is calculated using;

$$\begin{aligned} \text{Cost} &= -\left(\log(\Pr[\mathbf{d}|\hat{\mathbf{w}}]) + \log(\Pr[\hat{\mathbf{w}}]) \right) \\ &\approx -\left(\log(\Pr[g(x,y)]) + \eta \log(\Pr[\hat{\mathbf{w}}]) \right), \end{aligned} \tag{5.3}$$

where η is a constant balancing the effect of the two components. The effect of varying η is shown in Section 5.4. The cost is minimised by a multidimensional optimisation algorithm. This is discussed in Section 5.6.
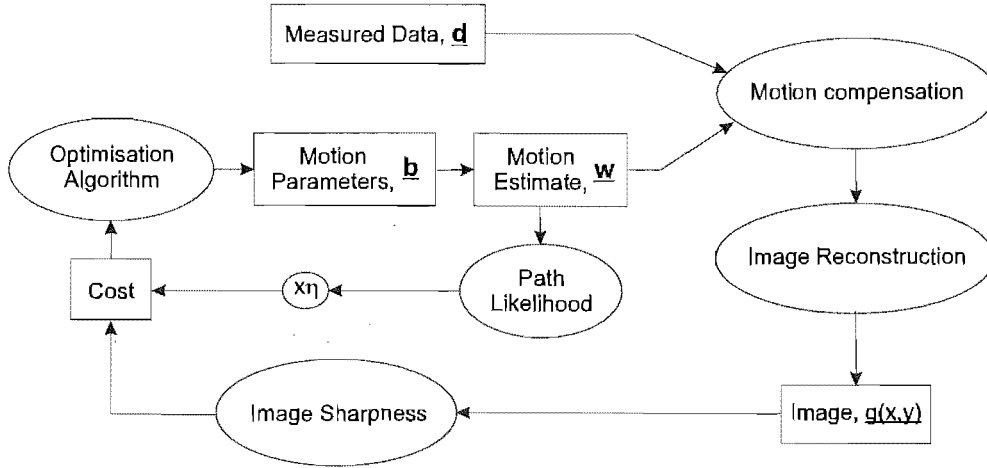
Figure 5.1 Statistical autofocus algorithm. The optimisation algorithm varies the motion parameters to minimise the cost.

5.3   MEASURING IMAGE SHARPNESS

This section describes a number of different methods of measuring image sharpness. The most successful sharpness measures from different fields are compared and the different functionality of different sharpness metrics is explained. The effect imaging errors have on sharpness is examined as well as the reason for, and ways to reduce, oversharpening. The different sharpness metrics are compared qualitatively and quantitatively on different scenes in Section 5.7.

5.3.1   Standard Sharpness Measures

Optical images are normally intensity only, but SAR and SAS have complex measurements. To be consistent, sharpness will shown as a function of intensity

$$I(x,y) = |g(x,y)|^2. \tag{5.4}$$

Maximum Intensity

A simple metric is the maximum intensity in the image,

$$S_{\max} = \max[I(x,y)]. \tag{5.5}$$

In astronomical imaging, it makes poor use of photon counting statistics but is satisfactory for bright objects [Muller and Buffington, 1974]. In coherent imagery, it has high variance due to speckle and is prone to oversharpening [Morrison, 2002; Fienup and Miller, 2003].

Power Law

A better approach is to use all pixels in the image. A popular metric is

$$S_\beta = \iint I^\beta(x,y)\,dx\,dy. \tag{5.6}$$

For optical images, it is proved to give maximum sharpness when the Fourier phase error is zero. [Muller and Buffington, 1974][Hamaker et al., 1977]. $S_{\beta=2}$, shortened to S_2, is the classic sharpness metric [Paxman and Marron, 1988; Fienup, 2000; Gough and Lane, 1998], generalised to S_β by Fienup and Miller [2003]. For $\beta < 1$, the sharpest image is obtained by minimising S_β, which is equivalent to maximising $-S_\beta$ or $1/S_\beta$.

Standard Deviation of Amplitude or Intensity

A commonly used metric is the ratio of the standard deviation to the mean of the amplitude

$$S_{va} = \frac{\sqrt{\iint [|g(x,y)| - \iint |g(x,y)|dx\,dy]^2 dx\,dy}}{\iint |g(x,y)|dx\,dy} \tag{5.7}$$

or intensity

$$S_{vi} = \frac{\sqrt{\iint [I(x,y) - \iint I(x,y)dx\,dy]^2 dx\,dy}}{\iint I(x,y)dx\,dy}. \tag{5.8}$$

Due to preservation of energy, the sum of intensity over an image should remain constant with varying Fourier phase error. Thus S_{vi} is a scaled and shifted version of S_2. Tests show they have similar performance [Fienup and Miller, 2003]. It is a popular sharpness metric and is used by many authors [Finley and Wood, 1985; Blacknell and Quegan, 1991; Blacknell et al., 1992; Berizzi and Corsini, 1996; Berizzi et al., 1996].

Similarly, S_{va} has the same performance as $1/S_{\beta=0.5}$. It is not as commonly used [Berizzi and Corsini, 1996]. Another metric involving amplitude not intensity is

$$S_I = \frac{\iint |g(x,y)|^2 \, dx\,dy}{[\iint |g(x,y)| \, dx\,dy]^2}. \tag{5.9}$$

This is also equivalent to using $1/S_{\beta=0.5}$.

Entropy

The negative entropy of the image is given by

$$S_{ent} = \iint I(x,y)\ln(I(x,y))\,dx\,dy. \tag{5.10}$$

It has been used successfully to focus ISAR [Xi et al., 1999] and SAR images [Morrison and Munson, 2002; Morrison, 2002; Fienup and Miller, 2003].

Reference Images

The metrics

$$S_{r1} = \iint I_0(x,y)I(x,y)\,dx\,dy \tag{5.11}$$

and

$$S_{r2} = \iint |I(x,y) - I_0(x,y)|^n\,dx\,dy \tag{5.12}$$

compare the image intensity to a standard reference image I_0 [Muller and Buffington, 1974]. In SAR or SAS imaging such a reference is not known ahead of time, so these are not useful metrics but mentioned for completeness. S_{r1} has been found useful in real-time optical correction, as it can be computed optically [Buffington et al., 1977].

5.3.2   Normalisation

Metrics can have arbitrary scaling. When comparing metrics, it is useful to have them scaled to the same range. Fixed tolerance levels within optimisation routines require metrics to be scaled to the same numerical range for valid comparisons. To use sharpness as an estimate of image likelihood as in (5.3), sharpness should be scaled to between zero and one. For a simulated image, the correct image is known thus the sharpness can be scaled to one for the correct image. If the correct image is unknown, the maximum sharpness is scaled to one. In addition, the minimum sharpness is scaled to zero. This can be done using the scaling

$$S_N(\hat{\mathbf{w}}) = \frac{S(\hat{\mathbf{w}}) - S^-}{S^+ - S^-}, \tag{5.13}$$

where S^+ is the sharpness of the correct image, or, if that is unknown, the maximum sharpness, and S^- is the minimum sharpness. A similar normalisation process was used by Fienup and Miller [2003].

For blind optimisation, the maximum and minimum sharpness values are unknown before the process begins. The minimum sharpness S^- can be estimated by measuring the sharpness of an image reconstructed with a large motion, or Fourier phase error. The maximum sharpness S^+ can be set by the initial image. The scaled sharpness will be higher than one as the image improves. It is important to ensure that $S^+ > S^-$. If the estimate of S^- is poor, and $S(\hat{\mathbf{w}}) < S^-$, then $S_N(\hat{\mathbf{w}}) < 0$.

Statistical autofocus minimises a cost function which includes a measure of the log-likelihood of the image (5.3). The log-likelihood of the image is estimated by

$$\log(\Pr[g(x,y)]) \approx \log(S_N(\hat{\mathbf{w}})). \tag{5.14}$$

The logarithm is a monotonic function for positive values. Thus maximising $S_N(\hat{\mathbf{w}})$ is the same as maximising $\log(S_N(\hat{\mathbf{w}}))$, so long as $S_N(\hat{\mathbf{w}}) > 0$. To ensure this, if $S_N < \epsilon$, where ϵ is a small positive number, set $S_N = \epsilon$. This also removes problems with the logarithm of zero. The path-cost component of (5.3) ensures that the gradient of the cost function is not zero if this occurs.

5.3.3 Effect of different sharpness metrics

It is useful to consider a generalised sharpness function $\Omega[I(x,y)]$, which is a nonlinear function of the intensity of a single image pixel. The sharpness of an image is the sharpness function averaged over all pixels:

$$S_\Omega = \frac{1}{N_x N_y} \sum_{x,y} \Omega[I(x,y)]. \tag{5.15}$$

Most of the metrics defined in Section 5.3.1 can be defined in this manner. For example, S_β uses the sharpness function

$$\Omega_\beta[I] = I^\beta. \tag{5.16}$$

A key to understanding the effect of different sharpness metrics is the nature of the second derivative of the sharpness function [Fienup and Miller, 2003],

$$\Omega''[I] = \frac{\partial^2 \Omega}{\partial I^2}. \tag{5.17}$$

Fig. 5.2 shows a plot of a sharpness function with a positive second derivative. Consider
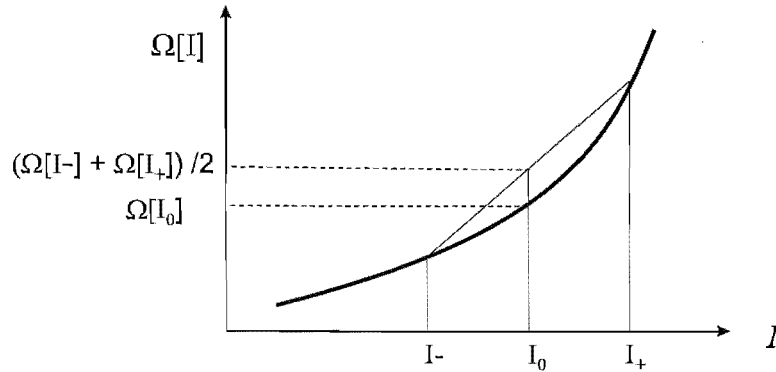
Figure 5.2 Sharpness function of intensity. If the second derivative is positive, the sharpness will increase as the values of I spread out.

two pixels in an image with the same intensity, I_0. The original sharpness is $S(1) = (\Omega[I_0] + \Omega[I_0])/2 = \Omega[I_0]$. Consider a change in phase estimate that causes an intensity change of ΔI to be transfered from one pixel to the other. The intensity of the two pixels is now $I_+ = I_0 + \Delta I$ and $I_- = I_0 - \Delta I$, so the sum of intensity remains unchanged (as required by conservation of energy). The new sharpness value is $S(2) =$

$(\Omega[I_+] + \Omega[I_-])/2$. By a second order Taylor approximation about I_0, this can be approximated by $S(2) \approx \Omega[I_0] + \Omega''[I_0]\Delta I^2/2 \approx S(1) + \Omega''[I_0]\Delta I^2/2$. Thus as the intensity values spread out, the sharpness value increases proportional to the second derivative of the sharpness function, $\Omega''[I_0]$. This can be seen in Fig. 5.2. The larger the curvature in the sharpness function, the larger the increase in sharpness as intensity values spread out.

Different sharpness functions affect how the second derivative changes with intensity. The sharpness metric has a larger effect on the brightest points if the second derivative goes up with increasing intensity, or has a larger effect on shadows if the second derivative is highest for a lower intensity. Fig. 5.3 shows several sharpness func-
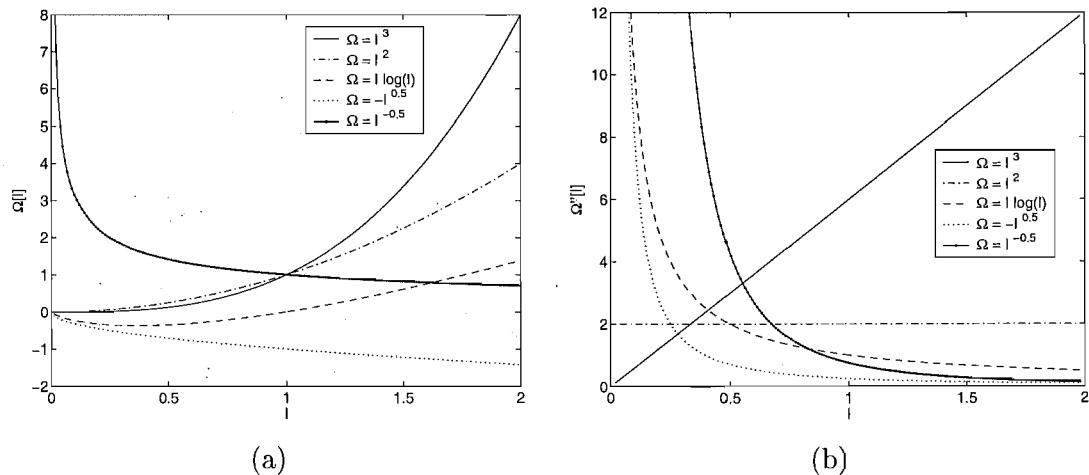
Figure 5.3   (a) Examples of different sharpness functions and (b) their second derivatives.

tions and their second derivatives. This shows the sharpness function $\Omega = I^3$ will emphasise bright points, as the second derivative is highest for a high intensity. Conversely, using $\Omega = -I^{0.5}$ will emphasise shadows, as the second derivative is highest for points of low intensity. This has been confirmed in trials, (Section 5.7). Notably, the classic sharpness metric, $\Omega = I^2$ has constant curvature with intensity, thus it treats shadows and bright points the same.

5.3.4   Effect of imaging errors

The image reconstruction process can introduce errors into the image, which can affect the sharpness measurement. For SAS, exact reconstruction is a slow, intensive process. It is impractical if a scene is reconstructed many times as in sharpness maximisation. If some approximations are made, significantly faster Fourier domain reconstruction techniques such as the wavenumber algorithm can be used [Hawkins, 1996].

A pivotal step of the wavenumber reconstruction is the Stolt transform, which involves interpolating the data onto another grid in the frequency domain. The inter-

polator used is important [Jakowatz *et al.*, 1996; Li, 1992], as an error in the frequency domain spreads an error all over the image. This has a large effect on contrast. For example, using a cubic interpolator on data sampled close to Nyquist rate causes large errors as shown in Fig. 5.4. Fig. 5.4(a) shows how errors in the image vary as the path estimate is changed. This causes variations in the sharpness measurement which are larger than the variation due to smearing of the target, making sharpness maximisation impossible. Fig. 5.4(c) shows the nature of the error in the image at one of the spikes in sharpness.
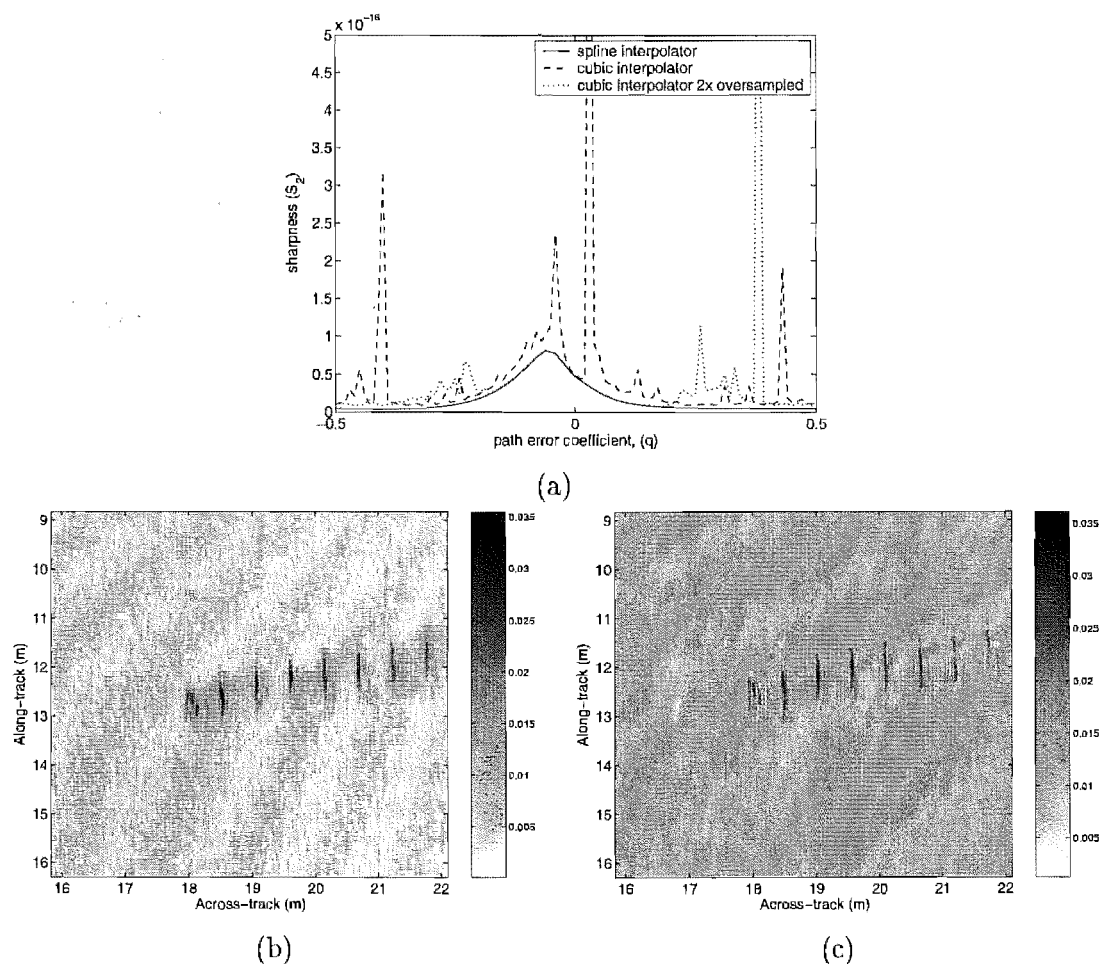
Figure 5.4 Effect of using different interpolation methods in the wavenumber reconstruction algorithm on field SAS image. (a) Effect on sharpness of field image as path error coefficient (q) is changed. (Path is sinusoidal with a 15 m period, q is rms error.) (b) Portion of image reconstructed using spline interpolation. (q=0.05) (c) Portion of image reconstructed using cubic interpolation. (q=0.05)

There may be other artefacts present in the image. At close range there is a strong return due to cross-talk between the transmitter and receiver in all KiwiSAS images. This needs to be removed before the sharpness is measured, so it does not bias the sharpness. This was achieved by cropping it out. Another artefact present in images are target sidelobes. These can be reduced by windowing during image reconstruction.

5.3.5   Oversharpening

It is possible for some images to have a higher sharpness measure with a particular path estimate than that of the ideal image with the correct path. One example of over-sharpening is shown in Fig. 5.5(b). The phase correction that maximises the sharpness makes the phase of the brightest range-bin linear, creating a single brighter point out of the two points in the same range-bin. Other range bins may be blurred and have low sharpness measures, but if this point is significantly brighter than the others, the whole image may have a higher sharpness measure than having each point correctly focused. In this case, a sharpness maximisation algorithm will give a poor result.

Some sharpness metrics are more likely to cause oversharpening than others. Those with emphasis on making a single point the brightest such as $S_{\max}$ or S_β for a large β are more likely to cause oversharpening than those with higher second derivative for lower intensity. Shadows cannot be produced by oversharpening, making sharp shadows a good indication of correct focus. This is one reason why Fienup found entropy and power-law metrics with low powers had better performance in trials on spotlight SAR [Fienup and Miller, 2003].

One way to reduce any oversharpening is to use the fact that oversharpening only occurs at a single dominant range bin, while the others are blurred. A range dependent weighting function can be used to reduce the dominance of brighter range bins [Fienup, 2000], using

$$S_\Omega = \frac{1}{N_x N_y} \sum_{x,y} W(x)\Omega[I(x,y)] \tag{5.18}$$

The weighting function $W(x)$ can be used to weight each range bin depending on how useful it is for phase-error estimation [Fienup, 2000]. The energy of each range bin can be normalised using the weighting function

$$W(x) = \frac{1}{[\sum_y I(x,y)]^2}. \tag{5.19}$$

Using this weighting, the brightest range bin will dominate less over the others. Fig. 5.5(c) shows that for this image, using a sharpness measure normalised in range reduced over-sharpening. Range weighting is discussed further in Section 6.2.2.

An oversharpened image often has a rapidly changing estimated phase as shown in Fig. 5.5(d). Although this may give a *sharper* image than the aberration free image, the estimated path can be judged to be less likely. By penalising unlikely paths, the correct path is more likely to be estimated. The next section develops a method for measuring the likelihood of the path and combining it with sharpness to reduce oversharpening.

(a)

(b)

(c)

(d)
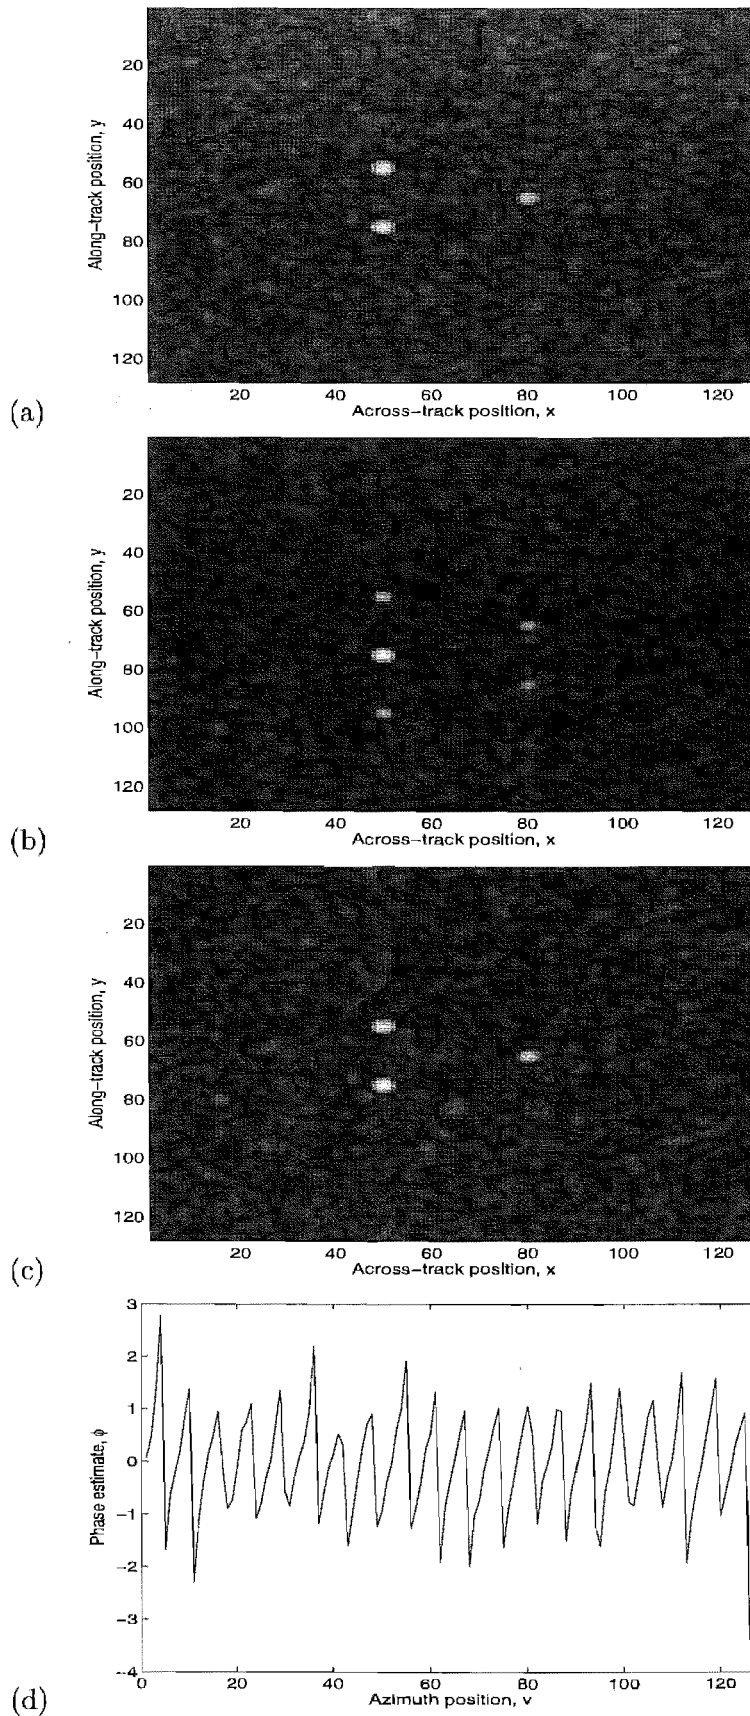
Figure 5.5 Example of oversharpening in simulated spotlight image. (a) Original image consisting of three points. (b) Image with maximum sharpness S_2. The two points in the same range-bin are oversharpened, smearing the third point. (c) Image with maximum range normalised sharpness $\overline{S_2} = \sum_x S_2/[\sum_y I(x,y)]^2$. There is now no oversharpening. (d) Phase estimate giving maximum sharpness S_2.

5.4 PATH PROBABILITY

A slant range sway is the motion that has the largest effect on blurring a single receiver SAS image [Johnson *et al.*, 1995]. A method for calculating the probability of an estimated sway will be shown. The same principle will then be extended to other motions. Before this step, any prior information on the path, for example INU data, or a non-coherent shear average estimate, should be corrected in the data.

5.4.1 Sway

To estimate the log-likelihood of an estimated sway $\hat{w}(u)$ at each ping u, one needs to measure, or make assumptions about, the statistics of the actual sway $w(u)$. This can be measured from previously collected data or platform measurements. Statistical autofocus is not highly sensitive to these statistics, so an accurate model is not essential. For results presented, the sway has been modelled with two parameters, the variance and the correlation length. This model appears sufficient, but if data is available, a more accurate model could be formed.

The variance is estimated from measurements as

$$\text{Var}\,[w(u)] = \sigma_w^2. \tag{5.20}$$

The autocorrelation of the sway, $R_w(\tau) = E\,[w(u+\tau)w(y)]$, can be estimated from measurements. Presented results have assumed a Gaussian correlation, with correlation length κ_w, i.e.,

$$R_w(\tau) = \sigma_w^2 \exp\left(-\pi(\tau/\kappa_w)^2\right). \tag{5.21}$$

Any constant offset will not blur the image, so $w(u)$ can be assumed zero mean. The correlation will then be equal to the covariance.

Let the estimated sway at each along-track position form a column vector $\hat{w}$. The covariance of this vector is a matrix given by

$$\mathbf{R_w} = E\,[\mathbf{w}\mathbf{w}^T] = \begin{bmatrix} R_w(0) & R_w(1) & \ldots & R_w(n-1) \\ R_w(1) & R_w(0) & \ldots & R_w(n-2) \\ \vdots & \vdots & \ddots & \vdots \\ R_w(n-1) & R_w(n-2) & \ldots & R_w(0) \end{bmatrix} \tag{5.22}$$

assuming $w(u)$ is stationary. The log-likelihood of the path depends on the distribution of $w(u)$. Assuming Gaussian distribution, the log-likelihood of an estimated path is given by [Beck and Arnold, 1977],

$$\log[\Pr\{\hat{w}(v)\}] = -\frac{1}{2}[\hat{\mathbf{w}}^T\mathbf{R_w}^{-1}\hat{\mathbf{w}} + N\log(2\pi) + \log|\mathbf{R_w}|]. \tag{5.23}$$

Only the first term of the above expression varies with the path estimate. A maximum likelihood estimator can ignore the second two terms and use

$$\log[\Pr\{\hat{w}(v)\}] = -\frac{1}{2}\hat{\mathbf{w}}^T \mathbf{R_w}^{-1}\hat{\mathbf{w}}. \tag{5.24}$$

5.4.2 Effect of path cost weighting

The path-likelihood acts as a form of regularisation by penalising large or rapid movements of the platform. Fig. 5.6(a) shows the log likelihood surface for two path parameters. The greater the movement from a straight line, or the more rapid the movement, the lower the probability. Fig. 5.6(c) shows an example of the path-cost combined with the image likelihood measure. In areas far from the peak, the sharpness measure is flat, making it difficult to maximise. The path-cost adds gradient to these areas, making a successful maximisation more likely, but not moving the position of the peak.

Increasing the size of η changes the weighting of the path cost. This will generally result in a smoother estimated path. This can be seen in Fig. 5.7. Increasing η results in a smoother phase estimate. In areas of low coherence, rather than give a random phase, the regulated estimate will give a smooth phase close to zero. This can be observed at the edge of the phase estimate for $\eta = 0.05$ in Fig. 5.8(e). The other effect of regularisation is the reduction of oversharpening. This can be seen in Fig. 5.8. The unregulated sharpened image ($\eta = 0$) suffers from oversharpening, with the double point target turned into a single bright point with sidelobes. Increasing the weight of the path-cost η has reduced the oversharpening. This is due to any rapid changes in phase being penalised by the path-cost component.

For a cascaded non-parametric optimisation (see Section 5.6), the phase at each azimuth position is adjusted one position at a time. Regularisation will penalise any movement, as it is out of line with all the other positions. Thus the phase can only change slowly. Many iterations are required to converge. This can be seen in Fig. 5.8(f). The larger the path-cost weighting η, the slower the convergence. The sharpness did not converge for $\eta = 0.02$ in 500 iterations, which is extremely slow for an image consisting of 128 pings. Due to this slow convergence, regularisation does not work well with a cascaded non-parametric optimisation method. An improved parametric method is developed in Section 5.5.

The performance of SAF for different path cost weightings η and different sharpness metrics on different images are evaluated in Section 5.7.

5.4.3 Other motions

Any motion parameter that degrades image sharpness and can be corrected in the image reconstruction process can be estimated by SAF. For example, mean platform

(a)

(b)

(c)
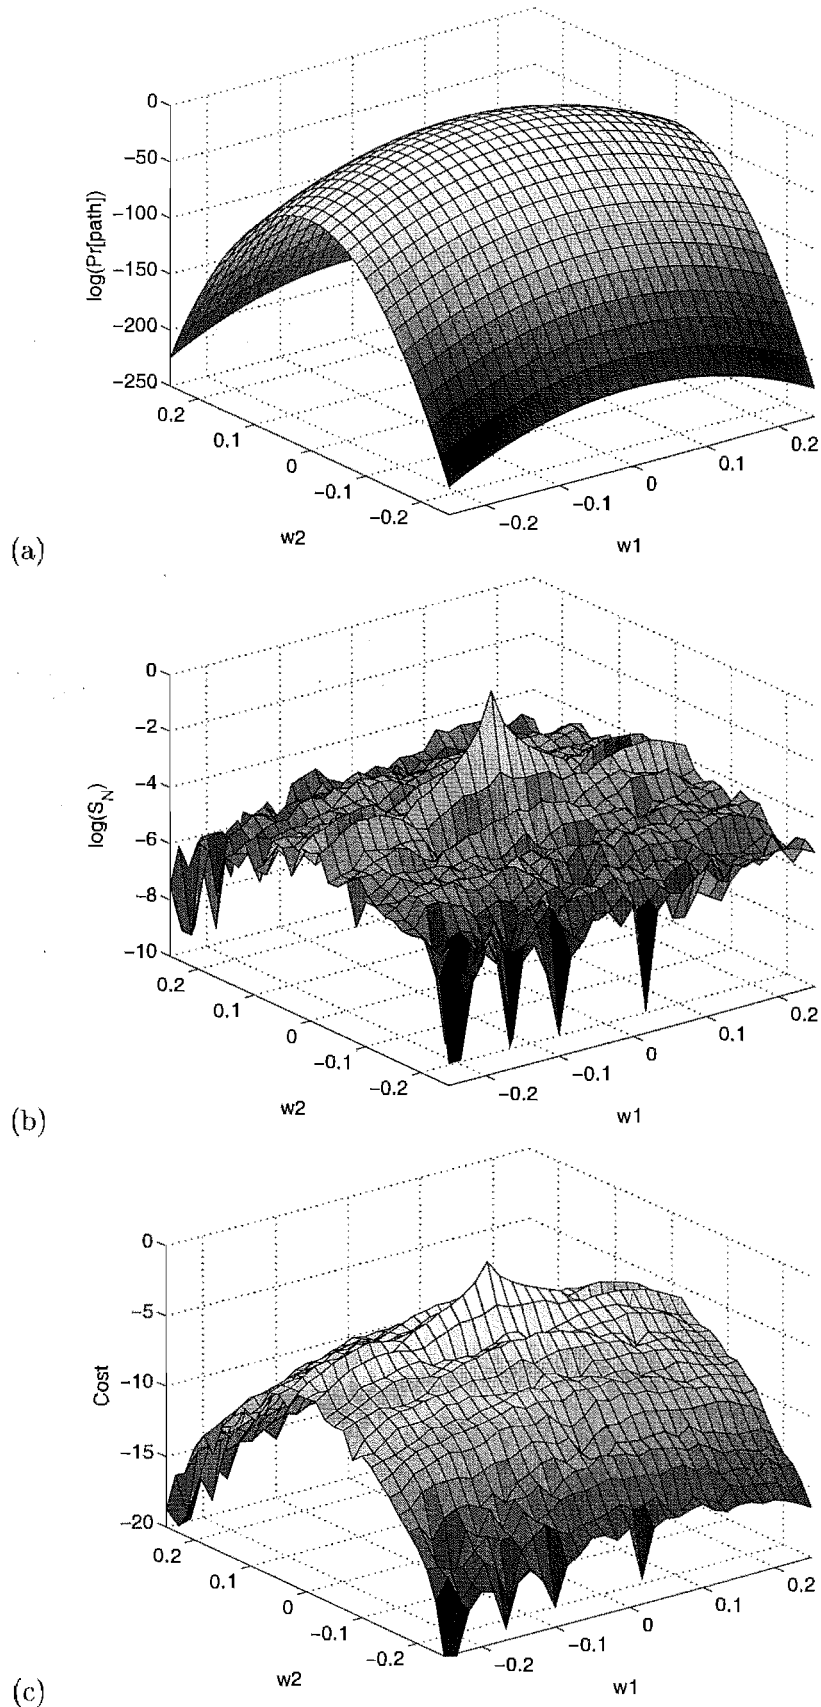
Figure 5.6 Effect of including path probability in cost function. The path is comprised of two parameters, set by path $= w_1 \sin(2\pi f_1 y) + w_2 \sin(2\pi f_2 y)$, with $f_1 = 1/4\,\mathrm{m}^{-1}$, $f_2 = 1\,\mathrm{m}^{-1}$. (a) Log likelihood of path. (b) Log of normalised sharpness of test image. (Simulated SAS image of mine-like objects). (c) Negative cost, -cost $= \log(S_N) + \eta \log(\Pr[\mathrm{path}])$ for $\eta = 0.05$.
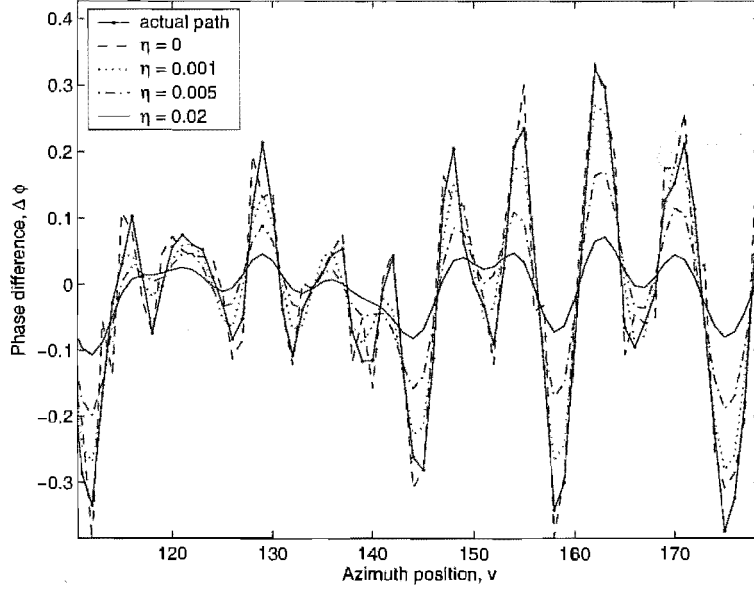
Figure 5.7 Effect of regularisation on phase difference estimate. Sharpness metric used is negative entropy, S_{ent}. Increasing regularisation reduces phase-difference estimate, smoothing phase estimate.

velocity, and variation around that mean (surge). In multiple-hydrophone SAS, yaw is an important motion parameter as it significantly degrades the quality of the imagery [Douglas and Lee, 1993; Christoff, 1998; Gough and Miller, 2004]. If different motions are independent, they should be estimated separately as this is more efficient. If they are coupled, different motions could be included in a common parameter vector $\hat{w}$ and covariance matrix R_w. For example, intuitively sway and yaw are likely to be coupled. Thus a path estimate that sways one direction, but yaws in the other could be judged less probable than one which the two motions were linked. Benefits may also arise from decoupling the motions with path parameterisation (see following section). This has not been attempted yet, but is an area of promising future work.

5.5 PARAMETERISING THE PLATFORM MOTION

Perturbing each individual position of the path is inefficient since each point is highly correlated with the point next to it. It would be much faster if the path could be parameterised by fewer, independent coefficients. This suggests representing the path as a series of basis functions and only perturbing those functions which have higher mean energy. However, the type and number of basis functions used is dependent on the statistics of the platform motion. For example, a platform which moves rapidly requires more terms to represent its path than one that moves smoothly.

If there are N_y pings, a non-parametric representation of the sway at each ping would require N_y different parameters. Since the sway of adjacent pings is correlated, this would require a full N_y dimensional optimisation. By representing the path by
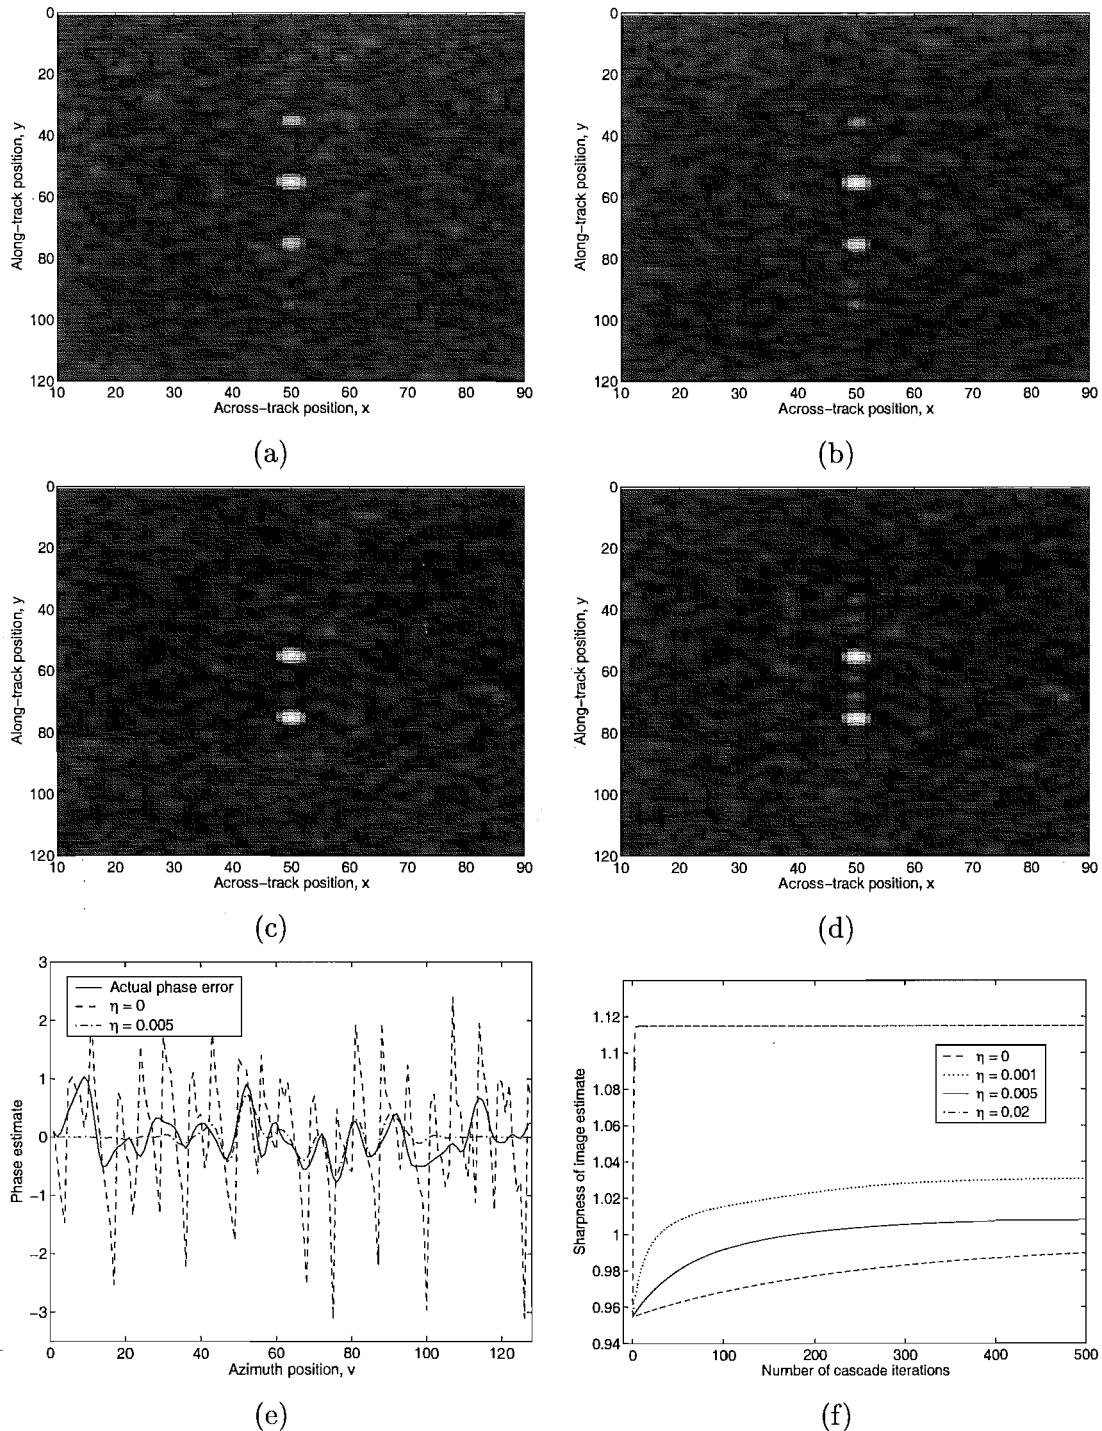
Figure 5.8   Effect of regularisation on oversharpnening. Images blurred by same random phase error, then sharpened using S_2 and different path cost weightings η. Cascaded, non-parametric optimisation is used. Original image is simulated spotlight image of two point targets at (50,55) and (50,75). (a) Sharpened image, $\eta = 0$. Suffers from oversharpening. (b) Sharpened image, $\eta = 0.001$. Oversharpening is significantly reduced. (c) Sharpened image, $\eta = 0.005$. No visible oversharpening. (d) Sharpened image, $\eta = 0.02$. Some residual blurring as sharpness has not yet converged. (e) Phase error estimates. For high coherence (centre of azimuth), regularised sharpness phase estimate is close to actual phase error. In areas of low coherence (edges of azimuth), phase estimate goes to zero. (f) Measured sharpness S_2, normalised by sharpness of original image, at each iteration. Regularisation slows convergence for non-parametric optimisation.

fewer parameters, the search space N is reduced. By parameterising the motion, the cost function is less coupled between parameters. Minimisation will generally require fewer iterations to converge as shown in Fig. 5.9(b).

5.5.1 Karhunen-Loeve decomposition

The optimal set of basis functions for the path is to form a Karhunen-Loeve decomposition. This consists of the eigenvectors of the covariance matrix of the path, $\mathbf{R_w}$ [Jayant and Noll, 1984]. The path $\hat{\mathbf{w}}$ can be represented by

$$\hat{\mathbf{w}} = \mathbf{U}\hat{\mathbf{b}}, \tag{5.25}$$

where $\mathbf{U}$ is assembled from columns of the eigenvectors of $\mathbf{R}$. The parameters $\hat{\mathbf{b}}$ is a column vector of coefficients, obtained using

$$\hat{\mathbf{b}} = \mathbf{U}^{-1}\hat{\mathbf{w}} = \mathbf{U}^{T}\hat{\mathbf{w}}. \tag{5.26}$$

This method allows full resolution representation of an arbitrary path. The parameters $\hat{\mathbf{b}}$ are less coupled than the path parameters $\hat{\mathbf{w}}$. Perturbing the estimates of $\hat{\mathbf{b}}$ will converge in fewer iterations as shown in Fig. 5.9(b). The difference is more significant for regularised sharpness maximisation as shown in Fig. 5.10(b).

5.5.2 Term reduction

To reduce the number of terms used to represent the path, it is possible to use fewer basis functions or fewer columns of $\mathbf{U}$, leading to fewer elements of $\hat{\mathbf{b}}$. The eigenvalues represent the mean-square energy of the corresponding term. By choosing the eigenvectors with the higher mean-square energy, the path can be accurately represented using the fewest terms. To use enough parameters to represent 95% of the energy of the path, use the number of eigenvalues with 95% of the total. Generally this means ignoring the higher order terms, since they tend to have less energy. However, even though high order terms are less likely, the effect they have on image quality is more severe. If a faster estimate of the platform motion is required, high-order terms can be dropped. If a faster reduction of blurring is required, low order terms can be dropped since they have smaller effect on image quality. If high quality imagery is required, all terms should be retained.

5.5.3 Convergence

Fig. 5.9 compares the result of a non-parametric and a parametric optimisation of the sharpness of the same aberrated image. Both methods converge to the same result,

but the parametric optimisation converges in fewer iterations ($\approx$ 4) than the non-parametric method ($\approx$ 20). Fig. 5.10 compares the result of a regularised optimisation with the path cost included. Again both methods converge to the same result, the small difference between the paths is due to the non-parametric method not having reached convergence. The difference between convergence rates is more significant. The parametric optimisation still converges in approximately 5 iterations, but the non-parametric method takes more than 500. Fig. 5.11 shows the convergence of the sharpness of an image consisting of a double point target as shown in Fig. 5.8. This can be compared to the convergence of a non-parametric optimisation in Fig. 5.8(f). This shows that regularisation does not slow the convergence of a parametric optimisation.
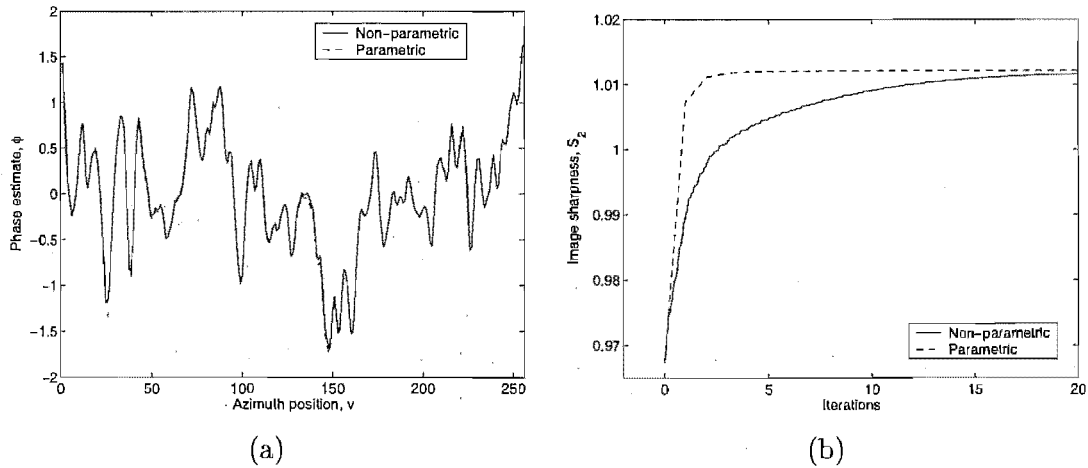
Figure 5.9 Comparison of non-parametric and full resolution parametric optimisation. Non-regularised ($\eta = 0$) cascade optimisation of sharpness measure S_2 of simulated spotlight image. (a) Phase estimates. Almost exact match between methods. (b) Convergence of sharpness S_2, normalised by sharpness of aberration free image.
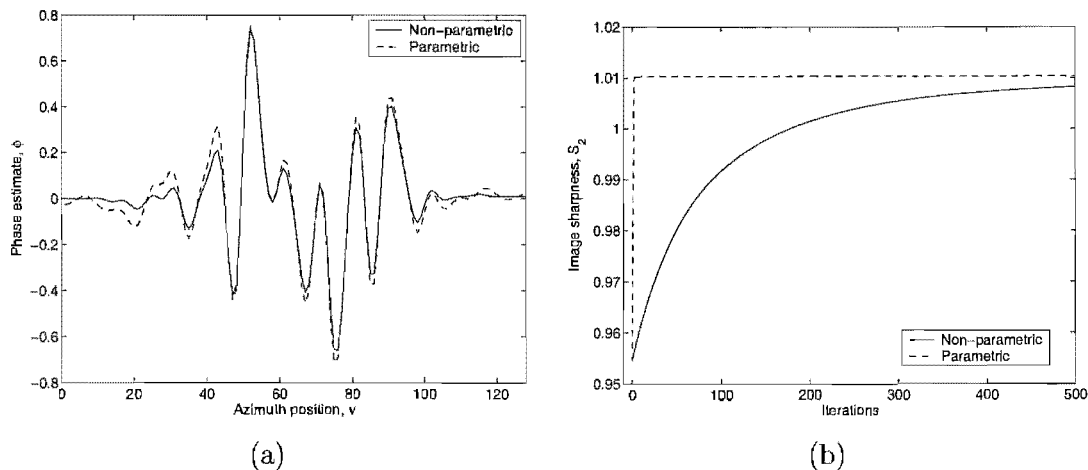
Figure 5.10 Comparison of non-parametric and full resolution parametric optimisation. Regularised ($\eta = 0.005$) cascade optimisation of sharpness measure S_2 of simulated spotlight image. (a) Phase estimates. (b) Convergence of sharpness S_2, normalised by sharpness of aberration free image.
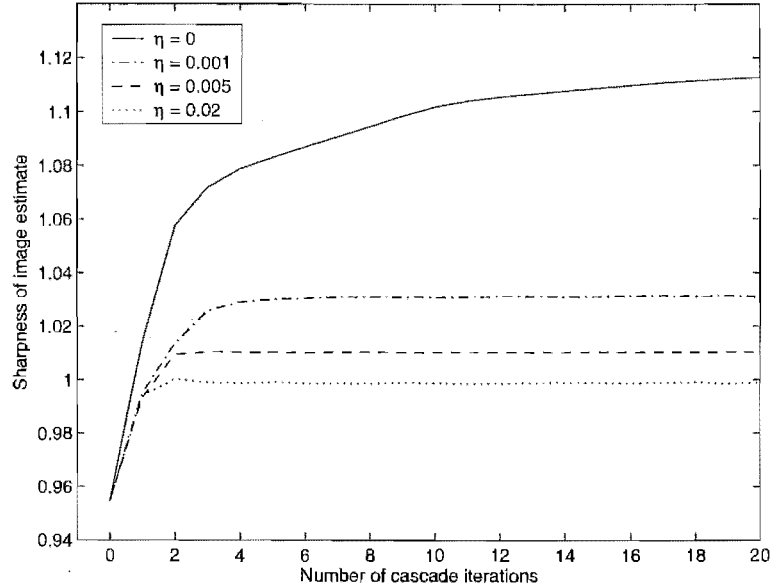
Figure 5.11 Convergence of parametric cascade optimisation for different path-cost weights η. Measured sharpness S_2, normalised by sharpness of original image, at each iteration. Regularisation does not slow convergence of parametric optimisation. Can be compared to results using identical image in Fig. 5.8(f) for non-parametric optimisation.

5.6 OPTIMISATION ALGORITHM

This section describes methods to perturb the motion parameters $\hat{\mathbf{w}}$ or $\hat{\mathbf{b}}$ to minimise the cost function C. The design of an efficient optimisation algorithm is a complex task. It is desirable to use a standard optimisation routine provided by a software package. Provided with a function that can calculate a cost function from a set of parameters, the optimisation routine will attempt to find the parameter values that minimises the cost function. This section will not go into detail of the operation of the various optimisations methods. The different categories of optimisation will be outlined, with how they relate to SAF.

Sharpness measures are subject to noise due to speckle as shown in Section 3.2. Thus an optimisation method that is robust in the presence of noise is important. The optimisation algorithms used will all find a local rather than global minimum. Thus they may not find the true minimum of a non-convex cost function. Algorithms finding a global minimum are possible but at a far higher computational cost.

5.6.1 One dimensional constrained optimisation

The simplest problem is to minimise the function of a single variable within known limits. An example of a single parameter optimisation is the estimation of the velocity, or along-track spacing, of the platform. If there are no firm bounds for the unknown velocity, it is reasonable to give wide bounds that the true velocity is certain to fall

between.
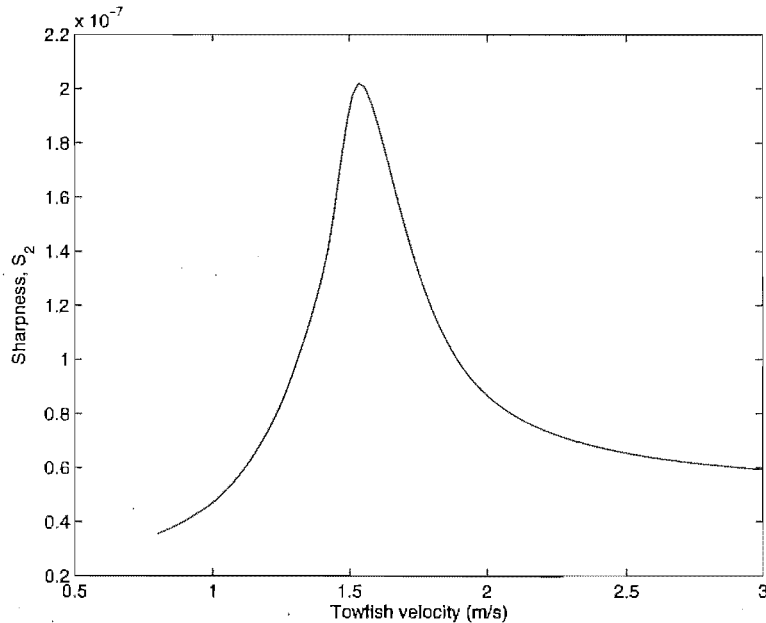
Figure 5.12 Variation of sharpness with velocity used in SAS reconstruction for field SAS image.

Fig. 5.12 shows how the measured sharpness varies with the velocity used in reconstructing a field SAS image. The velocity giving maximum sharpness can be found using a standard maximisation algorithm. Using a golden section search and parabolic interpolation [Forsythe et al., 1976] implemented by MATLAB [MathWorks, 1994], the peak at $v = 1.53$ m/s can be found within 0.01 m/s in 10 iterations.

5.6.2 Multidimensional unconstrained optimisation

For a noisy cost function, the usually recommended method [Elster and Neumaier, 1995; Lagarias et al., 1998; Jarvis, 1997] is the simplex method of Nelder and Mead [Nelder and Mead, 1965]. This method was implemented in MATLAB [MathWorks, 1994]. Tests have shown that the simplex method of Nelder and Mead is more robust that many possibly faster optimisation algorithms in the presence of noise [Jarvis, 1997]. It also requires fewer function evaluations per iteration than other methods [Lagarias et al., 1998]. It is known as a direct search method as gradients are not required. The Elster-Neumaier method [Elster and Neumaier, 1995] has been shown to be an improvement over Nelder-Mead for a noisy cost function.

The search space for an N-dimensional optimisation increases as $O(e^N)$. Thus the computational demands of a multidimensional search rapidly increase for large N. A full resolution optimisation of the sway of a moderately sized image, although possible, requires unreasonable computational demands. The Elster-Neumaier method has a computational order of $O(n^6)$ and is designed for a dimensional limit of $N \le 12$ [Elster

and Neumaier, 1995]. Convergence properties of the Nelder-Mead method are difficult to determine [Lagarias *et al.*, 1998]. There is no strict limit placed on the number of dimensions [MathWorks, 1994], but it is clearly not designed for the several hundred dimensions required for image phase estimation. For a problem of large dimensions, a more efficient method is to break an N-dimensional optimisation into N different 1-dimensional optimisations, then iterate to remove dependencies.

5.6.3 Iterated optimisation

Instead of optimising N parameters simultaneously, another approach is to optimise each parameter independently, resulting in N different 1-dimensional optimisations. If the cost function is not coupled between parameters, both methods will converge to the same solution. If the parameters are coupled, the process can be iterated to converge to the true minimum cost solution if the cost surface is convex.

There are two different schemes for an iterated optimisation. The first is to independently optimise each parameter, update the image estimate, then repeat until convergence. This will be termed *concurrent iterated optimisation*. The second is to optimise the first parameter, update the image with this estimate, optimise the second and so on for all parameters, then repeat. This will be termed *cascaded optimisation*. Fig. 5.13 shows for one example image, the two methods converge to the same result. The cascade optimisation method converges faster, so will be used in preference. SSA is a modified form of cascade optimisation [Xi *et al.*, 1999; Morrison, 2002].

(a)   (b)
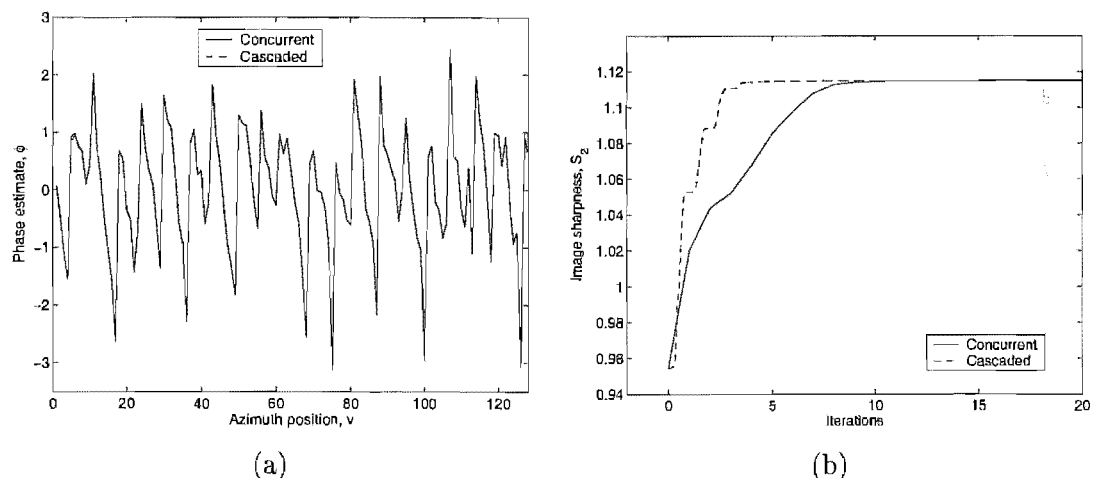
Figure 5.13 Comparison of concurrent and cascaded iterated optimisation methods. Non-regularised ($\eta = 0$) optimisation of sharpness measure S_2 of simulated spotlight image. (a) Phase estimates. Almost exact match between methods. (b) Convergence of sharpness S_2, normalised by sharpness of aberration free image.

Fig. 5.14 compares the result of a cascaded optimisation to a multi-dimensional optimisation using the Nelder-Mead method. This example was performed using 128

parameters. The Nelder-Mead method had not converged after 500000 iterations. It is clear that the cascaded optimisation converges at a faster rate. However, they did not converge to the same estimate. There is a difference in the phase estimate in areas of low coherence. This has little effect on the quality of the image, with both methods giving images of comparable visual quality.

                              (a)   (b)
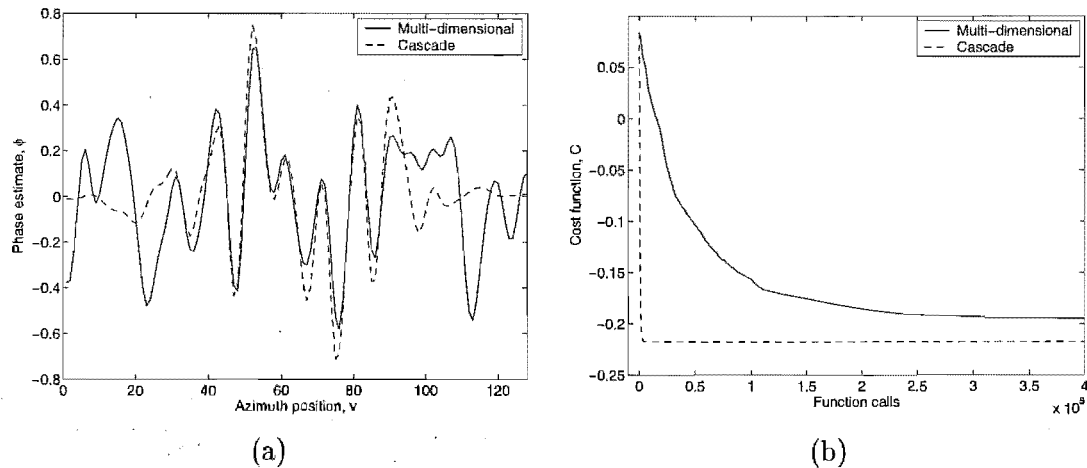
Figure 5.14   Comparison of multi-dimensional and cascaded optimisation methods. Regularised ($\eta = 0.05$) parametric optimisation of sharpness measure S_2 of simulated spotlight image. (a) Phase estimates. The estimates match in areas of high coherence (centre of azimuth) but differ for low-coherence estimates. (b) Convergence of cost function with number of function evaluations. Cascade optimisation is significantly faster. Both methods gave images of comparable quality.

Stopping criteria

A cascaded optimisation can be judged to have converged when the image sharpness, or cost function, does not change substantially between iterations. A set number of iterations can be made, or the optimisation could be stopped when the change in cost falls below a certain threshold between iterations.

5.6.4   Conjugate gradient search

For sharpness maximisation of a spotlight image, it is possible to determine an analytic measure of the gradient of the sharpness with respect to the phase estimate. The sharpness gradient, and its derivation, is shown in Appendix B. This allows, for a particular path estimate, not only a measure of the sharpness but the sharpness gradient. This allows the use of the highly efficient conjugate-gradient optimisation algorithm [Gough and Lane, 1998; Fienup, 2000]. A conjugate gradient search allows the use of a parametric path representation and range weighting [Fienup, 2000] but not regularisation.

No analytic measure of the sharpness gradient has been determined for a stripmap image. Thus a conjugate gradient search can only be used on a spotlight image.

5.6.5 Local minima

All the optimisation methods discussed so far could converge to a local rather than a global minima if the cost surface is non-convex. Fig. 5.15 shows an example of the sharpness surface when two parameters are varied. The same image is treated as a spotlight or stripmap image and the sharpness measured as the two parameters are varied. The surfaces shown are typical for the different imaging modes. The spotlight surface is well behaved. It is truly convex, with no dependence between the parameters. The peak is rounded with the gradient steepening further from the peak. The stripmap surface has different characteristics. There is some dependence between the parameters, visible as an asymmetry in the surface. Noise in the surface is more apparent. The peak is sharp, with the gradient reducing further from the peak. The stripmap sharpness surface is normally more difficult to optimise. If an estimate is located some distance from the peak, it is possible to be trapped in a local minima.

For spotlight sharpness maximisation, local minima do not appear to be as large a problem. Thus there is no advantage in using a slower global optimisation method. Stripmap imagery does not have a strictly convex cost surface when some distance from the peak. The optimisation algorithm could feasibly converge at a local minima and result in a poor estimated imagery.

There are several possible approaches to reduce the chances of converging to local minima. The algorithm could be restarted at a different starting point. A class of algorithm known as *stochastic* or *genetic algorithms* [Reeves and Rowe, 2003] could be employed. The most well known technique is simulated annealing [Otten and van Ginneken, 1989; Zomaya, 2001], where random steps of a reducing size are taken. A useful method is to choose a starting point close to the true peak. A fast (but less accurate) bulk-error removal algorithm, such as noncoherent shear-average [Callow *et al.*, 2001b], or phase-curvature estimation (Section 4.5) can be used to give an initial estimate in the optimisation algorithm. Stripmap sharpness maximisation is not an efficient method for bulk-error removal but is a useful tool for obtaining high-quality imagery if other techniques are insufficient.

5.7 RESULTS

To evaluate different sharpness measures, first a qualitative look at how different sharpness measures change with motion parameters is made. This is used to narrow the field of sharpness measures, which are then evaluated in more depth. A quantitative assessment of sharpness maximisation using different sharpness functions and path cost weights is made on different images. This adds to the comparison of metrics made by Fienup and Miller [2003]. In addition, stripmap imagery is included and regularised sharpness evaluated.

(a)

(b)

(c)
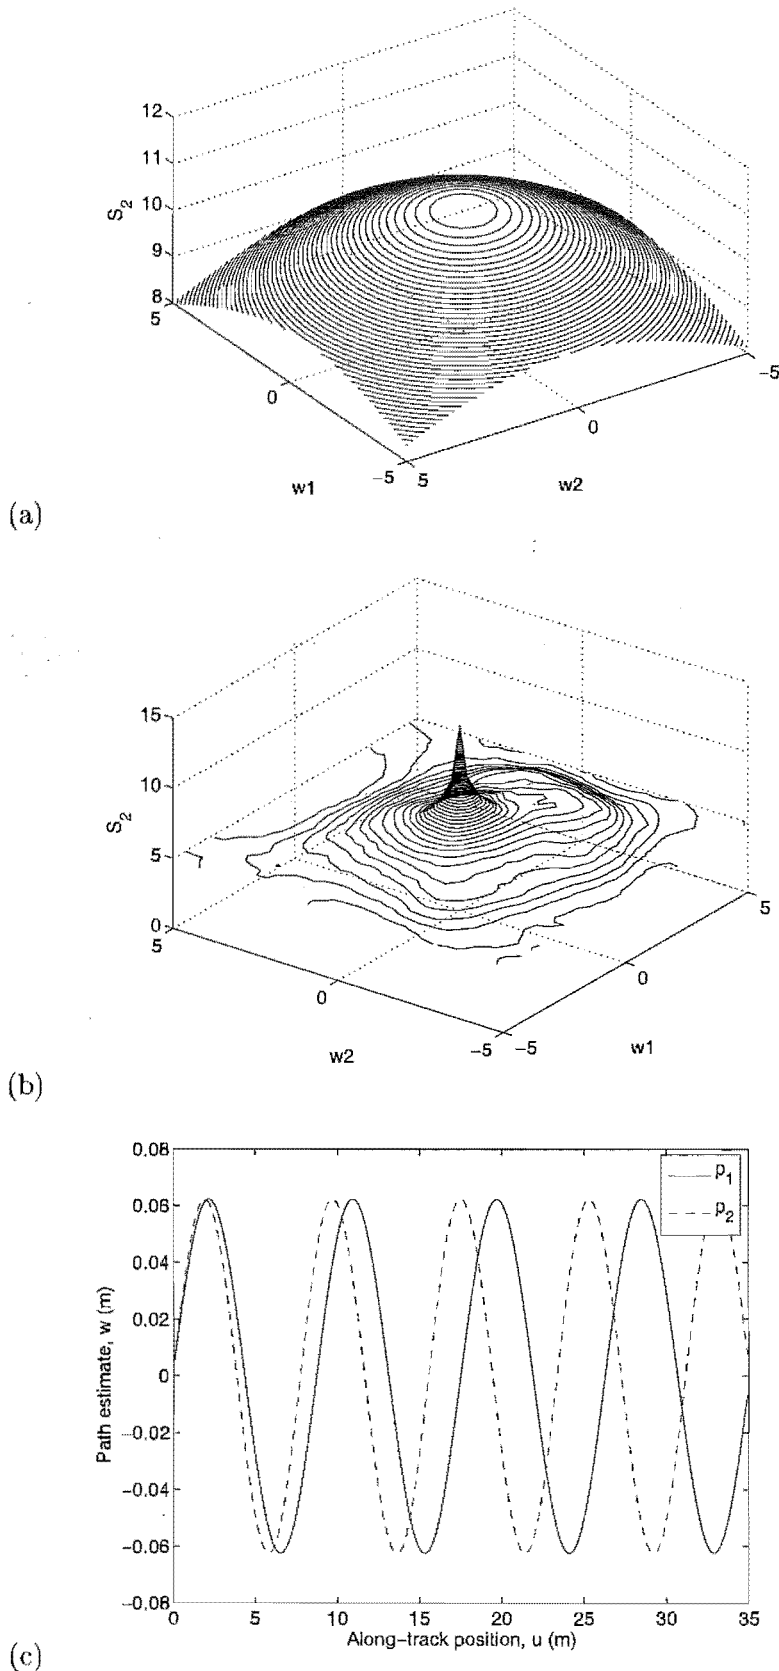
Figure 5.15   Sharpness surface over two parameters. Image is scene D in Fig. 5.17. Path is given by $w(u) = w_1\mathbf{p_1} + w_2\mathbf{p_2}$. (a) Spotlight image. (b) Stripmap image. (c) Path parameters $\mathbf{p}$.

5.7.1 Sharpness metric comparison

Some sharpness measures can be clearly shown to be unsuitable for statistical autofocus. A simple test was performed which varied a single path-error parameter for a field SAS image and measured various sharpness metrics of the resulting image. Fig. 5.16 shows an example of one such test. The actual path error is unknown for this image, but most measures peak for $w \approx 0.17$. A desirable sharpness measure will smoothly increase to a peak at close to this value. Those measures that do not smoothly peak can be regarded inappropriate since they will be difficult to maximise correctly.

Fig. 5.16(a) shows the measures that have increasing curvature with decreasing intensity. $S_{-0.5}$ is noisy with many peaks, thus would be a poor measure. $S_{0.5}$ has a smooth peak, so will be a good measure. As expected, S_{va} has similar performance to $S_{0.5}$. S_{ent} is the best looking measure, with a smooth peak, smaller variations, and a sharper peak than $S_{0.5}$ and S_{va}.

Fig. 5.16(b) shows the measures that have increasing curvature with increasing intensity. For S_m or a large β (S_{10}), the sharpness has a sharp peak but is noisy with many peaks. This makes it unsuitable for sharpness maximisation. S_2 has good performance, with a nice, smooth peak. As expected, S_{vi} has similar performance to S_2. S_4 has a sharper peak than S_2 but has large variations, with small extra peaks. From the result of this test (along with many others), the sharpness measures used will be limited to $S_{0.5}$, S_{ent}, S_2, and S_4.

5.7.2 Evaluation

There are two differing methods to measure the success of sharpness maximisation. Performance may be measured as the number of iterations taken to converge to within an acceptable error faster. Alternatively, the residual error could be measured after converging. As speed is not a prime consideration, the second method will be used.

So how can focusing error be measured? Using the sharpness measure itself is not a reliable measure of good focus since in the case of oversharpening, it can give false results. For a simulated image, the actual phase error is known. Thus the estimated phase error can be compared to the known phase error, with linear and constant phase errors removed. The phase-difference has a greater effect on image degradation than phase alone [Callow, 2003], since a linear phase error leads to an image translation. Thus a more useful comparison is the rms *phase-difference* error.

Although this is a useful comparison, it is not the complete picture. It is possible to have a large phase error but a reasonable quality image. It is necessary to look at not only how good the phase estimate is but how good the image produced by the phase estimate is. This can be performed subjectively (with the human eye) or with a metric that compares the image to the original. The subjective measure attempts

(a)

(b)
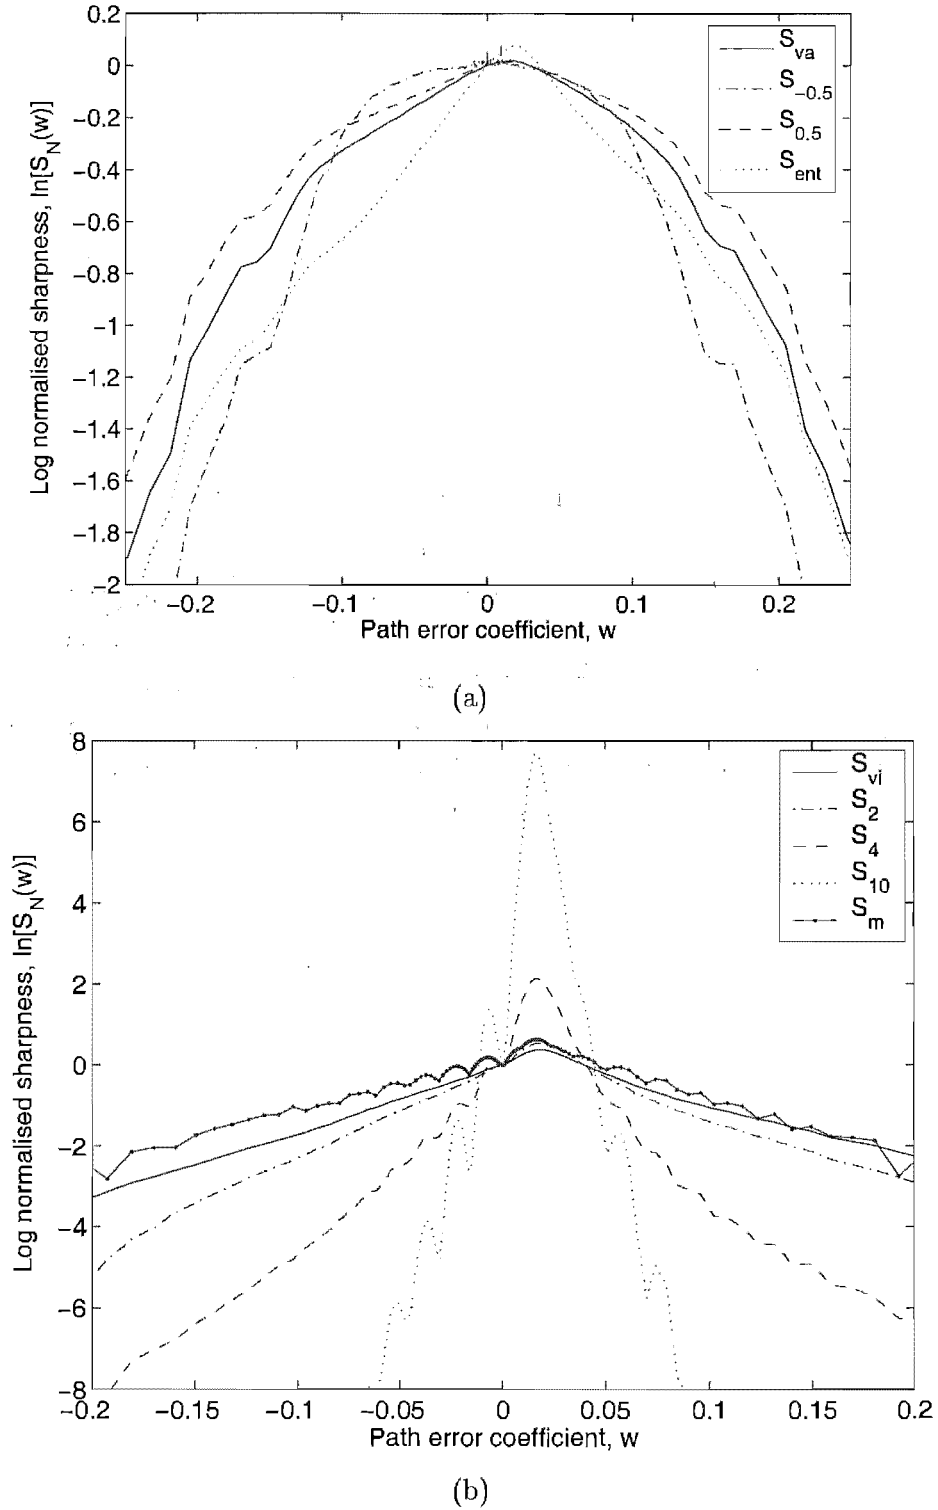
Figure 5.16   Variation of different sharpness metrics when a single path-error parameter is varied. Field SAS image contains a number of bright targets. Path error is sinusoidal, with 5 m period. (a) Sharpness measures with increasing curvature with decreasing intensity. (a) Sharpness measures with increasing curvature with increasing intensity.

to judge the *visual* extent of any blurring. The grades used are defined in Table 5.1. An image comparison measure that is independent of translations of the image and constant phase differences is given by [Fienup, 1997; Fienup and Miller, 2003] as

$$E^2 = \min_{a,x_0,y_0} \left\{ \frac{\sum_{x,y} \left| e^{ja} g(x - x_0, y - y_0) - g_0(x, y) \right|^2}{\sum_{x,y} |g_0(x, y)|^2} \right\}$$
$$= \frac{r_{gg}(0, 0) + r_{g_0 g_0}(0, 0) - 2 \max_{x_0, y_0} |r_{g_0 g}(x_0, y_0)|}{r_{g_0 g_0}(0, 0)}, \tag{5.27}$$

where $r_{g_0 g}$ is the cross correlation of $g_0(x, y)$ the ideal image, with $g(x, y)$, the estimated image and $r_{gg}(0, 0) = \sum |g(x, y)|^2$ is the energy in the image $g(x, y)$. E^2 is a measure of the normalised mean squared difference between the two images, allowing for image translations. The square root of the metric E is the normalised rms error.

Grade	Meaning
1	No blurring visible
2	Blurring difficult to see
3	Small amount of blurring visible
4	Large amount of blurring visible (similar to initial blurred image)
5	Image blurred worse than initial blurred image
#	Modifier. Image is sharp with spurious targets (oversharpened)

Table 5.1 Subjective grading of visual blurring of images. Fractional grades are permitted.

5.7.3 Metric comparison on images

The phase error and best image were estimated using SAF for a range of metrics and path-cost weights η. Results are compared using the rms error of the phase-difference estimate, the rms error of the image using the square root of (5.27) and a subjective measure of the appearance of the image. The subjective measure attempts to judge the visual extent of any noticeable blurring. The grades used are defined in Table 5.1. The original, aberration free images are shown in Fig. 5.17. The results are shown in Table 5.2, averaged over 10 different random initial phase errors.[1]

Comparing the results using different metrics without regularisation, it is clear that results depend on the nature of the image. For point targets (image A), S_{ent}, S_2 and S_4 all give good results with $S_{0.5}$ failing to focus the targets. For an extended speckle target (image C), S_{ent} and $S_{0.5}$ both focus the target well, with high powered metrics S_2 and S_4 doing poorly. For the double point target (image B), both S_2 and S_4 oversharpen the image without regularisation. $S_{0.5}$ fails to focus the points well, leaving S_{ent} as the only measure used that worked well. For the more complex scene

---

[1]The quality grade Q was measured for just a single initial phase error.

(a)

(b)

(c)

(d)
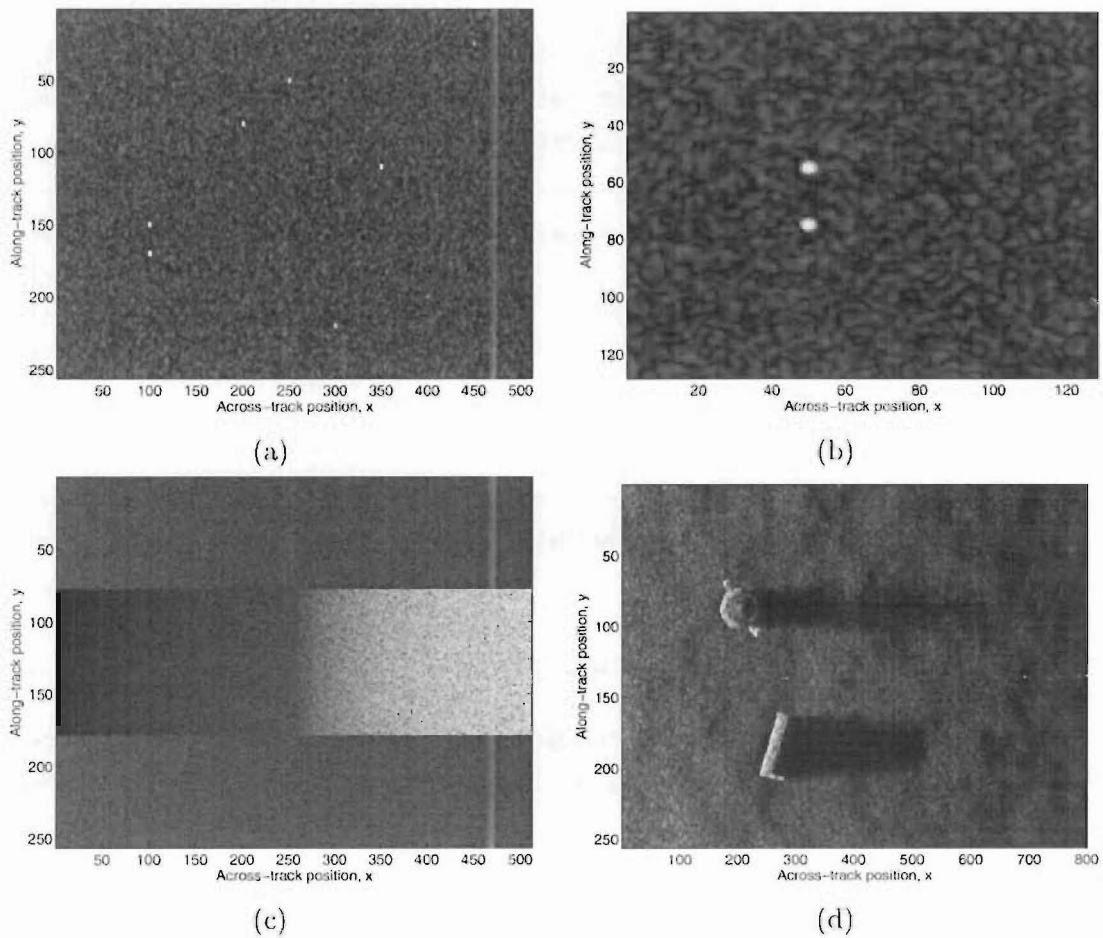
Figure 5.17   Test images. (a) Image A: Scattered point targets. (b) Image B: Double point targets. (c) Image C: Speckled block target and shadow. (d) Image D: Simulated SAS scene of mine-like objects with shadows.

Im.	Met.	$\eta = 0$			$\eta = 0.001$			$\eta = 0.005$			$\eta = 0.02$		
		PDE	IE	Q	PDE	IE	Q	PDE	IE	Q	PDE	IE	Q
A	S_{ent}	0.410	0.275	1	0.074	0.114	1	0.087	0.125	1	0.117	0.199	1
	$S_{0.5}$	1.260	0.757	5	0.332	0.588	4.5	0.207	0.400	3.5	0.137	0.331	3
	S_2	0.354	0.239	1	0.074	0.114	1	0.080	0.115	1	0.110	0.172	1
	S_4	0.369	0.264	1	0.095	0.158	1	0.081	0.135	1	0.096	0.141	1
B	S_{ent}	0.732	0.480	1	0.076	0.111	1	0.087	0.130	1	0.119	0.183	2
	$S_{0.5}$	1.255	0.745	5	0.280	0.419	4	0.157	0.267	3.5	0.131	0.282	3
	S_2	1.365	0.742	5#	0.289	0.331	3#	0.118	0.156	1	0.122	0.184	2
	S_4	1.508	0.757	5#	0.466	0.509	4#	0.328	0.393	4#	0.182	0.232	3#
C	S_{ent}	0.021	0.217	2	0.059	0.163	3	0.122	0.295	3.5	0.159	0.371	4
	$S_{0.5}$	0.009	0.067	1	0.039	0.091	2	0.103	0.267	3.5	0.154	0.364	4
	S_2	0.089	0.494	3.5	0.084	0.301	3.5	0.134	0.321	4	0.163	0.377	4
	S_4	1.288	1.145	5	0.538	1.020	5	0.366	0.931	5	0.241	0.729	4.5
D	S_{ent}	0.049	0.228	1	0.044	0.125	1	0.078	0.143	3	0.117	0.236	4
	$S_{0.5}$	0.031	0.246	1	0.034	0.075	1	0.060	0.117	2	0.106	0.228	3.5
	S_2	0.060	0.238	1	0.045	0.152	2	0.075	0.153	2	0.112	0.217	3
	S_4	0.102	0.217	2	0.049	0.205	2	0.063	0.161	2	0.095	0.176	3

Table 5.2 Statistical autofocus results of estimated image after 20 iterations of cascaded optimisation for different path-cost weightings η. Im.: Image, see Fig. 5.17. Met.: Sharpness metrics used are negative entropy (S_{ent}) and power law ($S_{0.5}$, S_2 and S_4). PDE: rms error of phase difference estimate. IE: rms invariant image error (5.27). Q: Image quality rating (see Table 5.1).

containing points, extended targets and shadows (image D), all metrics focused the scene well, S_4 leaving some small residual blurring.

Regularisation had the most dramatic effect if the image did not focus well to begin with. Regularisation was effective in reducing oversharpening of the double points (image B). Over-regularisation ($\eta = 0.02$) generally had a detrimental effect on image quality. A small amount of regularisation ($\eta = 0.001$) reduced image error for most cases and reduced path error for almost all cases. If the image was well focused (for example image C and $S_{0.5}$), increased regularisation adds errors to the path and image estimate.

5.8 CONCLUSIONS

The technique of sharpness maximisation has been presented in a general Bayesian framework. This varies from similar autofocus schemes, as it includes the statistics of the platform motion with the addition of a regularisation term (the log-likelihood of the path estimate). The effect of this regularisation is to smooth the path estimate. This generally leads to a more accurate estimation of the path. It also reduces the likelihood of oversharpening the image, or estimating a path giving higher sharpness than the original by combining targets or creating spurious ones.

Different sharpness measures have been discussed and their functionality compared using the second derivative of the sharpness function. Some measures ($S_{0.5}$ and S_{ent}) are expected to perform better on images with the contrasting areas darker than the background (shadows) and some measures expected to perform better on sharpening bright points (highlights). This was confirmed using some test images. However, the negative entropy measure S_{ent} was found to work well on highlights as well as shadows, making it a good all-round choice for measuring image sharpness.

Performing a point by point (non-parametric) optimisation of regularised sharpness was found to be inefficient and slow to converge. Representing the path by a set of basis-functions, for example polynomials or a Fourier series, and perturbing the parameters of the basis-functions leads to a significantly more efficient optimisation. A parametric optimisation is faster than a non-parametric optimisation and is significantly faster for regularised sharpness. The path can be optimally represented by the fewest parameters using the platform motion statistics and a Karhunen-Loeve decomposition.

Performing an iterated optimisation, i.e. optimising a single parameter at a time, was shown to be significantly faster than a full multi-dimensional optimisation of all the parameters. This is due to a significantly reduced search space. This is more effective for a parametric optimisation since the parameters are less dependent. A cascaded approach, where the image was updated after the optimisation of each parameter was shown to converge faster than a concurrent approach, where the image was updated

after optimising all parameters individually. A conjugate gradient optimisation would speed convergence more, but cannot be performed for a stripmap image.

The sharpness surface for a spotlight image is generally well behaved, so local optimisation techniques are generally successful. A spotlight surface generally has a sharp peak, but far from the peak the surface has a small gradient and is noisy. A local optimisation technique is more likely to converge to the real minimum if the initial estimate of the motion is close to the true value. Sharpness maximisation is not good at estimating large-scale errors in stripmap images. It is preferable if a prior motion estimation step is performed and the data corrected before performing sharpness maximisation.

Chapter 6

PHASE OF MAXIMUM SHARPNESS

It has been shown that a particular sharpness metric ($S_{\beta=2}$ hereafter notated S_2) of noncoherent optical images is maximum when the image has zero (or linear) Fourier phase error [Muller and Buffington, 1974; Hamaker *et al.*, 1977]. This proof is important as it shows the validity of sharpness maximisation as a method to estimate the Fourier phase error in an image. The proof is repeated, then extended to coherent imagery. The phase that maximises image sharpness is investigated, its variance determined, and a improved method of calculating it developed.

The effect of aberrations on both coherent and noncoherent imaging systems, and on the sharpness of their images, is shown. The phase estimated by maximising the sharpness is determined. It is shown that the phase that maximises the sharpness of a coherent image is the Fourier phase of the aberrated image. The Fourier phase of a coherent image is a random variable (due to speckle), thus the phase giving maximum sharpness at a single point is not a good estimate of the Fourier phase error. The Fourier phase error may be estimated over an ensemble of phase estimates.

A timing error in spotlight imagery results in a one-dimensional phase error. The sharpness of all range-bins may be summed together with an arbitrary weighting. This chapter discusses the optimal weighting and methods to estimate this weighting. As in echo-correlation, the signal coherence can be used to weight the sharpness measure by an estimate of the phase inverse variance.

Non-parametric sharpness maximisation is shown to be able to be performed directly in the signal domain. A single calculation can replace a separate optimisation at each azimuth position. This leads to a significantly more efficient calculation of the phase estimate.

Non parametric sharpness maximisation is developed in the signal domain and shown to be a method of high-order echo-correlation. The limits on the variance of the two methods are the same. The performance of the two methods are compared against each other and the limit for a number of image types.

6.1 OPTICAL IMAGING SYSTEM

Sharpness maximisation was first developed for a non-coherent optical imaging system [Muller and Buffington, 1974]. This section develops a model of imaging for both coherent and noncoherent illumination. The effect of phase aberrations are modelled and the effect on image sharpness is described. The analysis of the image in the Fourier domain follows an approach by Hamaker et al. [1977].

6.1.1 Coherent Imaging System

A coherent imaging system is linear in complex field *amplitude*. The amplitude mapping is given by a shift-invariant convolution equation [Goodman, 1968];

$$U_i(x_i, y_i) = \int\!\!\!\int\limits_{-\infty}^{\infty} h(x_i - x_o, y_i - y_o)U_g(x_o, y_o) \, \mathrm{d}x_o \, \mathrm{d}y_o, \qquad (6.1)$$

where U_i is the complex field in the image plane (x_i, y_i) and U_g is the complex field of the image predicted by geometric optics in the object plane (x_o, y_o), a scaled version of the actual object being imaged. The system impulse response $h(x_i, y_i)$ is the Fourier transform of the pupil function $P(x, y)$ given by

$$h(x_i, y_i) = \mathcal{F}_{\substack{x \to x_i \\ y \to y_i}} \left\{ P(\lambda d_i x, \lambda d_i y) \right\}, \qquad (6.2)$$

where λ is the wavelength and d_i is the separation of the imaging plane and the lens. The image intensity $I_i(x_i, y_i)$ is normally of concern. This is given by

$$I_i(x_i, y_i) = |U_i(x_i, y_i)|^2 = \left| \int\!\!\!\int\limits_{-\infty}^{\infty} h(x_i - x_o, y_i - y_o)U_g(x_o, y_o) \, \mathrm{d}x_o \, \mathrm{d}y_o \right|^2. \qquad (6.3)$$

Frequency response of coherent imaging system

Taking a 2D Fourier transform of (6.1), it follows that

$$G_i(f_x, f_y) = H(f_x, f_y)G_g(f_x, f_y), \qquad (6.4)$$

where the following Fourier pairs are defined:

$$G_i(f_x, f_y) \equiv \mathcal{F}_{\substack{x_i \to f_x \\ y_i \to f_y}} \left\{ U_i(x_i, y_i) \right\}, \qquad (6.5)$$

$$H(f_x, f_y) \equiv \mathcal{F}_{\substack{x_i \to f_x \\ y_i \to f_y}} \left\{ h(x_i, y_i) \right\}, \qquad (6.6)$$

$$G_g(f_x, f_y) \equiv \mathcal{F}_{\substack{x_o \to f_x \\ y_o \to f_y}} \left\{ U_g(x_o, y_o) \right\}. \qquad (6.7)$$

$H(f_x, f_y)$ is the *coherent transfer function* (CTF). It follows from (6.2) that

$$H(f_x, f_y) = P(-\lambda d_i f_x, -\lambda d_i f_y). \tag{6.8}$$

If $P(x, y)$ is symmetric, then $H(f_x, f_y) = P(\lambda d_i f_x, \lambda d_i f_y)$. The system transfer function is set by the pupil. The pupil band-limits the system, resulting in a diffraction limited image. The Fourier transform of the intensity $\mathcal{G}_i(f_x, f_y) \longleftrightarrow I_i(x, y)$, follows from (6.3) and (6.4) as

$$\mathcal{G}_i(f_x, f_y) = H(f_x, f_y)G_g(f_x, f_y) \star H(f_x, f_y)G_g(f_x, f_y). \tag{6.9}$$

6.1.2 Noncoherent imaging system

When the object illumination is perfectly noncoherent, the field amplitudes across the object vary in a statistically independent fashion. An idealised representation of an noncoherently illuminated object is given by [Goodman, 1968]

$$\Gamma_{U_g}(x_0, y_0, x_0', y_0', t) = \left\langle U_g(x_o, y_o; t)U_g^*(x_o', y_o'; t)\right\rangle = \kappa I_g(x_o, y_o)\delta(x_o - x_o', y_o - y_o'), \tag{6.10}$$

where κ is a real constant and $\langle\rangle$ represents an infinite time average.[1] With this assumption, the result

$$I_i(x_i, y_i) = \kappa \int\int\limits_{-\infty}^{\infty} |h(x_i - x_o, y_i - y_o)|^2 I_g(x_o, y_o)\, dx_o\, dy_o \tag{6.11}$$

is obtained [Goodman, 1968]. Thus the image intensity is the convolution of the ideal image intensity I_g with the impulse response $|h|^2$. For noncoherent illumination, the imaging system is linear with *intensity*.

Frequency response of noncoherent imaging system

Taking a Fourier transform of (6.11), it follows that within scaling,

$$\mathcal{G}_i(f_x, f_y) = T(f_x, f_y)\mathcal{G}_g(f_x, f_y), \tag{6.12}$$

---

[1] For limitations of this representation, see [Goodman, 1968].

where the following Fourier pairs are defined

$$\mathcal{G}_i(f_x, f_y) = \mathcal{F}_{\substack{x_i \to f_x \\ y_i \to f_y}} \{I_i(x_i, y_i)\} \tag{6.13}$$

$$T(f_x, f_y) = \mathcal{F}_{\substack{x_i \to f_x \\ y_i \to f_y}} \{|h(x_i, y_i)|^2\} \tag{6.14}$$

$$\mathcal{G}_g(f_x, f_y) = \mathcal{F}_{\substack{x_o \to f_x \\ y_o \to f_y}} \{I_g(x_o, y_o)\} . \tag{6.15}$$

$T(f_x, f_y)$ is the *optical transfer function* (OTF). Its magnitude $|T|$ is known as the *modulation transfer function* (MTF). It follows that

$$T(f_x, f_y) = H(f_x, f_y) \star H(f_x, f_y) \tag{6.16}$$

and substituting (6.8) gives;

$$T(f_x, f_y) = P(\lambda d_i f_x, \lambda d_i f_y) \star P(\lambda d_i f_x, \lambda d_i f_y). \tag{6.17}$$

Thus the transfer function of an noncoherent optical system is the scaled autocorrelation of the pupil function.

6.1.3 Effect of aberrations on the image

Consider an aberration, or departure of the wavefront from the ideal spherical wave. It can be considered as a phase only error at the aperture [Goodman, 1968]. Consider a path length error of $W(x, y)$, giving a phase error of $\epsilon(x, y) = \frac{2\pi}{\lambda} W(x, y) = kW(x, y)$. The complex transmittance of the pupil, known as the *generalised pupil function*, is given by

$$\widetilde{P}(x, y) = P(x, y) \exp\left[j\epsilon(x, y)\right]. \tag{6.18}$$

Coherent system

The effect of an aberration on the coherent transfer function is thus

$$\widetilde{H}(f_x, f_y) = P(\lambda d_i f_x, \lambda d_i f_y) \exp\left[j\epsilon(\lambda d_i f_x, \lambda d_i f_y)\right]. \tag{6.19}$$

This shows that aberrations cause no band limitation on the CTF but introduce Fourier phase distortions in the passband.

Noncoherent system

For noncoherent illumination, the aberrated OTF becomes

$$\widetilde{T}(f_x, f_y) = P(\lambda d_i f_x, \lambda d_i f_y) \exp\left[j\epsilon(\lambda d_i f_x, \lambda d_i f_y)\right] \star P(\lambda d_i f_x, \lambda d_i f_y) \exp\left[j\epsilon(\lambda d_i f_x, \lambda d_i f_y)\right]$$
(6.20)

An important property of the aberrated OTF, is that aberrations will never increase the magnitude of the OTF (the MTF), i.e.;

$$|T(f_x, f_y)|^2 \geq |\widetilde{T}(f_x, f_y)|^2 \quad \text{for all } (f_x, f_y).$$
(6.21)

This can be proved using Schwarz' inequality on (6.20). Thus aberrations filter the higher spatial-frequency components of an noncoherent image, further reducing resolution from the diffraction-limit. The expression (6.21) is equality if and only if $\exp\left[-j\epsilon(\zeta, \eta)\right] = \kappa_1(f_x, f_y) \exp\left[-j\epsilon(\zeta - \lambda d_i f_x, \eta - \lambda d_i f_y)\right]$ where κ_1 is a complex constant or $\epsilon(\zeta, \eta) - \epsilon(\zeta - \lambda d_i f_x, \eta - \lambda d_i f_y) = \kappa_2(f_x, f_y)$ where κ_2 is a real constant, i.e., the phase error is linear.

6.1.4   Effect of aberrations on sharpness

Consider the sharpness of the image given by

$$S_2 = \int\!\!\!\int\limits_{-\infty}^{\infty} I^2(x, y) \, dx \, dy,$$
(6.22)

then by Rayleigh's theorem [Bracewell, 1986],

$$S_2 = \int\!\!\!\int\limits_{-\infty}^{\infty} |\mathcal{G}_i(f_x, f_y)|^2 \, df_x \, df_y.$$
(6.23)

Noncoherent image

Substituting the Fourier transform of the image intensity (6.12) into the sharpness expression (6.23), gives

$$S_2 = \int\!\!\!\int\limits_{-\infty}^{\infty} |\widetilde{T}(f_x, f_y)|^2 |\mathcal{G}_g(f_x, f_y)|^2 \, df_x \, df_y.$$
(6.24)

The ideal image $\mathcal{G}_g(f_x, f_y)$ is not affected by a phase aberration, so maximising S_2 is the same as maximising $\int\!\!\int_{-\infty}^{\infty} |\widetilde{T}(f_x, f_y)|^2 \, df_x \, df_y$ for all (f_x, f_y). From (6.21), $|T(f_x, f_y)|^2 \geq |\widetilde{T}(f_x, f_y)|^2$. Thus S_2 is maximum when there is no aberration and the image is

diffraction-limited. The equality will exist only if $\epsilon(x, y)$ is linear over all (x, y). A linear phase error causes a translation of the image and no degradation.

Alternately, following the approach of Hamaker et al. [1977], the aberrated OTF (6.20) can be expanded into

$$\widetilde{T}(f_x, f_y) = \iint\limits_{-\infty}^{\infty} P(\zeta, \eta) P(\zeta + \lambda d_i f_x, \eta + \lambda d_i f_y)$$

$$\exp\left[j(\epsilon(\zeta, \eta) - \epsilon(\zeta + \lambda d_i f_x, \eta + \lambda d_i f_y))\right] \, d\zeta \, d\eta \quad (6.25)$$

The quantity $\iint\limits_{-\infty}^{\infty} |\widetilde{T}(f_x, f_y)|^2 \, df_x \, df_y$ will be maximum when $|\widetilde{T}|$ is maximum over all spatial frequencies. From (6.25), this will be maximum when all phasors in the integral are aligned, or of equal phase, i.e.,

$$\epsilon(\zeta, \eta) - \epsilon(\zeta + \lambda d_i f_x, \eta + \lambda d_i f_y) = \kappa_0(f_x, f_y). \quad (6.26)$$

Integrating over (ζ, η), (6.26) becomes

$$\epsilon(f_x, f_y) = \kappa_1 f_x + \kappa_2 f_y + \kappa_3. \quad (6.27)$$

This shows, again, that to maximise S_2, $\epsilon(f_x, f_y)$ must be linear.

Coherent image

As for the noncoherent case, the sharpness is given by

$$S_2 = \iint\limits_{-\infty}^{\infty} |\mathcal{G}_i(f_x, f_y)|^2 \, df_x \, df_y. \quad (6.28)$$

A coherent image differs from the noncoherent case in (6.12), since the Fourier transform of the image intensity is given by

$$\mathcal{G}_i(f_x, f_y) = H(f_x, f_y) G_g(f_x, f_y) \star H(f_x, f_y) G_g(f_x, f_y). \quad (6.29)$$

Thus the transfer function and object cannot be separated as in (6.24). Substituting the generalised pupil function and defining $\xi(f_x, f_y) = \angle\{G_g(f_x, f_y)\}$ gives

$$\mathcal{G}_i(f_x, f_y) = \iint\limits_{-\infty}^{\infty} P(\zeta, \eta) P(\zeta + \lambda d_i f_x, \eta + \lambda d_i f_y) |G_g(\zeta, \eta)| |G_g(\zeta + \lambda d_i f_x, \eta + \lambda d_i f_y)|$$

$$\exp\left[j(\epsilon(\zeta, \eta) - \epsilon(\zeta + \lambda d_i f_x, \eta + \lambda d_i f_y) + \xi(\zeta, \eta) - \xi(\zeta + \lambda d_i f_x, \eta + \lambda d_i f_y))\right] \, d\zeta \, d\eta.$$

$$(6.30)$$

The summation of (6.28) will be maximum when $|\mathcal{G}_i|$ is maximum over all frequencies. This will occur when all the phasors in the summation in (6.30) are aligned, or

$$\epsilon(\zeta, \eta) - \epsilon(\zeta + \lambda d_i f_x, \eta + \lambda d_i f_y) + \xi(\zeta, \eta) - \xi(\zeta + \lambda d_i f_x, \eta + \lambda d_i f_y) = \kappa_0(f_x, f_y) \quad (6.31)$$

Integrating over (ζ, η), (6.31) becomes

$$\epsilon(f_x, f_y) + \xi(f_x, f_y) = \kappa_1 f_x + \kappa_2 f_y + \kappa_3 \quad (6.32)$$

Thus maximising the sharpness of a coherent image does not directly estimate the phase error of the image $\epsilon(f_x, f_y)$, but the Fourier phase of the aberrated image $\epsilon(f_x, f_y) + \xi(f_x, f_y)$. The variance of the phase of maximum sharpness will depend on the variance of the image Fourier phase, $\xi(f_x, f_y)$.

Estimating the phase error of a coherent image

The phase of maximum sharpness estimates the Fourier phase of the image. Thus an estimate of the 2D phase error cannot be made separately at each pixel. An ensemble average over a patch is required. If $\mathrm{E}\left[\xi(f_x, f_y)\right] = 0$, the image Fourier phase component $\xi(f_x, f_y)$ is removed. For 2D phase errors, such as in astronomical images, the phase error of a segment of the image can be estimated [Muller and Buffington, 1974]. The larger the patch, the more accurate the phase estimate, but with a trade-off of lower estimator resolution. Another possibility is to form an ensemble average over many short-exposure images or 'looks' of the same scene, as is performed in speckle interferometry [Bates and McDonnell, 1986]. In a SAR spotlight image, it is common to deal with a 1D phase error. An ensemble average can then be made over the other dimension, removing the image phase component from the phase estimate. Estimating the phase error of a spotlight image using sharpness maximisation will be developed further in the following section.

6.2   SPOTLIGHT SYNTHETIC APERTURE SYSTEM

Consider a spotlight SAR imaging system. Motion of the platform, or medium, will corrupt the signal history with a phase error. It is usual to assume a 1D phase error which is a function of the azimuth position v, so the model of the measured signal history is

$$G[x, v] = \overline{G}[x, v] \exp[j \phi_e[v]], \quad (6.33)$$

where $\overline{G}[x, v]$ is the ideal range-compressed signal history and $\phi_e(v)$ is the phase error degrading the image. Once a Fourier phase estimate $\widehat{\phi}$ is made, it can be corrected for with

$$\widehat{G}[x, v] = G[x, v] \exp[-j \widehat{\phi}[v]]. \quad (6.34)$$

The corresponding complex image estimate is

$$\widehat{g}[x,y] = \mathcal{F}_{v \to y}^{-1}\left\{\widehat{G}[x,v]\right\} = \sum_{v} \widehat{G}[x,v]\exp\left(j2\pi yv/N\right). \tag{6.35}$$

Thus the image is corrupted with a Fourier phase error. This is similar to the optical imaging model but the Fourier phase error varies only in the along-track direction. Comparing the system model for a spotlight system (6.33) with a coherent optical system (6.4)(6.5)(6.19), they are identical, except the Fourier phase error occurs in two dimensions in the optical image. The pupil function, $P(x,y)$ is incorporated in the object signal $G[x,v]$. Since the Fourier transform is in the azimuth direction, each range bin can be considered separately as a 1D estimation problem and the estimates for each range can be combined later.

6.2.1 Single range bin

Consider a single range bin $g[x_0,y]$. One estimate of the Fourier phase error $\widehat{\phi}(x_0,v)$ is the phase error that maximises the sharpness measure for this range-bin;

$$S_2(x_0) = \sum_{y} |\widehat{g}[x_0,y]|^4. \tag{6.36}$$

From (6.32) applied in one dimension;

$$\widehat{\phi}(x_0,v) = \phi_e(v) + \xi(x_0,v) + \kappa_1 v + \kappa_2, \tag{6.37}$$

where $\xi(x_0,v)$ is the phase of $\overline{G}[x_0,v]$ and κ_1,κ_2 are arbitrary constants. Ignoring the linear offset, it is apparent that

$$\mathrm{E}\left[\widehat{\phi}(x_0,v)\right] = \phi_e(v) + \mathrm{E}\left[\xi(x_0,v)\right], \tag{6.38}$$

$$\mathrm{Var}\left[\widehat{\phi}(x_0,v)\right] \equiv \sigma_\phi^2(x_0) = \mathrm{Var}\left[\xi(x_0,v)\right]. \tag{6.39}$$

Thus the phase estimate from a single range bin depends on the statistics of the Fourier phase of the object ξ. The statistics of the image phase are developed in Section 4.3.1. The phase of maximum sharpness is a direct way of estimating image phase $\xi(x_0,v)$, as opposed to echo-correlation which estimates phase difference $\theta(x_0,v)$. From (4.72), the variance of the phase difference is

$$\mathrm{Var}\left[\theta_x\right] = 2\,\mathrm{Var}\left[\xi_x\right]. \tag{6.40}$$

The phase variance of a constant point target can be estimated from the signal coherence γ_x using (4.71);

$$\text{Var}\,[\xi_x] \approx \frac{1 - |\widehat{\gamma_x}|^2}{4\,|\widehat{\gamma_x}|^2}.$$ (6.41)

The phase variance of a rough extended target can be calculated from the pdf using (4.79);

$$\text{Var}\,[\xi] = \int_{-\pi}^{\pi} \xi^2\, f_\xi(\xi + \xi_0)\, d\xi,$$ (6.42)

where the probability density function $f_\xi(\xi)$ of the phase is given by (4.78) and $\xi_0 = \text{E}\,[\xi]$. Fig. 6.1 shows the variance of the phase of maximum sharpness of a single range-bin ϕ for a variety of signal to clutter ratios. For a constant point target (Fig. 6.1(a)), there is a close match with the image phase ξ and the model (6.41). For a rough extended target (Fig. 6.1(b)), the variance is slightly above that of the image phase ξ and the model (6.42).
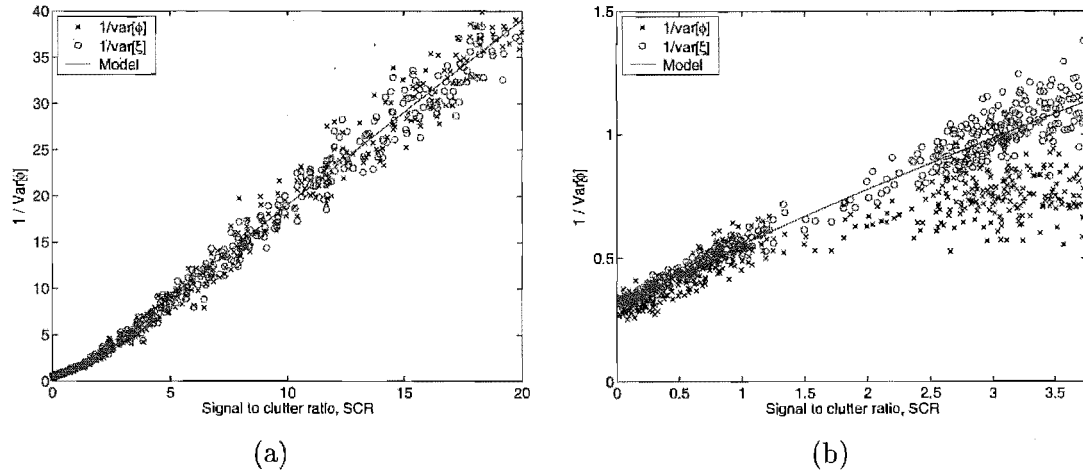
Figure 6.1 The variance of the phase of a single range-bin of the image ξ, compared to model and the variance of the phase giving maximum sharpness ϕ. (a) Constant point target. Model given by (6.41). (b) Rough block target. Model given by (6.42).

6.2.2 Combining range bins

For a spotlight imaging system, an estimate of the image Fourier phase can be obtained for each range bin using

$$\widehat{\phi}(x, v) = \max_{\phi(v)} [S(x, \phi(v))].$$ (6.43)

An estimate of the phase error ϕ_e can then be made by combining all range estimates. Each estimate will have a different variance, depending on the signal-clutter ratio of the range-bin. The information from the separate range bins can be combined in a number of ways. The phase estimates could be combined, or the sharpness values combined.

Sum of phase estimates

A simple approach is to average the phase of all range-bins using

$$\widehat{\phi}_{\text{OLS}}(v) = \frac{1}{N_x} \sum_x \widehat{\phi}(x, v). \tag{6.44}$$

This is the ordinary least squares (OLS) estimator. It minimises $\sum_x \left[\widehat{\phi}_{\text{OLS}}(v) - \widehat{\phi}(x,v) \right]^2$ assuming the errors are independent, zero mean, and constant variance (homoscedasticity) [Beck and Arnold, 1977]. The variance of the estimated value is reduced by

$$\text{Var}\left[\widehat{\phi}_{\text{OLS}}\right] = \frac{\sigma_\phi^2}{N_x}. \tag{6.45}$$

A problem with this method is that the phase estimate at each range has a different variance, but each is weighted the same. If a range-bin has a coherent target, then $\text{E}\left[\widehat{\phi}(x,v)\right] = \phi_e(v)$. If a range-bin has speckle only, the phase estimate from that range-bin is a uniformly distributed random variable in the interval $[-\pi, \pi]$. Due to phase wrapping, $\text{E}\left[\widehat{\phi}(x,v)\right] = 0$. Thus any range-bins without coherent targets will bias the estimator.

Assuming the errors are independent with varying variance (heteroscedasticity), then a maximum likelihood (ML) approach is better [Beck and Arnold, 1977]. Each phase estimate is scaled by the inverse variance of the estimate $\sigma_\phi^{-2}(x)$, using

$$\widehat{\phi}_{\text{ML}}(v) = \frac{\sum_x \widehat{\phi}(x, v) \sigma_\phi^{-2}(x)}{\sum_x \sigma_\phi^{-2}(x)}. \tag{6.46}$$

The maximum likelihood estimator minimises $\sum_x \left[\widehat{\phi}_{\text{ML}}(v) - \widehat{\phi}(x,v)\right]^2 / \sigma_\phi^2(x)$. Thus measurements with a high variance are weighted less, resulting in a more accurate estimation overall. The variance of the estimated phase is then

$$\text{Var}\left[\widehat{\phi}_{\text{ML}}\right] = \left[\sum_x \sigma_\phi^{-2}(x)\right]^{-1}. \tag{6.47}$$

This maximum likelihood approach minimises the variance of the estimated phase $\text{Var}\left[\widehat{\phi}_{\text{ML}}\right]$.

This method is inefficient to calculate, as it requires a separate optimisation at each range-bin. Another major problem is caused by the phase wrapping in the interval $[-\pi, \pi]$. Directly averaging phase values is generally not practical for these reasons.

One possibility is to approximate the phase averaging procedure using

$$\tilde{\phi}_{\mathrm{ML}}(v) = \angle\left\{\sum_x \sigma_\phi^{-2}(x)\exp\left(j\widehat{\phi}(x,v)\right)\right\}.$$ (6.48)

which does not suffer from wrapping problems.

Sum of sharpness

A more efficient approach is to sum sharpness over all range-bins, then maximise the sum, i.e.,

$$\widehat{\phi_s}(v) = \max_{\phi(v)}\left[\frac{1}{N_x}\sum_x S(x,\phi(v))\right].$$ (6.49)

This is the standard sharpness measure. It is simpler to compute than $\widehat{\phi}_{\mathrm{OLS}}$ since it only requires one maximisation. However, it will give a different result to $\widehat{\phi}_{\mathrm{OLS}}$ as shown in Fig. 6.2. The estimator $\widehat{\phi}_{\mathrm{OLS}}$ weights each range bin equally, while $\widehat{\phi}_s$ sums together the sharpness measures from each range-bin and maximises the result. Thus the result will depend on the relative size of the peak sharpness of each range-bin. Those ranges with larger sharpness will dominate. This may lead to over-sharpening as discussed in Section 5.3.5.
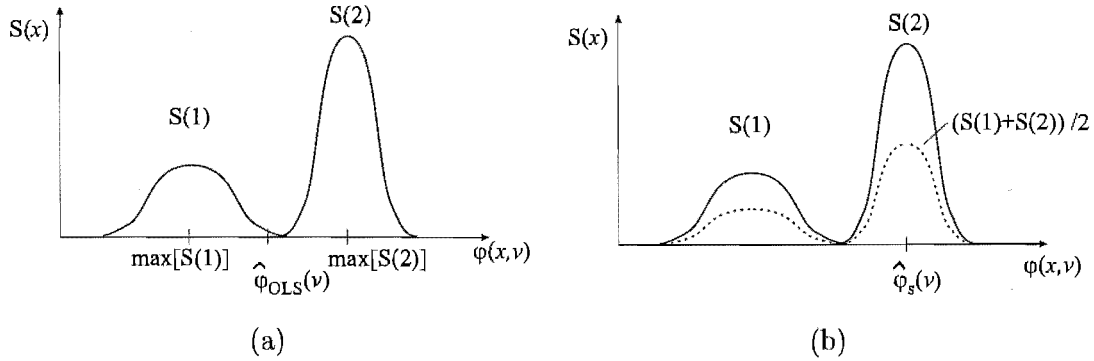
Figure 6.2   Estimating Fourier phase error from sharpness measure for two separate range bins. (a) $\widehat{\phi}_{\mathrm{OLS}}(v) = (\max[S(1)] + \max[S(2)])/2$. (b) $\widehat{\phi}_s(v) = \max[(S(1) + S(2))/2]$.

Range normalised sharpness

If the sharpness of each range-bin is summed, the contribution of each range bin is weighted by the magnitude of the sharpness of that bin. However, the sharpness of a range-bin is weighted by the average return. For a radar or sonar, the average return will vary with range. This can be compensated but this compensation is not always performed. Also, the topography can vary the average backscatter with range. An area of strong shadow has lower overall return, so will have lower sharpness. However, if it has strong coherence, it will have a lower variance of the phase estimate. It is desirable

to make the contribution from each range independent of a multiplicative scale factor. Normalising each range bin with average intensity gives a *range-normalised sharpness measure*, given by

$$\overline{S_\Omega} = \frac{1}{N_x} \sum_x \frac{\frac{1}{N_y} \sum_y \Omega[I(x,y)]}{\Omega\left[\frac{1}{N_y} \sum_y I(x,y)\right]}, \tag{6.50}$$

where the generalised sharpness function $\Omega[I(x,y)]$ is a nonlinear function of the image intensity (see Section 5.3.3). Each range-bin is weighted independent of any multiplicative scale factor, resulting in a fairer estimator. An example is shown in Fig. 6.3, where the range-normalised sharpness measure has a more equitable weighting between range-bins and a lower variance of resulting phase estimate than the unweighted sharpness measure. Note that $\sum_y I(x,y)$ is independent of Fourier phase errors, so can be calculated once prior to the optimisation routine.

The range-normalised measure is an example of a weighted sharpness measure, given by

$$S_W = \frac{1}{N_x} \sum_x W(x) S(x, \phi(v)). \tag{6.51}$$

This weighting should be independent of the estimated phase. Those range-bins judged to have higher SCR should be weighted more heavily. A weighting suggested by Fienup [2000] is

$$W(x) = \frac{1}{\left[\sum_y I(x,y)\right]^2}, \tag{6.52}$$

which for S_2 gives the range-normalised sharpness measure $\overline{S_2}$. However, the range-normalised range-weighting (6.50) is a more general expression.

Fig. 6.3(b) shows the sharpness weighting for various $W(x)$ for a simulated image with mean intensity $\mu_I(x)$. For $W(x) = 1$, the area of shadow is weighted lower than the speckle only. For $W(x) = \mu_I^{-1}(x)$, the shadow is weighted the same as the speckle. For $W(x) = \mu_I^{-2}(x)$, each range-bin is weighted proportional to the inverse variance σ_ϕ^{-2} as desired.

Maximum-likelihood weighting

The best range weighting would approximate the phase weighting of (6.46) by weighting each range-bin by the inverse variance of the phase estimate. One method would be to normalise the sharpness by the sharpness of the original image, then weight by an estimate of the phase inverse variance at each range-bin. The weighted sharpness (6.51) could be maximised with $W(x) = \frac{1}{S(x,0)\widehat{\sigma}_\phi^2(x)}$ giving

$$\widehat{\phi_w}(v) = \max_{\phi(v)} \left[\sum_x \frac{1}{S(x,0)\widehat{\sigma}_\phi^2(x)} S(x, \phi(v)) \right], \tag{6.53}$$
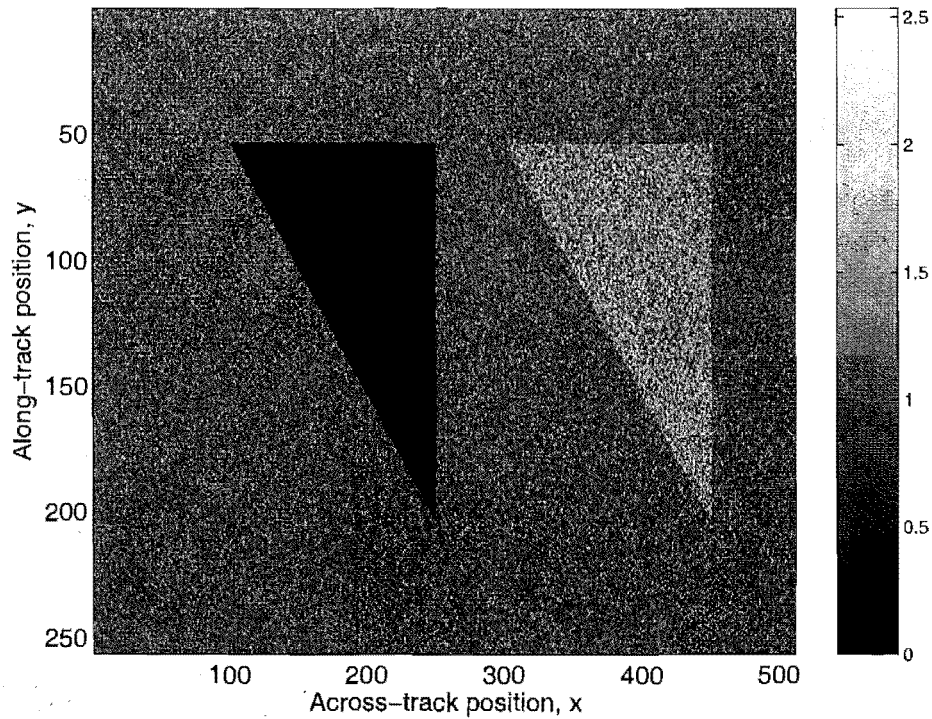
(a)
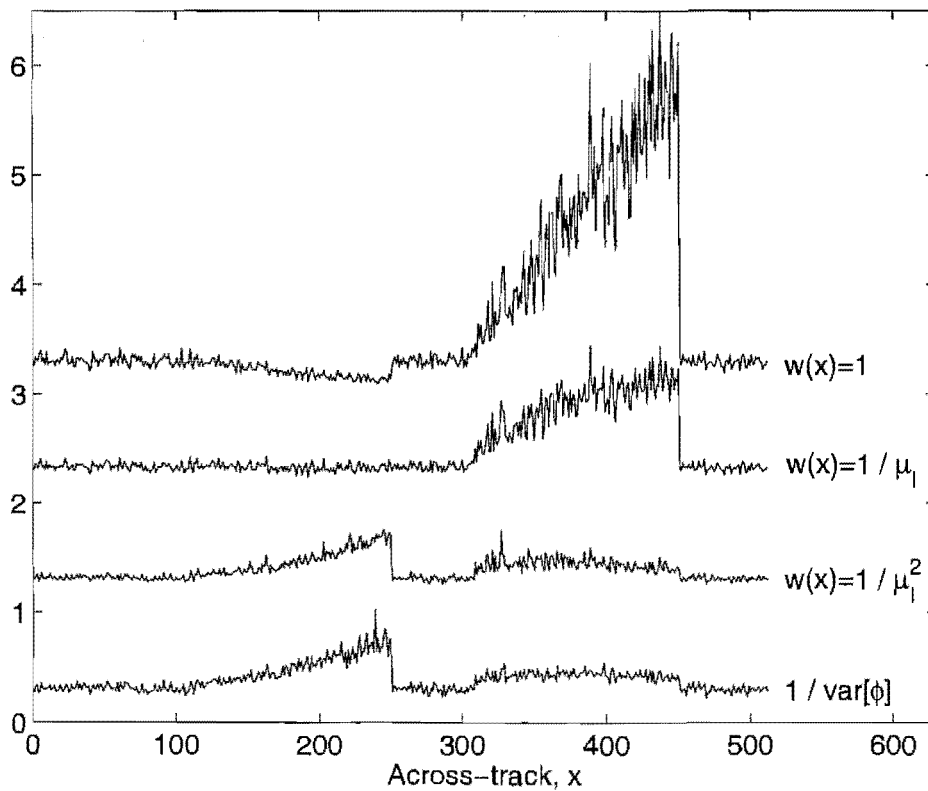
(b)

Figure 6.3 Example of the effect of range weighting on sharpness. (a) Root intensity of simulated image. Contains increasing shadow $100 < x < 250$ and increasing target $300 < x < 450$. (b) Range bin weighting $W(x).\max\{S_2\}$ for various $W(x)$ compared to phase estimate inverse variance σ_ϕ^{-2}. Plots offset for clarity. $W(x) = 1$ corresponds to weighting of S_2 and resulted in Var $[\phi] = 0.0322$. $W(x) = \mu_I^{-2}$ corresponds to $\overline{S_2}$ and resulted in Var $[\phi] = 0.0161$.

where the phase variance estimate $\widehat{\sigma}_\phi^2(x)$ is made using the measured coherence (6.41).

Modified weighting

For echo-correlation, better results are obtained if a modified inverse variance weighting was used, as shown in Section 4.3.2. For clutter only, the phase is uniformly distributed with a finite phase variance of $\pi^2/3$. The modified weighting reduces the estimated variance caused by phase wrapping by using a weighting of

$$W(x) = \frac{\sigma_\phi^{-2}(x) - 0.3}{S(x,0)},\qquad(6.54)$$

Results have shown this also improves the results for sharpness maximisation.

6.2.3   Effect of different weightings

A number of images of different types were tested. A list of the images used is shown in Table 6.1 and the images are shown in Fig. 6.4 and Fig. 6.5. The phase at each azimuth position v was varied and the sharpness of the image measured at each range bin. The sharpness was then combined over all range-bins by different methods and the phase giving maximum sharpness $\widehat{\phi}$ found for each azimuth position v. The mean squared error of this phase estimate is shown in Table 6.2.

Image	Description
A	Rectangular target and shadow
B	Clutter with increasing gradient
C	Combination of block targets and shadows
D	Triangular shadow and target
E	Point target of increasing strength with range
F	Several point targets (band-limited scene).
G	Real image of shipwreck with seafloor ripples and large shadow
H	Simulated band-limited image of mine-like objects with shadows

Table 6.1   Description of images used in spotlight trials. Images are shown in Fig. 6.4 and Fig. 6.5.

The variance of the phase estimate $\widehat{\phi}$ in each range-bin is compared to the weighting of sharpness measures S_2 and S_W in Fig. 6.4 and Fig. 6.5. The weighted sharpness S_W uses the modified weighting (6.54). S_W matches the measured phase inverse variance well for all image types. This corresponds to a lower phase estimate error in Table 6.2. The unmodified sharpness S_2 weights by signal energy, not coherence and performs poorly in some image regions. Areas of shadow (images A, C, D and H) are weighted too low. For point targets (E and F) and block targets of varying intensity (A) and size (D), the magnitude of he weighting does not match the inverse variance well over
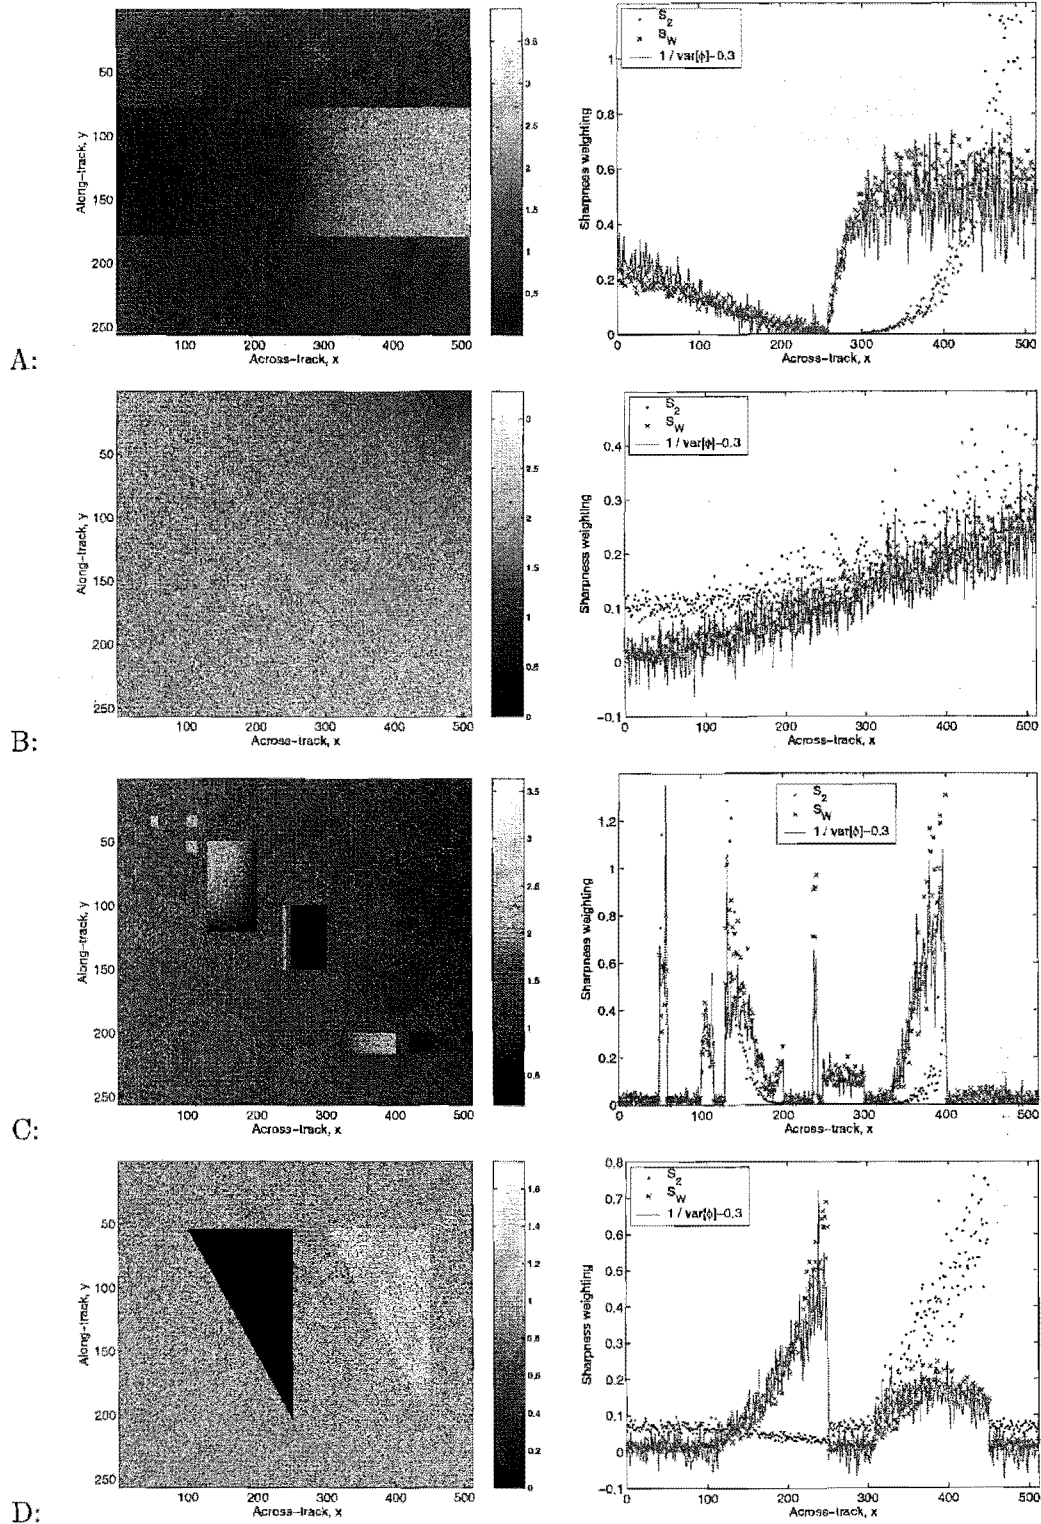
Figure 6.4 Simulated images, with the sharpness weighting for S_2 and S_W using the modified weighting (6.54) at each range-bin compared to the measured variance of the estimated phase.

Image A: a rectangular target 101 pixels wide, strengthening for $x > 256$ and rectangular shadow 101 pixels wide, strengthening for $x < 256$.

Image B: clutter with along-track intensity gradient increasing for larger x.

Image C: a combination of block targets, point targets and shadows, of varying strengths.

Image D: a triangular patch of zero return and triangular patch with return double that of background.

E:

F:

G:

H:
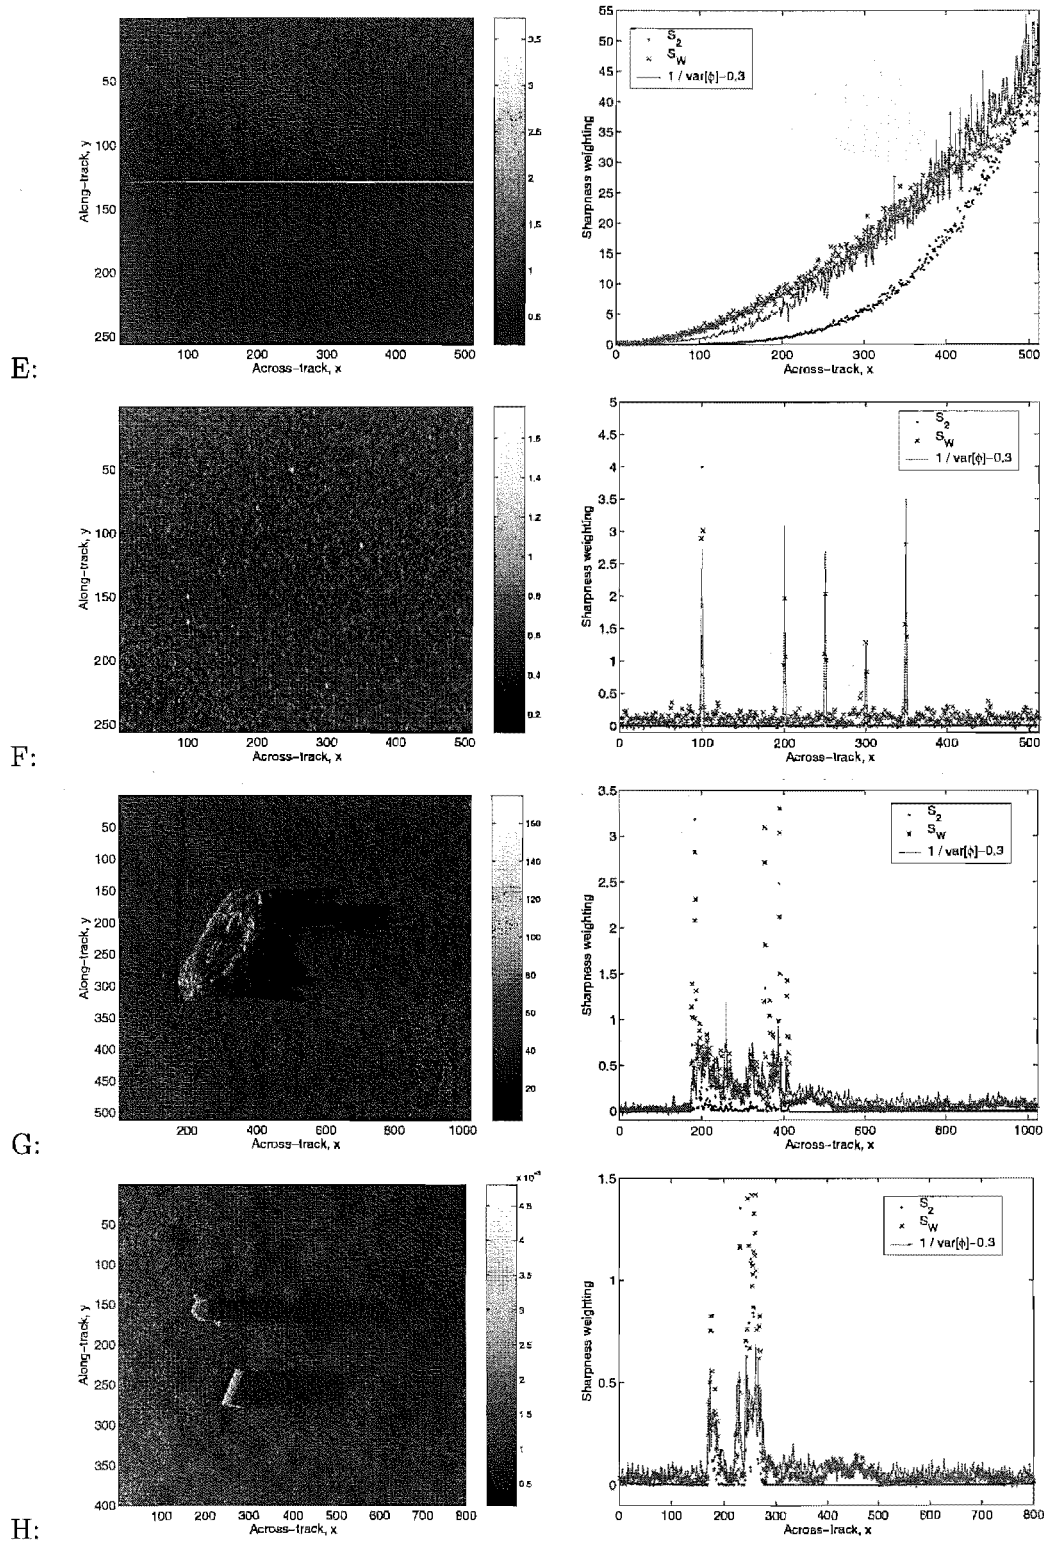
Figure 6.5 Test images, with the sharpness weighting for S_2 and S_W using the modified weighting (6.54) at each range-bin compared to the measured variance of the estimated phase.
Image E: a point target in center of each range-bin. Intensity increases for larger x.
Image F: a band-limited scene of clutter with several point targets.
Image G: real image of shipwreck.
Image H: simulated image of teapot and cylinder with shadows.

Image	Sharpness measure					CRLB
	S_2	$\overline{S_2}$	$\sqrt{\overline{S_2}}$	S_W	S_C	
A	0.00619	0.00253	0.00264	0.00241	0.00260	0.00080
B	0.01102	0.01171	0.01233	0.01004	0.01034	0.00212
C	0.01503	0.00633	0.00643	0.00575	0.00684	0.00177
D	0.03220	0.01611	0.01825	0.01177	0.01280	0.00227
E	0.000151	0.000144	0.000190	0.000115	0.000124	0.000117
F	0.06304	0.05519	0.14724	0.05369	0.05004	0.00593
G	0.03837	0.00947	0.00754	0.00915	0.01135	0.00047
H	0.01280	0.00428	0.00390	0.00622	0.00804	0.00047

Table 6.2 Mean squared error of phase estimates of various simulated images for different sharpness weightings. The images are described in Table 6.1 and shown in Fig. 6.4 and Fig. 6.5. S_W uses modified weighting (6.54) and rough surface pdf model (6.42). S_C uses point-target coherence model (6.41). The CRLB is given by (6.93).

all intensities. The weighted sharpness measure S_W resulted in the lowest phase error for all images except G and H.

A constant point target and rough surface target have a different effect on the phase variance. The same model cannot accurately predict the phase variance for both image types. This makes it difficult to use a single sharpness weighting which is best for all images. An improved model, which can account for both image types, is required. S_W works well for all images tested, with $\sqrt{\overline{S_2}}$ giving slightly better performance for images G and H. It may be posible to improve performance by characterising the image type using methods discussed in [Fienup and Miller, 2003, Sec 8].

6.3 DIRECT SHARPNESS MAXIMISATION

The previous section showed that the phase estimate that maximises the sharpness S_2 of a single range-bin is the Fourier phase of the aberrated image with a linear offset. This leads to the possibility that the phase giving maximum sharpness can be measured directly in the Fourier, or signal domain. This would be preferable, since optimisation is slow and needs to be repeated at every azimuth position. This section develops a method for measuring the sharpness in the Fourier domain. The effect of altering the phase of a single azimuth position on the intensity squared sharpness S_2 is then developed. It is shown that the phase that maximises this sharpness measure can be directly measured from the signal, removing the need to perform an optimisation. In addition, the phase estimate can be calculated at all azimuth positions simultaneously leading to a far more efficient phase estimation method. The proof of this result is developed in Appendix A, with the major results shown here.

Since direct sharpness-maximisation replaces a 1D optimisation, it requires recur-

sion when applied over many parameters. Typically, between 3 and 30 iterations are required to converge, depending on image type. Tests comparing the estimate from direct sharpness-maximisation to a full optimisation in the image domain show that both techniques do converge to the same result. Results applying different range weightings are shown.

An alternate derivation of direct sharpness maximisation shows the technique can be considered an extension of the conjugate gradient optimisation method. This derivation allows a general sharpness metric to be used.

6.3.1 Sharpness in Fourier domain

A common measure of image sharpness for a single range-bin x is calculated using

$$S_2[x] = \sum_y |g_x[y]|^4 . \tag{6.55}$$

The range-bin index x, will be dropped from the notation for clarity. The sharpness can be reformulated using the Fourier autocorrelation and energy theorems to be

$$S_2 = N \sum_d \left| \sum_v G[v] \, G^*[v - d] \right|^2 , \tag{6.56}$$

$$= N \sum_d |GG[d]|^2 , \tag{6.57}$$

where $G[v]$ is the discrete Fourier transform of $g[y]$, i.e.,

$$G[v] = \frac{1}{N} \sum_y g[y] \exp\left(-j2\pi yv/N\right) , \tag{6.58}$$

and where $GG[d]$ is the autocorrelation of $G[v]$ at lag d, i.e.,

$$GG[d] = \sum_v G[v] \, G^*[v - d] = \frac{1}{N} \sum_y |g[y]|^2 \exp\left(-j2\pi yv/N\right) . \tag{6.59}$$

6.3.2 Non-parametric sharpness maximisation

To estimate the phase at azimuth position v_0, the phase estimate at that position can be varied until the sharpness is maximised, i.e.,

$$\widehat{\phi}[v_0] = \arg\max_\phi \left\{ \tilde{S}_2[v_0, \phi] \right\} \tag{6.60}$$

where $\tilde{S}_2[v_0, \phi]$ is the sharpness of the image with a phase correction of ϕ at azimuth position v_0. It can be shown by straightforward but tedious manipulation (see

Appendix A), that the adjusted sharpness depends on ϕ by (A.42)

$$
\begin{aligned}
\tilde{S}_2[v_0, \phi] = {} & S_2 + 8E \left| G[v_0] \right|^2 - 4N\Re\left\{K[v_0]\right\} + 2N\Re\left\{L[v_0]\right\} - 6N \left| G[v_0] \right|^4 \\
& + 4N\Re\left\{K[v_0]\exp\left(-j\phi\right)\right\} \\
& - 4N\Re\left\{L[v_0]\exp\left(-j\phi\right)\right\} + 2N\Re\left\{L[v_0]\exp\left(-j2\phi\right)\right\} \\
& - 2N \left| G[v_0] \right|^4 \cos(2\phi) + 8N \left| G[v_0] \right|^4 \cos(\phi) \\
& - 8E \left| G[v_0] \right|^2 \cos(\phi).
\end{aligned}
\tag{6.61}
$$

where the following quantities are defined;

$$
E = N \sum_v |G[v]|^2 ,
\tag{6.62}
$$

$$
K[v_0] = G[v_0] \sum_d GG[d]\, G^*[v_0 + d],
\tag{6.63}
$$

$$
L[v_0] = G^2[v_0] \sum_d G^*[v_0 - d]\, G^*[v_0 + d].
\tag{6.64}
$$

E is the total energy in the signal. $K[v_0]$ can be considered a high order, weighted, phase difference estimate at v_0 while $L[v_0]$ is an high order, weighted, phase curvature estimate at v_0. To maximise the aberrated sharpness (6.61), one only need consider the terms that vary with ϕ. Neglecting terms with negligible amplitude a further simplification can be made, giving (A.44)

$$
\tilde{S}'_2[v_0, \phi] \approx 4N\Re\left\{ \left(K[v_0] - 2\frac{E}{N} |G[v_0]|^2 \right) \exp\left(-j\phi\right) \right\},
\tag{6.65}
$$

or alternatively,

$$
\tilde{S}'_2[v_0, \phi] \approx 4N\Re\left\{\chi[v_0]\exp\left(-j\phi\right)\right\}.
\tag{6.66}
$$

where

$$
\begin{aligned}
\chi[v_0] &= G[v_0] \sum_d GG[d]\, G^*[v_0 + d] - 2GG[0] \left| G[v_0] \right|^2 \\
&\approx G[v_0] \sum_d GG[d]\, G^*[v_0 + d],
\end{aligned}
$$

which may be calculated using

$$
\chi[v_0] = G[v_0] \left(\mathcal{F}_{y \to v} \left\{ |g[y]|^2\, g[y] \right\} \right)^*.
\tag{6.67}
$$

Equation (6.66) shows that the sharpness depends on the phase estimate ϕ by a sinusoidal factor. When altering the phase of a single azimuth position on any image, the sharpness is always a single sinusoid period 2π. The total sharpness can thus be considered a sum of sinusoids.

6.3.3   Direct phase estimation

The phase estimate $\hat{\phi}$ that maximises the image sharpness is equivalent to maximising the real part of χ, i.e.,

$$\hat{\phi}[v_0] = \arg\max_{\phi} \left\{ \mathrm{Re} \left\{ \chi[v_0] \exp\left(-j\phi\right) \right\} \right\}. \tag{6.68}$$

The real component of a complex number is maximum when it has zero phase. The phase estimate $\hat{\phi}$ can thus be directly calculated using

$$\hat{\phi}[v_0] = \angle \left\{ \chi[v_0] \right\}. \tag{6.69}$$

All the phase estimates can be estimated with one calculation:

$$\hat{\phi}[v_0] = \angle \left\{ G[v_0] \left(\mathcal{F}\left\{ |g[y]|^2\, g[y] \right\} \Big|_{[v_0]} \right)^* - 2\,|G[v_0]|^2 \sum_v |G[v]^2| \right\}. \tag{6.70}$$

This one calculation replaces N_y separate 1D optimisation routines, of a concurrent iterated optimisation, in the image domain.

Combining range-bins

It has been shown that a weighted combination of range-bins can improve results (see Section 6.2.2). A weighted sharpness measure can be calculated using (6.51)

$$\hat{\phi}[v] = \arg\max_{\phi} \left\{ \sum_x W(x)\tilde{S}[x, v_0, \phi] \right\}. \tag{6.71}$$

The identical operation can be performed for direct sharpness maximisation using

$$\hat{\phi}[v] = \angle \left\{ \sum_x W[x]\chi[x, v] \right\}, \tag{6.72}$$

where

$$\chi[x, v] = G_x[v] \left(\mathcal{F}_{y\to v}\left\{ |g_x[y]|^2\, g_x[y] \right\} \right)^* - 2\,|G_x[v]|^2 \sum_v |G[v]^2| \tag{6.73}$$

$$\approx G_x[v] \left(\mathcal{F}_{y\to v}\left\{ |g_x[y]|^2\, g_x[y] \right\} \right)^*. \tag{6.74}$$

Possible weightings include using no weighting (6.49),

$$W_N[x] = 1, \tag{6.75}$$

a range normalised weighting (6.52)

$$W_R[x] = \frac{1}{\left(\sum_v |G_x[v]|^2\right)^2},$$ (6.76)

and a weighting based on the measured coherence estimate of phase variance (6.53)

$$W_C[x] = \frac{1}{S(x,0)\widehat{\sigma}_\phi^2(x)} \approx \frac{4\,|\gamma_d[x]|^2}{(1 - |\gamma_x[x]|^2)\frac{1}{N_v}\sum_v |\chi[x,v]|}.$$ (6.77)

Fig. 6.6 compares the variance of the phase difference error for both a point and clutter target with varying SCR. If the SCR is constant with range, the performance of sharpness maximisation is the same, regardless of the weighting used. There is a small performance gain in using a weighted measure at high SCR for the block target. If the SCR varies with range, there is a clear performance improvement using the coherence weighted measure W_C, compared to the normalised measure W_R, which in turn performs better than the unweighted measure W_N. The normalised and coherence weighted measure have similar performance for a point target of low SCR ($SCR < 2$) and a block target.

Interpretation

Consider the direct sharpness maximisation given by (6.74) expanded to

$$\widehat{\phi}[v_0] = \angle \left\{ \sum_x \sum_d GG_x[d]G_x[v_0]G_x^*[v_0 + d] \right\}.$$ (6.78)

For a single separation d, this becomes

$$\widehat{\phi}[v_0, d] = \angle \left\{ \sum_x GG_x[d]G_x[v_0]G_x^*[v_0 + d] \right\},$$ (6.79)

which is a weighted phase difference estimation, or shear average. Thus direct maximisation of S_2 corrects the phase of a single echo by a combined weighted phase difference estimation over all separations. This could be considered a method of high-order echo-correlation. By measuring the phase difference to all other echos, then correcting by a weighted average of that difference, the mean difference between the phase of this echo and all others is minimised. The mean echo phase is *straightened*, increasing the sharpness of the image.
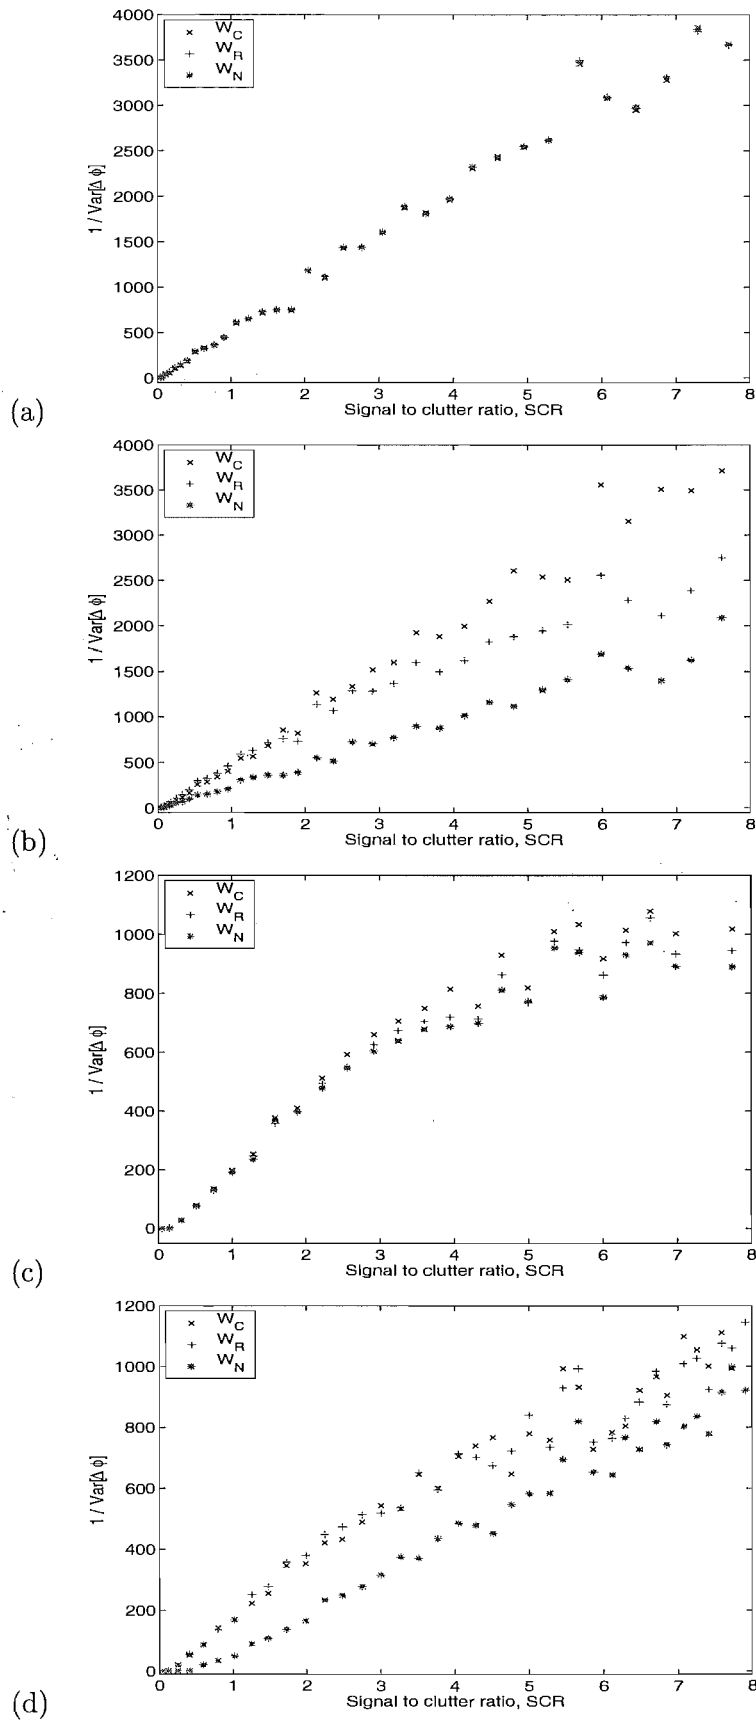
Figure 6.6 Inverse variance of error in phase difference for direct sharpness maximisation using different range-weightings. Weightings are described in text. (a) Point target, constant SCR with range. (b) Point target, background intensity varies with range. (c) Rough block target, constant SCR with range. (d) Rough block target, background intensity varies with range.

Iterating

Direct sharpness maximisation can directly calculate the solution of a single ping sharpness maximisation. To perform a maximisation over all pings, each ping can be corrected by this estimate and the process repeated. This allows for the coupling of the sharpness measure between the phase estimates at each ping. This recursive approach is the equivalent of the *concurrent iterated optimisation* (see Section 5.6). Fig. 6.7 shows the sharpness converging to the maximum in approximately 3 iterations for an image consisting of several point targets and clutter but taking approximately 30 iterations to converge for an image consisting of a rippled surface. Fig. 6.8 shows the estimate from maximising the sharpness with an optimisation algorithm closely matches the estimate using direct sharpness maximisation. Direct sharpness maximisation converges to the same phase estimate as sharpness optimisation in the image domain, but it can be performed significantly faster as the optimisation has been replaced by a quicker calculation.
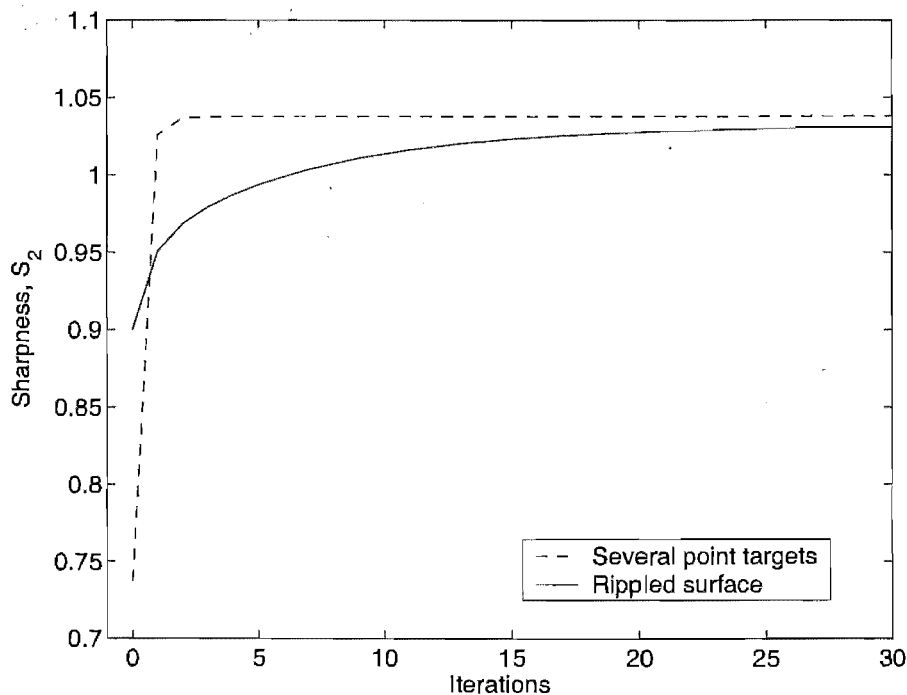
Figure 6.7 Convergence of sharpness using direct phase estimation. Uses image consisting of several point targets and image consisting of rippled surface. Sharpness is normalised to unity for aberration free image.

6.3.4 Direct sharpness maximisation from sharpness gradient

The rate of change of a sharpness measure, with respect to a change in a phase estimate $\widehat{\phi}(v)$, can be determined for a spotlight image [Gough and Lane, 1998; Fienup, 2000;
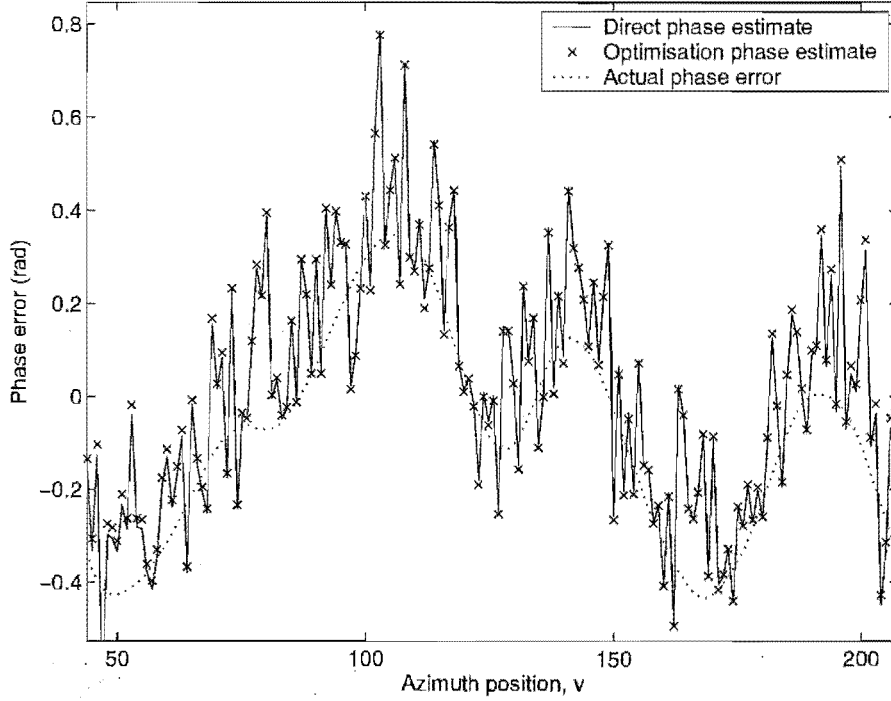
Figure 6.8 The phase estimated by 5 iterations of direct phase estimation, compared to a full cascade sharpness maximisation of the scene. (50 iterations of 256 separate optimisations.) Not all estimates shown for clarity.

Fienup and Miller, 2003]. This gradient is derived for a generalised sharpness measure

$$S_\Omega = \sum_{x,y} W(x)\,\Omega\,[I(x,y)]$$

in Appendix B (B.8) as

$$\frac{\partial S_\Omega}{\partial \widehat{\phi}(v)} = \frac{2}{N_v} \sum_x W(x)\,\mathrm{Im}\left\{ \widehat{G}(x,v)\left(\mathcal{F}_{y\to v}\left\{ \widehat{g}(x,y)\frac{\partial\Omega}{\partial I} \right\} \right)^* \right\}, \qquad (6.80)$$

where $\widehat{G}(x,v) = G(x,v)\exp\left[-j\widehat{\phi}(v)\right]$ is the estimated signal history. The sharpness S_Ω will be maximum (or minimum) when the gradient is zero, giving

$$\sum_x W(x)\,\mathrm{Im}\left\{ G(x,v)\exp\left[-j\widehat{\phi}(v)\right]\left(\mathcal{F}_{y\to v}\left\{ \widehat{g}(x,y)\frac{\partial\Omega}{\partial I} \right\} \right)^* \right\} = 0. \qquad (6.81)$$

Rearranging gives

$$\mathrm{Im}\left\{ \exp\left[-j\widehat{\phi}(v)\right]\sum_x W(x)\,G(x,v)\left(\mathcal{F}_{y\to v}\left\{ \widehat{g}(x,y)\frac{\partial\Omega}{\partial I} \right\} \right)^* \right\} = 0. \qquad (6.82)$$

This can be solved directly giving

$$\widehat{\phi}(v) = n\pi + \angle \left\{ \sum_x W(x)\, G(x,v) \left(\mathcal{F}_{y \to v} \left\{ \widehat{g}(x,y) \frac{\partial \Omega}{\partial I} \right\} \right) * \right\}, \qquad (6.83)$$

for any integer n. Generally $n = 0$ will maximise the sharpness, whereas $n = 1$ will minimise it. Replacing the estimated image $\widehat{g}(x,y)$ with the measured image $g(x,y)$ and using recursion to converge, for $n = 0$ this becomes

$$\widehat{\phi}(v) = \angle \left\{ \sum_x W(x)\, G(x,v) \left(\mathcal{F}_{y \to v} \left\{ g(x,y) \frac{\partial \Omega}{\partial I} \right\} \right) * \right\}, \qquad (6.84)$$

For $S_2 = \sum_{x,y} I^2(x,y)$, $\frac{\partial \Omega}{\partial I} = 2I$, giving

$$\widehat{\phi}(v) = \angle \left\{ \sum_x W(x)\, G(x,v) \left(\mathcal{F}_{y \to v} \left\{ 2g(x,y)\, |g(x,y)|^2 \right\} \right) * \right\}, \qquad (6.85)$$

which is the same result as the direct sharpness maximisation derived earlier (6.72) using the approximation (6.74).

This is a far simpler derivation of the earlier result, confirming its validity. This result also adds useful insight into the method of direct sharpness maximisation. It can be considered an extension of the conjugate gradient method, allowing the phase giving maximum sharpness to be calculated directly. If the sharpness gradient is known at a single point, and the sharpness function is sinusoidal with a set period, it is intuitively possible to deduce where the peak in sharpness will be without having to measure the sharpness at many phase estimates. The need for recursion is also confirmed.

This alternate derivation also allows a significant extension to direct sharpness maximisation. Different sharpness measures, other than S_2, may be used. From (6.83), the phase maximising sharpness measure S_Ω is given by

$$\widehat{\phi}(v) = \angle \left\{ \sum_x W(x)\, G(x,v) \left(\mathcal{F}_{y \to v} \left\{ g(x,y) \frac{\partial \Omega}{\partial I} \right\} \right) * \right\}. \qquad (6.86)$$

Examples of the partial derivative of the sharpness function are

$$\frac{\partial \Omega}{\partial I} = \beta\, [I(x,y)]^{\beta-1} \qquad (6.87)$$

for the power law sharpness measure $\Omega[I(x,y)] = I(x,y)^\beta$, and

$$\frac{\partial \Omega}{\partial I} = \log[I(x,y)] + 1 \qquad (6.88)$$

for the negative entropy sharpness measure $\Omega[I(x,y)] = I(x,y)\log[I(x,y)]$. Results using direct sharpness maximisation with different sharpness measures are shown in

Fig. 6.11 and Table 6.4. Direct sharpness maximisation was successful using a small power-law metric $(S_{1,1})$ using (6.87) and negative entropy (S_{ent}) using (6.88). Successful results for a power-law metric, with $\beta < 1$ (for example $S_{0.5}$), were not obtained.

6.3.5 Extensions to direct sharpness maximisation

The gradient of the sharpness can be measured with respect to a path parameter for a parametric maximisation of sharpness, as shown in Appendix B. It may be possible to derive a method of direct parametric sharpness maximisation. This may have better performance (faster convergence) than a point by point maximisation. This has not been attempted. It is also not apparent if a regularised sharpness maximisation (see Chapter 5) can be performed directly. The method of direct sharpness maximisation has not been attempted on stripmap images.

6.4 SHARPNESS MAXIMISATION PERFORMANCE

This chapter shows that maximising S_2 may be considered a high-order phase difference estimation of the aberrated image Fourier phase. The limit of performance is set by the variance of the image phase. As shown in Chapter 4, this is set by the level of signal coherence, or signal to clutter ratio (SCR). This section develops the performance bound for sharpness maximisation and compares the performance of sharpness maximisation to this bound for various images and different range-weightings. The performance of sharpness maximisation is then directly compared to echo-correlation. The performance of different sharpness metrics are also compared.

6.4.1 Cramér-Rao lower bound

Echo-correlation methods estimate the Fourier phase difference. To make valid comparisons, the *phase-difference* $\Delta\widehat{\phi}[v] = \widehat{\phi}[v+1] - \widehat{\phi}[v]$ of the sharpness-maximisation (SM) estimate should be compared to the performance limit and performance of echo-correlation. The phase-difference has a greater effect on image degradation than phase alone [Callow, 2003], since a linear phase error leads to an image translation. The phase difference variance can be calculated using (6.40)

$$\text{Var}\left[\Delta\widehat{\phi}[v]\right] = 2\,\text{Var}\left[\widehat{\phi}[v]\right]. \tag{6.89}$$

The CRLB of an order M phase difference estimation, which uses the phase difference between M pings to estimate the phase error, is given by (4.100)

$$\text{Var}\left[\Delta\widehat{\phi}[v]\right] \geq \frac{1}{N_x'}\left[\frac{1}{\text{SCR}} + \frac{1}{M\,\text{SCR}^2}\right], \tag{6.90}$$

where $N_x' = N_x B_c / f_s$ is the number of independent range samples. Thus the bound for phase difference estimation $(M = 2)$ is given by

$$\text{Var}\left[\Delta\widehat{\phi}[v]\right] \geq \frac{1}{N_x'}\left[\frac{1}{\text{SCR}} + \frac{1}{2\text{SCR}^2}\right]. \tag{6.91}$$

For sharpness maximisation, the limit as $M \to \infty$ is given by

$$\text{Var}\left[\Delta\widehat{\phi}[v]\right] \geq \frac{1}{N_x'}\left[\frac{1}{\text{SCR}}\right], \tag{6.92}$$

thus the limit on the phase estimate is

$$\text{Var}\left[\widehat{\phi}[v]\right] \geq \frac{1}{2N_x'}\left[\frac{1}{\text{SCR}}\right]. \tag{6.93}$$

6.4.2 Results compared to bound

To compare the performance of sharpness maximisation to the CRLB and echo-correlation, a scene consisting of either a line (point target in each range-bin), or rough (speckled) block was simulated with varying SCR. Direct sharpness maximisation of S_2 and phase-difference estimation was used, and the variance of the phase difference measured. Fig. 6.6 compares different sharpness weightings. The coherence estimate of phase variance W_C performed the best, so results using this weighting are shown in Fig. 6.9. Also shown are the results of using weighted phase difference estimation (WPDE) as developed in Chapter 4. Fig. 6.9(a) and (b) shows that for a constant point target, sharpness-maximisation (SM) results meet the CRLB (6.92), while WPDE meets the CRLB for WPDE (6.91), slightly below. The difference in the bounds means that for low SCR, SM outperforms WPDE. For high SCR, the performance is similar. However, for a rough speckled block (Fig. 6.9(c) and (d)), SM performs significantly worse than WPDE. For an extended target, SM performance is well below the CRLB.

The reason for the poor performance of sharpness maximisation of an extended target can be explained by the nature of the coherence measured at larger separations. Echo-correlation estimates the phase difference between adjacent pings. Sharpness maximisation combines the phase difference between all pings. For a point target, the coherence between pings is constant at larger separations, as shown in Fig. 6.10(a). As the coherence is still high at large separation d, sharpness maximisation performs slightly better than echo-correlation as shown in Fig. 6.9. For an extended target, the coherence drops rapidly for larger separations between pings, as shown in Fig. 6.10(b). Sharpness maximisation does not combine the information at larger separations in an optimal way, as the performance of the technique is below the bound. If the coherence drops off, the performance of sharpness maximisation drops too. A reliable prediction of the performance of sharpness maximisation based on the variation of the coherence

(a)

(b)

(c)

(d)
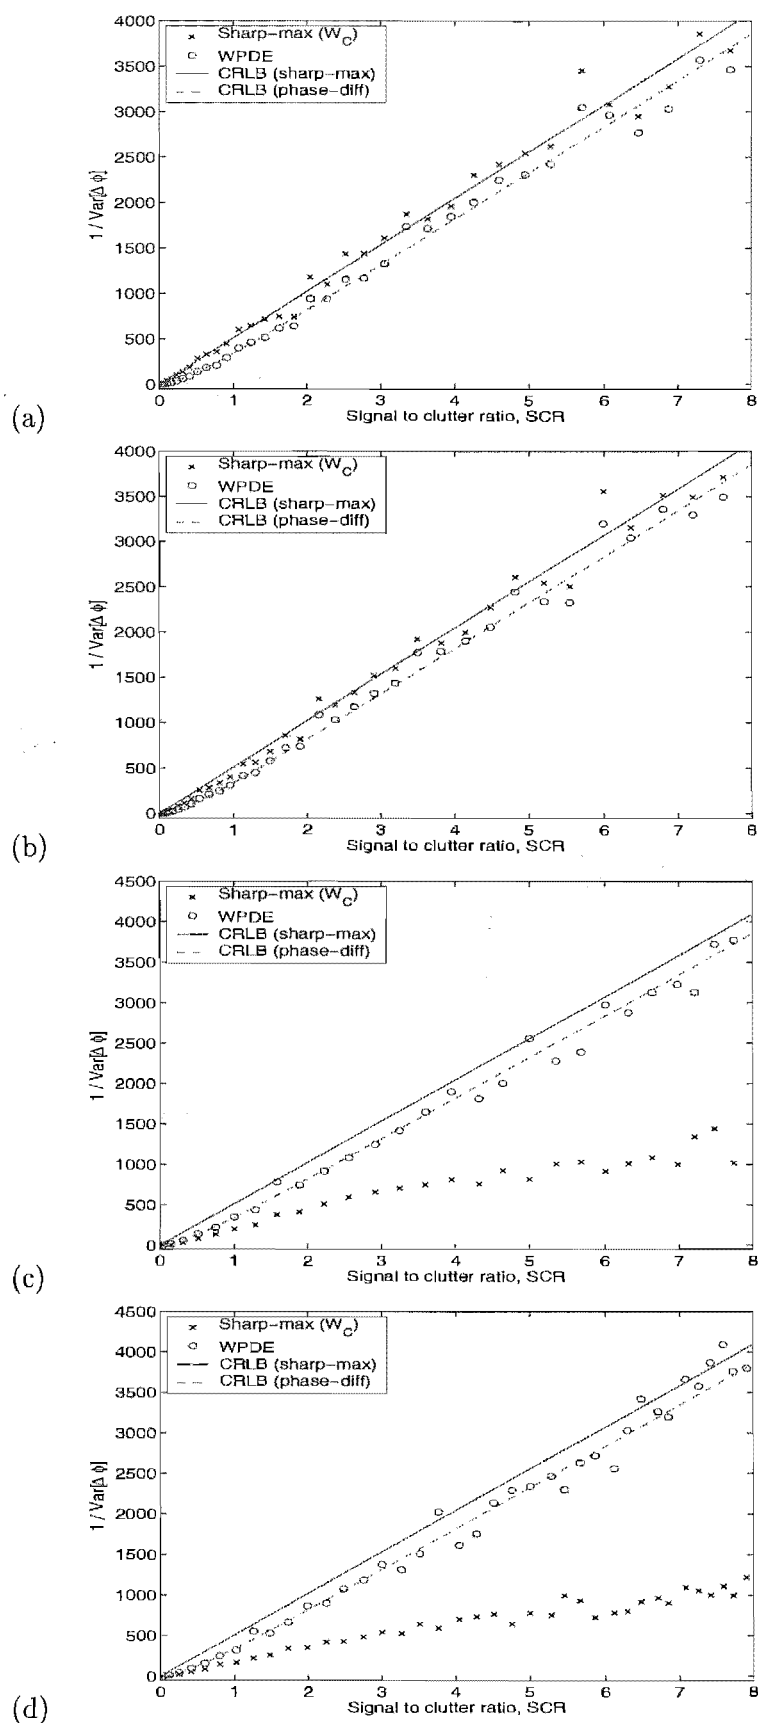
Figure 6.9   Inverse variance of error in phase difference for direct sharpness maximisation compared to weighted phase-difference estimation (WPDE) and the CRLB. Sharpness maximisation uses direct calculation with 30 iterations, metric S_2 and weighting W_C. WPDE (4.89) uses the same coherence weighting (4.73). (a) Point target, constant SCR with range. (b) Point target, background intensity varies with range. (c) Rough block target, constant SCR with range. (d) Rough block target, background intensity varies with range.

measure with range appears possible, but has not yet been determined.

6.4.3 Results using different sharpness metrics

The same experiment was performed using a direct maximisation of different sharpness metrics (6.86). The power-law metrics S_2 and $S_{1.1}$ and negative entropy S_{ent} were used. A weighted phase difference estimation (WPDE) was performed for comparison. All phase estimates used a coherence range-weighting W_C. Fig. 6.11(a) shows all techniques show similar performance for a point target. For a rough block target (Fig. 6.11(b)), as previously shown, the performance of S_2 falls well below the bound and WPDE. $S_{1.1}$ performs better, but still below the bound. The negative entropy measure (S_{ent}) performs better than the bound and WPDE. The bound was described for a high-order echo-correlation (or S_2), however this improvement in performance by negative entropy maximisation compared to echo-correlation is surprising. The reason for it is not clear. The experiment was repeated with additive white noise in the image, with similar results.

6.4.4 Results on different images

A different experiment was performed on a fixed set of images. The images used are shown in Figures 6.4 and 6.5 and described in Table 6.1. A random, known phase error was introduced into the images and the phase error estimated with both direct sharpness-maximisation using S_2 and phase difference estimation. The mean squared error in the phase difference was measured. This was repeated for 40 random initial path errors. The mean error is shown in Table 6.3.

Image	SCR	Direct sharpness maximisation				Phase difference estimation		
		W_N	W_R	W_C	CRLB	W_N	W_C	CRLB
A	1.687	0.0381	0.0068	0.0061	0.0012	0.0022	0.0013	0.0015
B	0.462	0.0184	0.0193	0.0155	0.0042	0.0104	0.0088	0.0088
C	0.552	0.0771	0.0075	0.0070	0.0035	0.0120	0.0058	0.0067
D	0.431	0.0827	0.0260	0.0310	0.0045	0.0200	0.0088	0.0098
E	8.366	0.00027	0.00025	0.00022	0.00023	0.00022	0.00022	0.00025
F	0.165	0.2714	0.1214	0.1547	0.0372	0.3241	0.1066	0.1230
G	1.057	0.0162	0.00860	0.00724	0.00307	0.00833	0.00716	0.00453
H	1.367	0.00465	0.00339	0.00285	0.00136	0.00660	0.00379	0.00187

Table 6.3 Mean squared error of phase difference estimates of various simulated images for different estimators. The images are described in Table 6.1 and shown in Fig. 6.4 and Fig. 6.5. Results are averaged over 40 trials with random initial phase error. W_N is with no weighting W_R is with range-normalised weighting (6.76) and W_C is with coherence inverse variance weighting (6.77).

(a)

(b)

(c)

(d)
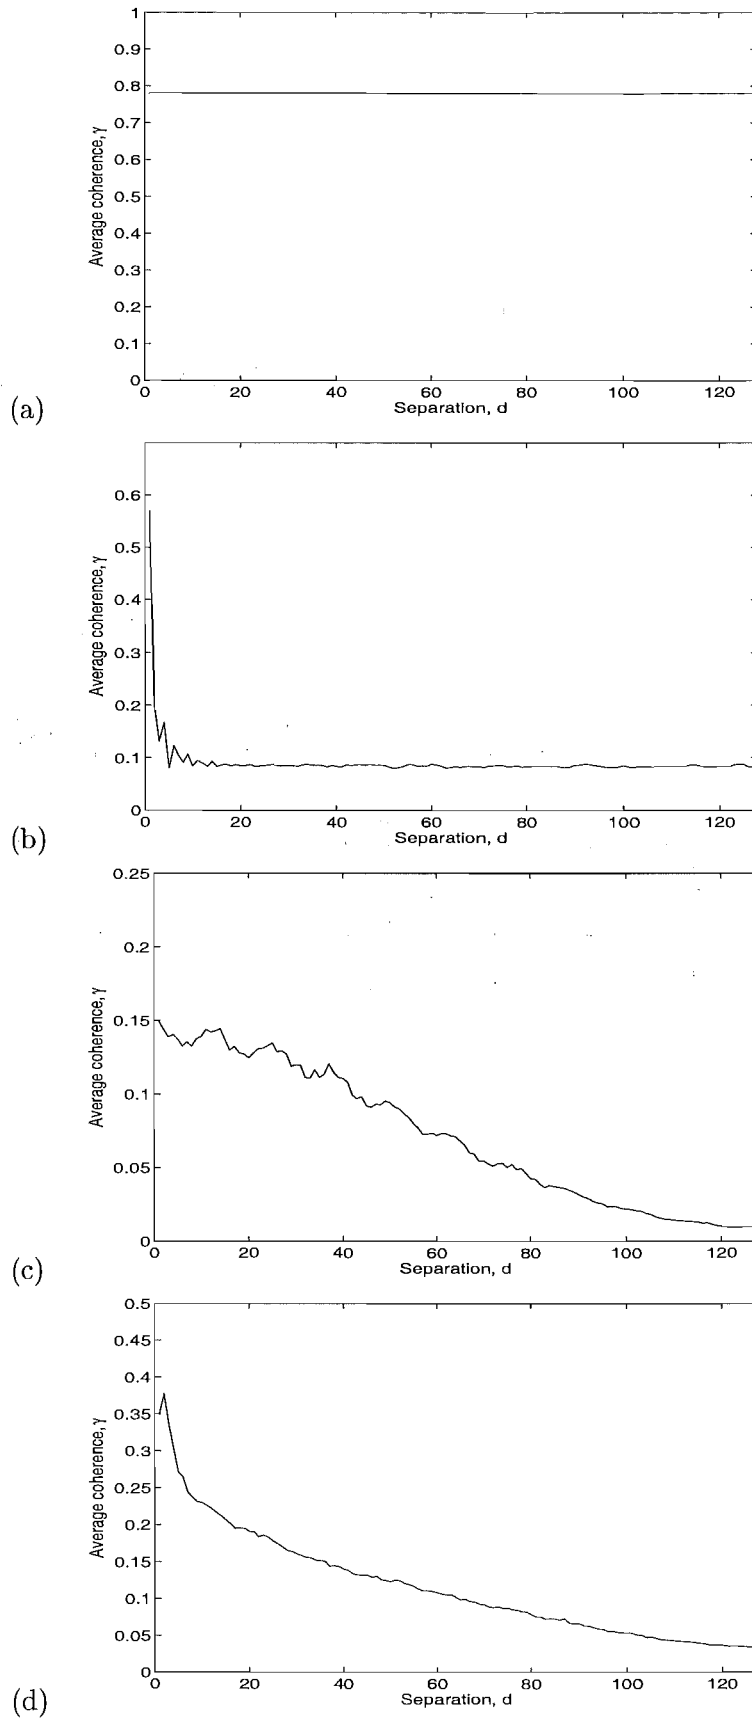
Figure 6.10   The variation of the average coherence between pings γ versus the separation between pings d. The images are described in Table 6.1 and shown in Fig. 6.4 and Fig. 6.5. (a) Point target in every range-bin (image E). (b) Rough, block target (image A). (c) Band-limited, scattered point targets (image F). (d) Simulated scene of mine-like objects (image H).
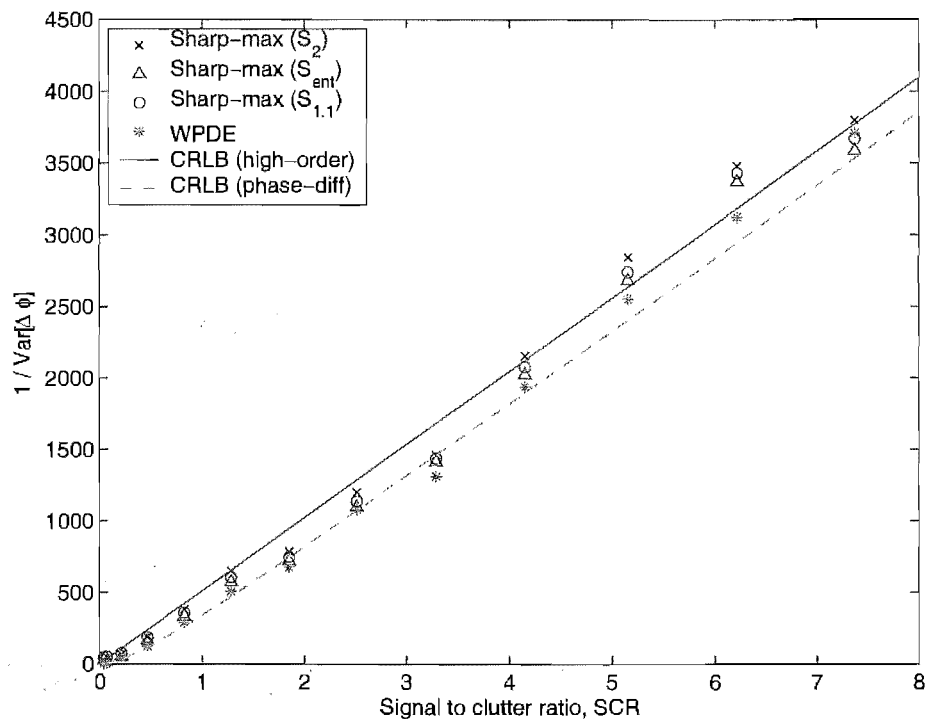
(a)
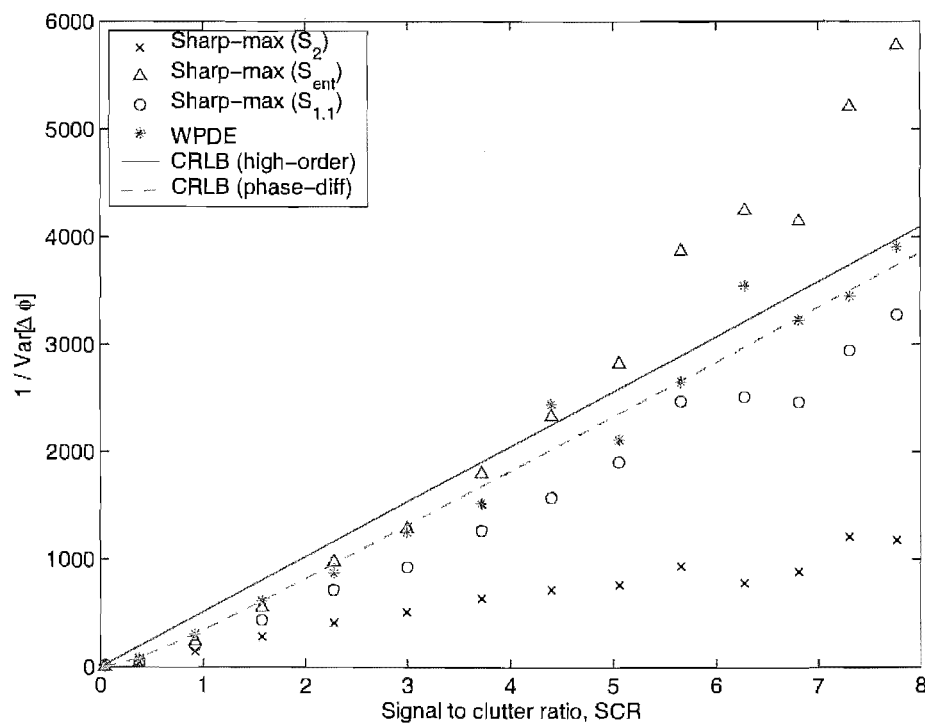
(b)

Figure 6.11 Results using different sharpness metrics. Inverse variance of error in phase difference for direct sharpness maximisation compared to weighted phase-difference estimation (WPDE) and the CRLB. Sharpness maximisation uses direct calculation with 30 iterations and weighting W_C. WPDE (4.89) uses the same coherence weighting (4.73). (a) Point target, background clutter intensity varies with range. (b) Rough block target, background clutter intensity varies with range.

Comparing the different weightings for sharpness maximisation (SM), the coherence measurement W_C gives the best performance for all images except D and F, for which W_R has slightly better performance. Using no range-weighting (W_N) has significantly worse performance, especially in images with significant shadow regions (images A, C, and D). Weighted phase difference estimation (WPDE) with W_C has significantly better performance than the unweighted version (W_N), again reinforcing the performance advantage of weighting for phase-difference estimation.

Comparing WPDE and SM with W_C, WPDE has better performance for the majority of images. SM has slightly better performance for images with dominant point targets (images E and H). This is expected, as SM was shown to perform poorer than WPDE for images consisting of extended speckled targets (images A,B, and D). The performance of the two techniques are similar for the most realistic images (F, G and H).

Different metrics

The previous experiment was repeated using direct sharpness maximisation with different sharpness metrics (6.86). The error in the estimated phase difference is shown in Table 6.4. These results show the negative entropy measure failed to converge for images F and H, performed similar to S_2 for image D, but performed significantly better than S_2 for all other images. The low power law metric $S_{1.1}$ performed significantly better than other measures for image D, and slightly better for most other images.

Image	S_2	$S_{1.1}$	S_{ent}	CRLB
A	0.0061	0.0014	0.0004	0.0012
B	0.0155	0.0181	0.0079	0.0042
C	0.0070	0.0050	0.0048	0.0035
D	0.0310	0.0021	0.0323	0.0045
E	0.000215	0.000225	0.000223	0.000233
F	0.1547	0.3274	-	0.0372
G	0.00724	0.00545	0.00518	0.00307
H	0.00285	0.00221	-	0.00136

Table 6.4 Mean squared error of phase difference estimates of various simulated images for different sharpness metrics. The images are described in Table 6.1 and shown in Fig. 6.4 and Fig. 6.5. Results are averaged over 40 trials with random initial phase error. All metrics used direct sharpness maximisation with coherence weighting W_C (6.77). A dash shows the estimator failed to converge.

Direct sharpness maximisation of negative entropy can fail to converge for images with sparse information. This includes images F and H in the above test. It has been shown that an iteration optimisation of sharpness does give a good solution for these images (see Section 5.7). So what went wrong? By looking at the phase of the expression inside the summation of the calculation (6.86), it appears that the

phase of the clutter biases the result when no target is present. When a target is present, the expression estimates the phase correctly. If the image contains sparse information, the bias can overwhelm the estimation, causing the estimator to fail to converge. The bias causes more problems for direct maximisation of entropy because the clutter phase is biased towards $\pm\pi$, causing the estimation to diverge. A bias towards zero slows convergence. A similar effect can be seen using the approximation for direct maximisation of S_2 (6.74). If no target is present, the phase of the approximate expression is biased to zero. Using the full expression, the phase of the clutter is uniformly distributed. More work is required to remove the clutter bias of direct sharpness maximisation.

6.5 CONCLUSIONS

This chapter has investigated the phase that maximises the sharpness of a coherent image. It has been shown to be equal to the Fourier phase of the aberrated image for a single range-bin. An ensemble of image phase estimates can be averaged to form an estimate of the Fourier phase error of the image. An aberrated spotlight image has a 1D Fourier phase error in the azimuth direction, so the estimated phase can be averaged over the range direction to form an estimate of the phase error. To improve this estimate, each range-bin may be weighted by the inverse variance of the phase estimate from that range-bin. A phase variance estimate can be made from the measured coherence between adjacent echos. As in echo-correlation, the variance of each phase estimate can be estimated using the measured signal coherence.

Image sharpness is normally measured in the image domain. This chapter has shown that the sharpness may be also measured in the Fourier, or signal domain. Furthermore, it shows that the phase that maximises image sharpness can be directly calculated in the signal domain. Direct sharpness-maximisation (SM) removes the need for multi-dimensional optimisation, or many iterated 1D optimisations. This removes the large computational hurdle of SM.

Maximisation of S_2 can be considered a form of high-order echo-correlation. Both SM and echo-correlation estimate the average Fourier phase of the aberrated image. They will have the same limit in performance, set by the variance of the image phase. This is determined by the signal coherence, or signal to clutter ratio (SCR). SM has a slightly better performance limit than a weighted phase-difference estimator (WPDE), since it is a higher-order estimator. On simulated images with delta-correlation SM meets the CRLB for a point target. However, for an extended rough (speckled) target, maximisation of S_2 falls considerably below the theoretical limit. This is due to the coherence between pings dropping away for higher separations. As echo-correlation uses adjacent pings, the performance is based on the coherence at a separation of $d = 1$. WPDE thus performs better than SM for an image without dominant point targets.

If the image has strong point targets, the performance of SM and WPDE are similar. WPDE is also simpler to compute. WPDE is thus preferred over SM to estimate phase errors in a spotlight synthetic aperture image.

By extending the method of conjugate gradient sharpness maximisation, it is possible to calculate a direct sharpness maximisation for a generalised sharpness function. Maximisation of a lower-power sharpness function, such as $S_{1.1}$ or negative entropy S_{ent}, gives improved performance compared to S_2 for extended targets and shadows. Maximisation of negative entropy exceeds the performance bound of S_2 and echo-correlation for a rough block, whereas S_2 fails to reach the bound. However, direct maximisation of negative entropy failed to converge for images with sparse information. This is due to a phase bias in the clutter. Future work is required to remove this bias.

Chapter 7

CONCLUSIONS

This thesis describes different methods of estimating the phase error in a synthetic aperture image. Clutter (speckle) in coherent imagery causes the echo from a rough surface to be a random variable. The phase of clutter is uniformly distributed. Clutter with a phase error also has uniformly distributed phase, so the phase error cannot be determined from clutter alone. The phase of a signal plus clutter has a variance determined by the signal to clutter ratio (SCR). Thus the variance of the phase estimate is limited by the signal to clutter ratio of the image. This thesis compares the results of different phase estimation techniques to this limit and to each other. Methods of improving the phase estimate are described.

7.1 ECHO CORRELATION

A narrow-band timing error results in a 1-D phase error over an image. This error can be estimated by averaging a measure of the image phase over a number of range-bins. Any bias caused by the phase of the signal should first be removed, leaving the phase error and a zero-mean random component caused by image speckle. The variance of this component is determined by the signal to clutter ratio (SCR) of that range-bin. The resulting estimate, when averaged over all range-bins, has a variance determined by the total variance of each estimate and the number of independent speckles averaged in range.

The variance of the estimate can be improved in two ways. The SCR can be increased by windowing the image spatially in the along-track direction. This works well for isolated point-like targets since they are localised in space. It does not work well for shadows, or dispersed targets such as trees or ripples. Another method is to weight the signal in range. Range-bins with a higher SCR have a lower phase variance and should contribute more to the phase error estimation. The standard method of range weighting is with the energy of the signal. This does not give the optimum result for images with shadows or dispersed targets. Weighting each phase estimate by the inverse variance of the unwrapped phase estimate gives the lowest variance estimate.

The phase variance can be estimated using a measure of the mean signal coherence between pings.

Phase bias can be removed from a spotlight image by centre shifting the image, then taking the phase difference. Directly estimating and removing the mean phase difference results in less phase bias than shifting the brightest point to the centre of the image, thus should be used in preference. Bias in stripmap images can be reduced by decreasing the weight of bright targets in the phase-difference image. Alternatively, the phase can be differentiated again, and the phase curvature estimated. Using a weighted phase curvature estimate (WPCE) increases performance by weighting in range. WPCE performs better than noncoherent shear-average. For images with a low SCR, amplitude weighted shear-average has better performance than WPCE. At higher SCR, WPCE performs better. WPCE is likely to be less accurate than other stripmap phase error estimation methods, such as SPGA. However it is a fast method to compute.

7.2 SHARPNESS MAXIMISATION

Speckle in a coherent image causes the image intensity and the sharpness to be a random variable. Chapter 3 describes the distribution of the image sharpness. The mean value of sharpness is proportional to the sharpness of the noncoherent image of the scene $V(x,y)$. The variance is proportional to the sharpness of $V^2(x,y)$ and inversely proportional to the number of independent speckles in the image.

A regularised form of sharpness maximisation is proposed. Regularisation reduces oversharpening by penalising unlikely phase error estimates. This results in a smoother phase estimate and generally a more accurate one.

Sharpness metrics can be characterised by the curvature, or second derivative of the sharpness function with respect to intensity. High power metrics, such as S_4, have increasing curvature with intensity, which weights the brightest points (highlights) more heavily. High power metrics were found to be noisier and more prone to oversharpening than other measures. Low power metrics, such as the square-root metric $S_{0.5}$ and negative entropy S_{ent}, have decreasing curvature with intensity, which weights the darkest points (shadows) more heavily. The classical intensity squared sharpness measure S_2 weights all values with the same curvature. Low-power metrics do not oversharpen, so are not improved with regularisation. Negative entropy (S_{ent}) worked well on all image types tested, making it the preferred measure of sharpness.

Different methods may be used to minimise the measured cost function. A cascaded, iterated, parametric optimisation was found to converge significantly faster than other methods tested. The path can be optimally represented by the fewest parameters using the platform motion statistics and a Karhunen-Loeve decomposition. A local optimisation method is deemed suitable for spotlight images. It is suitable for stripmap

images if the start-point is reasonably close to the solution. It is preferable that a prior motion estimation step is performed and the data corrected before performing sharpness maximisation of a stripmap image.

The phase that maximises the sharpness of a single range-bin is shown to be the Fourier phase of the aberrated image. Thus the variance of the phase estimate is equal to the variance of the image phase. Weighting each range bin by the inverse variance of the image phase results in a more accurate estimate of phase error. The same weighting as developed for echo-correlation may be used.

The phase that maximises sharpness can be calculated directly from the signal. Iterating this calculation is the same as performing an iterated non-parametric sharpness maximisation. Direct sharpness maximisation is significantly faster to compute since no optimisations are required. By extending the method of conjugate gradient sharpness maximisation, it is possible to calculate a direct sharpness maximisation for a generalised sharpness function. However, direct sharpness maximisation of negative entropy can be biased by clutter and fail to converge.

Maximisation of the classical intensity squared sharpness, S_2, is shown to be a form of high-order echo-correlation. The variance of the phase estimate is then bounded by the same limit as high-order echo-correlation. Maximisation of S_2 is shown to meet this limit for point-like objects, but perform below this limit for extended objects. Echo correlation is simpler to compute than sharpness maximisation and has better performance on an extended target. Thus a weighted phase-difference estimation is preferred.

Maximising negative entropy matches the performance limit for point-like objects, but is shown to exceed it for extended objects. Thus maximisation of negative entropy, weighted by the inverse variance of image phase, is the most accurate method of phase error estimation that was tested.

7.3 FUTURE RESEARCH

Suggestions for future research to extend, improve or verify methods discussed in this thesis are listed below.

Recommendations for **weighted phase difference estimation (WPDE)**:

- Implement WPDE within the framework of PGA and DPCA.

- Test performance of WPDE on real SAR or SAS data.

Recommendations for **weighted phase curvature estimation (WPCE)** of stripmap imagery:

- Compare WPCE to phase curvature autofocus (PCA).

- Develop a wideband version of WPCE.

- Test if sub-banding the data before WPCE reduces footprint shift decorrelation.

- For stripmap imagery, measure coherence over a local patch rather than all echos.

Recommendations for improving the method of **direct sharpness maximisation (DSM)**:

- Extend DSM to a parametric optimisation.

- Apply DSM to stripmap imagery.

- Remove the clutter bias that can cause DSM of negative entropy to fail.

- Extend DSM to power-law metric S_β for $\beta < 1$.

Recommendations for increasing understanding of **sharpness maximisation**:

- Determine the optimal metric, depending on measured image statistics.

- Determine why negative entropy maximisation outperforms other methods.

- Predict the performance of intensity squared metric based on coherence drop-off at higher separations.

- Estimate the yaw of multiple-hydrophone SAS using sharpness maximisation.

- Determine the statistical coupling of different motions such as yaw and sway, and use within the statistical autofocus framework.

Appendix A

DIRECT PHASE ESTIMATION FROM SHARPNESS

A common measure of image sharpness is calculated using

$$S_2 = \sum_y |g[y]|^4 .$$
(A.1)

This can be reformulated using the Fourier autocorrelation and energy theorems to be

$$S_2 = N \sum_{v_2} \left| \sum_{v_1} G[v_1]G^*[v_1 - v_2] \right|^2 ,$$
(A.2)

$$= N \sum_d \left| \sum_v G[v]G^*[v - d] \right|^2 ,$$
(A.3)

$$= N \sum_d |GG[d]|^2 ,$$
(A.4)

where $G[v]$ is the discrete Fourier transform of $g[y]$, i.e.,

$$G[v] = \frac{1}{N} \sum_y g[y] \exp\left(-j2\pi yv/N\right)$$
(A.5)

and where $GG[d]$ is the autocorrelation of $G[v]$ at lag d, i.e.,

$$GG[d] = \sum_v G[v]G^*[v - d] = \frac{1}{N} \sum_y |g[y]|^2 \exp\left(-j2\pi yv/N\right).$$
(A.6)

Note the image energy E is given by

$$E = NGG[0] = N \sum_v |G[v]|^2 = \sum_y |g[y]|^2 .$$
(A.7)

Consider adjusting the phase at $v = v_0$, so that

$$\tilde{G}[v, v_0, \phi] = \begin{cases} G[v] \exp\left(-j\phi\right) & v = v_0, \\ G[v] & \text{otherwise,} \end{cases}$$
(A.8)

then the adjusted image sharpness is

$$\tilde{S}_2[v_0, \phi] = N \sum_d \left| \tilde{GG}[d, v_0, \phi] \right|^2 . \tag{A.9}$$

The adjusted image spectrum can be written

$$\tilde{G}[v, v_0, \phi] = G[v] + G[v] \left(\exp\left(-j\phi\right) - 1 \right) \delta[v - v_0], \tag{A.10}$$

where

$$\delta[v] = \begin{cases} 1 & v = 0, \\ 0 & \text{otherwise.} \end{cases} \tag{A.11}$$

Using this definition, the autocorrelation of the adjusted image spectrum is

$$\tilde{GG}[d, v_0, \phi] = \sum_v \tilde{G}[v, v_0, \phi]\tilde{G}^*[v - d, v_0, \phi]. \tag{A.12}$$

Dropping the dependence on v_0 and ϕ, then

$$\begin{aligned} \tilde{GG}[d] = &\sum_v G[v]G^*[v - d] \\ &+ G[v_0]G^*[v_0 - d] \left(\exp\left(-j\phi\right) - 1 \right) \\ &+ G[v_0 + d]G^*[v_0] \left(\exp\left(j\phi\right) - 1 \right) \\ &+ |G[v_0]|^2 \left(\exp\left(-j\phi\right) - 1 \right) \left(\exp\left(j\phi\right) - 1 \right) \delta[d]. \end{aligned} \tag{A.13}$$

Equivalently,

$$\begin{aligned} \tilde{GG}[d] &= \sum_v G[v]G^*[v - d] + H[d] \left(1 - \delta[d] \right), \tag{A.14} \\ &= GG[d] + H[d] \left(1 - \delta[d] \right), \tag{A.15} \end{aligned}$$

where

$$H[d] = G[v_0]G^*[v_0 - d] \exp\left(-j\phi\right) + G[v_0 + d]G^*[v_0] \exp\left(j\phi\right) - G[v_0]G^*[v_0 - d] - G[v_0 + d]G^*[v_0]. \tag{A.16}$$

The sharpness is now

$$\begin{aligned} \tilde{S}_2[v_0, \phi] &= N \sum_d \Big[|GG[d]|^2 + |H[d]|^2 \left(1 - \delta[d] \right) \\ &\quad + 2\Re\left\{ GG[d]H^*[d] \right\} \left(1 - \delta[d] \right) \Big] \tag{A.17} \\ &= N \sum_d |GG[d]|^2 + \sum_{d'} |H[d']|^2 + \sum_{d'} 2\Re\left\{ GG[d]H^*[d] \right\}, \tag{A.18} \end{aligned}$$

where $d' = d, d \neq 0$. Now the first term is simply the sharpness of the original image,

S_2. The second term of (A.18) is

$$\sum_{d'} |H[d']|^2 = \sum_{d} |H[d]|^2 - |H[0]|^2, \qquad (A.19)$$

where

$$
\begin{aligned}
|H[d]|^2 = {} & 2\,|G[v_0]|^2\,|G[v_0 - d]|^2 \\
& + 2\,|G[v_0 + d]|^2\,|G[v_0]|^2 \\
& + 2\Re\left\{G[v_0]G^*[v_0 - d]G[v_0]G^*[v_0 + d]\exp\left(-j2\phi\right)\right\} \\
& - 2\Re\left\{G[v_0]G^*[v_0 - d]G[v_0 - d]G^*[v_0]\exp\left(-j\phi\right)\right\} \\
& - 2\Re\left\{G[v_0]G^*[v_0 - d]G[v_0]G^*[v_0 + d]\exp\left(-j\phi\right)\right\} \\
& - 2\Re\left\{G[v_0 + d]G^*[v_0]G[v_0 - d]G^*[v_0]\exp\left(j\phi\right)\right\} \\
& - 2\Re\left\{G[v_0 + d]G^*[v_0]G[v_0]G^*[v_0 + d]\exp\left(j\phi\right)\right\} \\
& + 2\Re\left\{G[v_0]G^*[v_0 - d]G[v_0]G^*[v_0 + d]\right\}
\end{aligned}
\qquad (A.20)
$$

Using the result that $\Re\{Z\} = \Re\{Z^*\}$, then

$$
\begin{aligned}
|H[d]|^2 = {} & 2\,|G[v_0]|^2\,|G[v_0 - d]|^2 \\
& + 2\,|G[v_0 + d]|^2\,|G[v_0]|^2 \\
& + 2\Re\left\{G[v_0]G^*[v_0 - d]G[v_0]G^*[v_0 + d]\exp\left(-j2\phi\right)\right\} \\
& - 2\Re\left\{G[v_0]G^*[v_0 - d]G[v_0 - d]G^*[v_0]\exp\left(-j\phi\right)\right\} \\
& - 2\Re\left\{G[v_0]G^*[v_0 - d]G[v_0]G^*[v_0 + d]\exp\left(-j\phi\right)\right\} \\
& - 2\Re\left\{G[v_0]G^*[v_0 + d]G[v_0]G^*[v_0 - d]\exp\left(-j\phi\right)\right\} \\
& - 2\Re\left\{G[v_0]G^*[v_0 + d]G[v_0 + d]G^*[v_0]\exp\left(-j\phi\right)\right\} \\
& + 2\Re\left\{G[v_0]G^*[v_0 - d]G[v_0]G^*[v_0 + d]\right\}.
\end{aligned}
\qquad (A.21)
$$

This can be simplified to

$$
\begin{aligned}
|H[d]|^2 = {} & 2\,|G[v_0]|^2\,|G[v_0 - d]|^2 \\
& + 2\,|G[v_0 + d]|^2\,|G[v_0]|^2 \\
& + 2\Re\left\{G^2[v_0]G^*[v_0 - d]G^*[v_0 + d]\exp\left(-j2\phi\right)\right\} \\
& - 2\,|G[v_0]|^2\,|G[v_0 - d]|^2\cos\phi \\
& - 2\Re\left\{G^2[v_0]G^*[v_0 - d]G^*[v_0 + d]\exp\left(-j\phi\right)\right\} \\
& - 2\Re\left\{G^2[v_0]G^*[v_0 - d]G^*[v_0 + d]\exp\left(-j\phi\right)\right\} \\
& - 2\,|G[v_0]|^2\,|G[v_0 - d]|^2\cos\phi \\
& + 2\Re\left\{G^2[v_0]G^*[v_0 - d]G^*[v_0 + d]\right\},
\end{aligned}
\qquad (A.22)
$$

and after grouping common factors,

$$
\begin{aligned}
|H[d]|^2 = {} & 2\,|G[v_0]|^2 \left(|G[v_0 - d]|^2 + |G[v_0 + d]|^2 \right) \\
& + 2\Re \left\{ G^2[v_0] G^*[v_0 - d] G^*[v_0 + d] \exp\left(-j2\phi\right) \right\} \\
& - 4\,|G[v_0]|^2 \, |G[v_0 - d]|^2 \cos\phi \\
& - 4\Re \left\{ G^2[v_0] G^*[v_0 - d] G^*[v_0 + d] \exp\left(-j\phi\right) \right\} \\
& + 2\Re \left\{ G^2[v_0] G^*[v_0 - d] G^*[v_0 + d] \right\}.
\end{aligned}
\tag{A.23}
$$

Summing over d yields,

$$
\begin{aligned}
\sum_d |H[d]|^2 = {} & 4\,|G[v_0]|^2\, GG[0] - 4\,|G[v_0]|^2\, GG[0]\cos\phi \\
& + 2\Re \left\{ G^2[v_0] \sum_d G^*[v_0 - d] G^*[v_0 + d] \exp\left(-j2\phi\right) \right\} \\
& - 4\Re \left\{ G^2[v_0] \sum_d G^*[v_0 - d] G^*[v_0 + d] \exp\left(-j\phi\right) \right\} \\
& + 2\Re \left\{ G^2[v_0] \sum_d G^*[v_0 - d] G^*[v_0 + d] \right\},
\end{aligned}
\tag{A.24}
$$

then after factoring

$$
\begin{aligned}
\sum_d |H[d]|^2 = {} & 4\,|G[v_0]|^2\, GG[0]\,(1 - \cos\phi) \\
& + 2\Re \left\{ G^2[v_0] \sum_d G^*[v_0 - d] G^*[v_0 + d] \left(\exp\left(-j2\phi\right) - 2\exp\left(-j\phi\right) + 1\right) \right\}.
\end{aligned}
\tag{A.25}
$$

But at $d = 0$,

$$
\begin{aligned}
|H[0]|^2 & = |G[v_0]|^4 \left[6 + 2\Re \left\{ \exp\left(-j2\phi\right) - 4\exp\left(-j\phi\right) \right\} \right], \tag{A.26} \\
& = |G[v_0]|^4 \left[6 + 2\cos(2\phi) - 8\cos\phi \right], \tag{A.27}
\end{aligned}
$$

so

$$
\begin{aligned}
\sum_{d'} |H[d']|^2 = {} & 4\,|G[v_0]|^2\, GG[0]\,(1 - \cos\phi) \\
& + 2\Re \left\{ G^2[v_0] \sum_d G^*[v_0 - d] G^*[v_0 + d] \left(\exp\left(-j\phi\right) - 1\right)^2 \right\} \\
& - |G[v_0]|^4 \left[6 + 2\Re \left\{ \exp\left(-j2\phi\right) - 4\exp\left(-j\phi\right) \right\} \right].
\end{aligned}
\tag{A.28}
$$

The third term of (A.18) is

$$\sum_{d'} 2\Re\left\{GG[d']H^*[d']\right\} = \sum_{d} 2\Re\left\{GG[d]H^*[d]\right\} - 2\Re\left\{GG[0]H^*[0]\right\}. \tag{A.29}$$

This is equivalent to

$$\sum_{d'} 2\Re\left\{GG[d']H^*[d']\right\} = 2\Re\left\{\sum_{d} GG[d]H^*[d]\right\} - 2GG[0]\Re\left\{H[0]\right\}, \tag{A.30}$$

where

$$\begin{aligned}
2\Re\left\{\sum_{d} GG[d]H^*[d]\right\} &= 2\Re\left\{G^*[v_0]\sum_{d}GG[d]G[v_0-d]\exp\left(j\phi\right)\right\} \\
&\quad + 2\Re\left\{G[v_0]\sum_{d}GG[d]G^*[v_0+d]\exp\left(-j\phi\right)\right\} \\
&\quad - 2\Re\left\{G^*[v_0]\sum_{d}GG[d]G[v_0-d]\right\} \\
&\quad - 2\Re\left\{G[v_0]\sum_{d}GG[d]G^*[v_0+d]\right\}.
\end{aligned} \tag{A.31}$$

After conjugating the first and third terms, then

$$\begin{aligned}
2\Re\left\{\sum_{d} GG[d]H^*[d]\right\} &= 2\Re\left\{G[v_0]\sum_{d}GG^*[d]G^*[v_0-d]\exp\left(-j\phi\right)\right\} \\
&\quad + 2\Re\left\{G[v_0]\sum_{d}GG[d]G^*[v_0+d]\exp\left(-j\phi\right)\right\} \\
&\quad - 2\Re\left\{G[v_0]\sum_{d}GG^*[d]G^*[v_0-d]\right\} \\
&\quad - 2\Re\left\{G[v_0]\sum_{d}GG[d]G^*[v_0+d]\right\}.
\end{aligned} \tag{A.32}$$

Then since $|g[y]|^2$ is real, $GG[v]$ is Hermitian, i.e., $GG[-v] = GG^*[v]$, and thus

$$\begin{aligned}
2\Re\left\{\sum_{d} GG[d]H^*[d]\right\} &= 4\Re\left\{G[v_0]\sum_{d}GG[d]G^*[v_0+d]\exp\left(-j\phi\right)\right\} \\
&\quad - 4\Re\left\{G[v_0]\sum_{d}GG[d]G^*[v_0+d]\right\}.
\end{aligned} \tag{A.33}$$

Combining terms,

$$2\Re\left\{\sum_d GG[d]H^*[d]\right\} = 4\Re\left\{G[v_0]\sum_d GG[d]G^*[v_0+d]\left(\exp\left(-j\phi\right)-1\right)\right\}. \quad \text{(A.34)}$$

At $d = 0$,

$$2\Re\left\{GG[0]H^*[0]\right\} = 2GG[0]\left|G[v_0]\right|^2 \Re\left\{\exp\left(j\phi\right)+\exp\left(-j\phi\right)-2\right\}, \quad \text{(A.35)}$$

$$= 4GG[0]\left|G[v_0]\right|^2\left(\cos\phi-1\right), \quad \text{(A.36)}$$

so

$$\sum_{d'} 2\Re\left\{GG[d']H^*[d']\right\} = 4\Re\left\{G[v_0]\sum_d GG[d]G^*[v_0+d]\left(\exp\left(-j\phi\right)-1\right)\right\}$$
$$- 4GG[0]\left|G[v_0]\right|^2\left(\cos\phi-1\right). \quad \text{(A.37)}$$

Combining the three terms of (A.18) yields

$$\tilde{S}_2[v_0,\phi] = S_2 + 8E\left|G[v_0]\right|^2\left(1-\cos\phi\right)$$
$$+ 4N\Re\left\{G[v_0]\sum_d GG[d]G^*[v_0+d]\left(\exp\left(-j\phi\right)-1\right)\right\}$$
$$+ 2N\Re\left\{G^2[v_0]\sum_d G^*[v_0-d]G^*[v_0+d]\left(\exp\left(-j\phi\right)-1\right)^2\right\} \quad \text{(A.38)}$$
$$- 2N\left|G[v_0]\right|^4\left[3+\Re\left\{\exp\left(-j2\phi\right)-4\exp\left(-j\phi\right)\right\}\right].$$

Let's define the following quantities,

$$K[v_0] = G[v_0]\sum_d GG[d]G^*[v_0+d], \quad \text{(A.39)}$$

$$L[v_0] = G^2[v_0]\sum_d G^*[v_0-d]G^*[v_0+d], \quad \text{(A.40)}$$

so that (A.38) becomes,

$$\tilde{S}_2[v_0,\phi] = S_2 + 8E\left|G[v_0]\right|^2\left(1-\cos\phi\right)$$
$$+ 4N\Re\left\{K[v_0]\left(\exp\left(-j\phi\right)-1\right)\right\}$$
$$+ 2N\Re\left\{L[v_0]\left(\exp\left(-j\phi\right)-1\right)^2\right\} \quad \text{(A.41)}$$
$$- 2N\left|G[v_0]\right|^4\left[3+\Re\left\{\exp\left(-j2\phi\right)-4\exp\left(-j\phi\right)\right\}\right].$$

Expanding out,

$$\tilde{S}_2[v_0, \phi] = S_2 + 8E \, |G[v_0]|^2 - 4N\Re\{K[v_0]\} + 2N\Re\{L[v_0]\} - 6N \, |G[v_0]|^4$$
$$+ 4N\Re\{K[v_0]\exp(-j\phi)\}$$
$$- 4N\Re\{L[v_0]\exp(-j\phi)\} + 2N\Re\{L[v_0]\exp(-j2\phi)\} \qquad \text{(A.42)}$$
$$- 2N \, |G[v_0]|^4 \cos(2\phi) + 8N \, |G[v_0]|^4 \cos(\phi)$$
$$- 8E \, |G[v_0]|^2 \cos(\phi).$$

Considering only the terms that vary with ϕ, this becomes

$$\tilde{S}'_2[v_0, \phi] = + 4N\Re\{K[v_0]\exp(-j\phi)\} \qquad \text{(A.43a)}$$
$$- 4N\Re\{L[v_0]\exp(-j\phi)\} + 2N\Re\{L[v_0]\exp(-j2\phi)\} \qquad \text{(A.43b)}$$
$$- 2N \, |G[v_0]|^4 \cos(2\phi) + 8N \, |G[v_0]|^4 \cos(\phi) \qquad \text{(A.43c)}$$
$$- 8E \, |G[v_0]|^2 \cos(\phi). \qquad \text{(A.43d)}$$

Consider the magnitude of the terms of each line separately. If N is large and $G[v]$ is wide-band, then $\sum_v |G[v]|^2 \gg |G[v_0]|^2$ and term (d) is much larger than term (c). Thus term (c) can be ignored. If $\angle\{G[v]\}$ varies with v, then $\sum_v |G[v]|^2 \gg \sum_d G^*[v_0 - d]G^*[v_0 + d]$. Thus term (d) is much larger than term (b) and term (b) can be ignored. A further simplification is possible as term (a) is much larger than term (b), making only term (a) a reasonable approximation of the varying sharpness. However this approximation is not necessary, so this term will be included.

Thus the varying sharpness can be approximated by

$$\tilde{S}'_2[v_0, \phi] \approx 4N\Re\left\{\left(K[v_0] - 2\frac{E}{N} \, |G[v_0]|^2\right)\exp(-j\phi)\right\}, \qquad \text{(A.44)}$$

or alternatively,

$$\tilde{S}'_2[v_0, \phi] \approx 4N\Re\{\chi[v_0]\exp(-j\phi)\}, \qquad \text{(A.45)}$$

where

$$\chi[v_0] = G[v_0]\sum_d GG[d]G^*[v_0 + d] - 2GG[0] \, |G[v_0]|^2. \qquad \text{(A.46)}$$

This can also be expressed as a weighted, higher-order, shear-average:

$$\chi[v_0] \approx \sum_d GG[d]G[v_0]G^*[v_0 + d]. \qquad \text{(A.47)}$$

This can be calculated using

$$\chi[v_0] \approx G[v_0]\mathcal{F}\left\{|g[y]|^2 \, g[y]\right\}^*. \qquad \text{(A.48)}$$

The phase estimate $\hat{\phi}$ that maximises the image sharpness is equivalent to maximising the real part of χ, i.e.,

$$\hat{\phi}[v_0] = \arg\max_{\phi} \left\{ \Re \left\{ \chi[v_0] \exp\left(-j\phi\right) \right\} \right\}, \tag{A.49}$$

The phase estimate $\hat{\phi}$ can thus be directly calculated using

$$\hat{\phi}[v_0] = \angle \left\{ \chi[v_0] \right\}. \tag{A.50}$$

All the phases can be estimated with one calculation:

$$\widehat{\phi}[v_0] = \angle \left\{ \left(\left. \mathcal{F}\left\{ |g[y]|^2 g[y] \right\} \right|_{[v_0]} \right)^* G[v_0] - 2\sum_v |G[v]^2| \, |G[v_0]|^2 \right\} \tag{A.51}$$

$$\approx \angle \left\{ \left(\left. \mathcal{F}\left\{ |g[y]|^2 g[y] \right\} \right|_{[v_0]} \right)^* G[v_0] \right\}. \tag{A.52}$$

Appendix B

SPOTLIGHT GRADIENT CALCULATION

Consider a spotlight image with (corrupted) signal history $G(x, v)$. x is the range co-ordinate, v is the slow-time co-ordinate. A phase-error estimate $\phi(v)$ is used to correct the phase history with a 1-D correction,

$$\widehat{G}(x, v) = G(x, v) \exp\left[-j\phi(v)\right]. \tag{B.1}$$

The corresponding complex-valued image is computed via a 1-D inverse discrete Fourier transform in the v direction,

$$
\begin{aligned}
\widehat{g}(x, y) &= \mathcal{F}_{v \to y}^{-1}\left\{ \widehat{G}(x, v) \right\} \\
&= 1/N_v \sum_v \widehat{G}(x, v) \exp\left[j2\pi vy/N_v\right] \\
&= 1/N_v \sum_v G(x, v) \exp\left[-j\phi(v)\right] \exp\left[j2\pi vy/N_v\right].
\end{aligned}
\tag{B.2}
$$

Taking the derivative with respect to $\phi(v)$,

$$
\begin{aligned}
\frac{\partial \widehat{g}(x, y)}{\partial \phi(v)} &= -j\frac{1}{N_v} G(x, v) \exp\left[-j\phi(v)\right] \exp\left[j2\pi vy/N_v\right] \\
&= -j\frac{1}{N_v} \widehat{G}(x, v) \exp\left[j2\pi vy/N_v\right].
\end{aligned}
\tag{B.3}
$$

Similarly, for the conjugate of the image $\widehat{g}^*(x, y)$,

$$
\begin{aligned}
\frac{\partial \widehat{g}^*(x, y)}{\partial \phi(v)} &= j\frac{1}{N_v} G^*(x, v) \exp\left[j\phi(v)\right] \exp\left[-j2\pi vy/N_v\right] \\
&= \left(\frac{\partial \widehat{g}(x, y)}{\partial \phi(v)} \right)_*
\end{aligned}
\tag{B.4}
$$

The derivative of the image intensity $I(x, y) = |\widehat{g}(x, y)|^2 = \widehat{g}(x, y)\widehat{g}^*(x, y)$ with

respect to $\phi(v)$ is given by

$$
\begin{aligned}
\frac{\partial I(x,y)}{\partial \phi(v)} &= \widehat{g}^*(x,y)\frac{\partial \widehat{g}(x,y)}{\partial \phi(v)} + \widehat{g}(x,y)\frac{\partial \widehat{g}^*(x,y)}{\partial \phi(v)} \\
&= -j\frac{1}{N_v}\widehat{g}^*(x,y)\widehat{G}(x,v)\exp\left[j2\pi vy/N_v\right] + j\frac{1}{N_v}\widehat{g}(x,y)\widehat{G}^*(x,v)\exp\left[-j2\pi vy/N_v\right] \\
&= \frac{2}{N_v}\mathrm{Im}\left\{\widehat{g}^*(x,y)\widehat{G}(x,v)\exp\left[j2\pi vy/N_v\right]\right\}.
\end{aligned}
$$

$$(B.5)$$

Now consider a sharpness measure of the image which consists of the sum of a non-linear point transform of the image intensity pixels,

$$
S_\Omega = \sum_{x,y} W(x)\,\Omega\left[I(x,y)\right],
\tag{B.6}
$$

where $W(x)$ is a range-weighting function. Taking the derivative with respect to $\phi(v)$,

$$
\begin{aligned}
\frac{\partial S_\Omega}{\partial \phi(v)} &= \sum_{x,y} W(x)\frac{\partial \Omega\left[I(x,y)\right]}{\partial I(x,y)}\frac{\partial I(x,y)}{\partial \phi(v)} \\
&= \frac{2}{N_v}\sum_{x,y} W(x)\frac{\partial \Omega}{\partial I}\mathrm{Im}\left\{\widehat{g}^*(x,y)\widehat{G}(x,v)\exp\left[j2\pi vy/N_v\right]\right\}.
\end{aligned}
$$

$$(B.7)$$

Rearranging and taking the summation over y,

$$
\begin{aligned}
\frac{\partial S_\Omega}{\partial \phi(v)} &= \frac{2}{N_v}\sum_{x} W(x)\,\mathrm{Im}\left\{\widehat{G}(x,v)\left(\sum_{y}\frac{\partial \Omega}{\partial I}\widehat{g}(x,y)\exp\left[-j2\pi vy/N_v\right]\right)^*\right\} \\
&= \frac{2}{N_v}\sum_{x} W(x)\,\mathrm{Im}\left\{\widehat{G}(x,v)\left(\mathcal{F}_{y\to v}\left\{\widehat{g}(x,y)\frac{\partial \Omega}{\partial I}\right\}\right)^*\right\}.
\end{aligned}
$$

$$(B.8)$$

For $S_2 = \sum_{x,y} I^2(x,y)$, $\frac{\partial \Omega}{\partial I} = 2I$, giving

$$
\frac{\partial S_2}{\partial \phi(v)} = \frac{2}{N_v}\sum_{x} W(x)\,\mathrm{Im}\left\{\widehat{G}(x,v)\left(\mathcal{F}_{y\to v}\left\{2\widehat{g}(x,y)\,|\widehat{g}(x,y)|^2\right\}\right)^*\right\}.
\tag{B.9}
$$

This derivation is also shown by Fienup [2000].

B.0.1   Parametric optimisation

Consider representing the phase error estimate as a set of basis functions $U_k(v)$, with coefficients b_k. The phase error is then represented as

$$
\phi(v) = \sum_{k=1}^{N_p} b_k U_k(v).
\tag{B.10}
$$

The gradient of the sharpness metric with respect to the parameters b_k is given by the chain rule as

$$\frac{\partial S_\Omega}{\partial b_k} = \sum_{x,y} W(x) \frac{\partial \Omega\left[I(x,y)\right]}{\partial I(x,y)} \sum_v \frac{\partial I(x,y)}{\partial \phi(v)} \frac{\partial \phi(v)}{\partial b_k} \quad \text{(B.11)}$$

$$= \sum_v U_k(v) \frac{\partial S_\Omega}{\partial \phi(v)}, \quad \text{(B.12)}$$

which is the projection of the nonparametric gradient onto the basis set [Fienup and Miller, 2003].

REFERENCES

[Attia and Steinberg, 1989] E. H. Attia and B. Steinberg. Self-cohering large antenna arrays using the spatial correlation properties of radar clutter. *IEEE Transactions on Antennas and Propagation*, 37(1):30–38, January 1989.

[Bamler and Hartl, 1998] R. Bamler and P. Hartl. Topical review: Synthetic aperture radar interferometry. *Inverse Problems*, 14(3):R1–R54, August 1998.

[Bamler, 1992] R. Bamler. A comparison of range-doppler and wavenumber domain SAR focusing algorithms. *IEEE Transactions on Aerospace and Electronic Systems*, 30(4):706–713, July 1992.

[Barclay et al., 2005] P. J. Barclay, M. P. Hayes, and P. T. Gough. MAP estimation of seafloor topography using multi-frequency synthetic aperture sonar. In *Oceans 2005 Europe*, Brest, France, July 2005. IEEE.

[Bates and McDonnell, 1986] R. H. T. Bates and M. J. McDonnell. *Image Restoration and Reconstruction*. Clarendon Press, Oxford, 1986.

[Beck and Arnold, 1977] J. V. Beck and K. J. Arnold. *Parameter Estimation in Engineering and Science*. John Wiley & Sons, 1977.

[Bellettini and Pinto, 2002] A. Bellettini and M. Pinto. Theoretical accuracy of synthetic aperture sonar micronavigation using a displaced phase-center antenna. *IEEE Journal of Oceanic Engineering*, 27(4):780–789, October 2002.

[Benesty et al., 2004] J. Benesty, J. Chen, and Y. Huang. Time-delay estimation via linear interpolation and cross correlation. *IEEE Transactions on Speech and Audio Processing*, 12(5):509–519, September 2004.

[Berizzi and Corsini, 1996] F. Berizzi and G. Corsini. Autofocusing of inverse synthetic aperture radar images using contrast optimization. *IEEE Transactions on Aerospace and Electronic Systems*, 32(3):1185–1191, July 1996.

[Berizzi et al., 1996] F. Berizzi, G. Corsini, M. Diani, and M. Veltroni. Autofocus of wide azimuth angle SAR images by contrast optimisation. In *International Geoscience and Remote Sensing Symposium*, volume 2, pages 1230–1232, 1996.

[Blacknell and Quegan, 1991] D. Blacknell and S. Quegan. SAR motion compensation using autofocus. *International Journal Remote Sensing*, 12(2):253–275, 1991.

[Blacknell et al., 1992] D. Blacknell, A. P. Blake, C. J. Oliver, and R. G. White. A comparison of SAR multilook registration and contrast optimisation autofocus alogrithms applied to real SAR data. In *International Radar Conference*, pages 363–366, 1992.

[Bracewell, 1986] R. N. Bracewell. *The Fourier transform and its applications*. McGraw-Hill Book Company, 1986.

[Buffington et al., 1977] A. Buffington, F. S. Crawford, R. A. Muller, A. J. Schwemin, and R. G. Smits. Correction of atmospheric distortion with an image-sharpening telescope. *Journal of the Optical Society of America*, 67(3):298–303, March 1977.

[Cafforio et al., 1991] C. Cafforio, C. Pratti, and F. Rocca. SAR data focussing using seismic migration techniques. *IEEE Transactions on Aerospace and Electronic Systems*, 27(2):194–207, March 1991.

[Callow et al., 2001a] H. J. Callow, M. P. Hayes, and P. T. Gough. Autofocus of multi-band, shallow-water, synthetic aperture sonar imagery using shear-averaging. In *IGARSS 2001, International Geoscience and Remote Sensing Symposium*, pages 1601–1603. IEEE, July 2001.

[Callow et al., 2001b] H. J. Callow, M. P. Hayes, and P. T. Gough. Noncoherent autofocus of single receiver, broad-band synthetic aperture sonar imagery. In *Oceans 2001, Marine Technology and Ocean Science Conference*, volume 1, pages 157–162, November 2001.

[Callow et al., 2003] H. J. Callow, M. P. Hayes, and P. T. Gough. Stripmap phase gradient autofocus. In *Oceans 2003, Marine Technology and Ocean Science Conference*, volume 1, pages 2414–2421, 2003.

[Callow et al., 2004] H. J Callow, M. P. Hayes, and P. T. Gough. Motion compensation improvements for widebeam multiple transducer SAS systems. *IEEE Journal of Oceanic Engineering*, 2004. Accepted for publication, submitted July 2002.

[Callow, 2003] H. J. Callow. *Signal Processing for Synthetic Aperture Sonar Image Enhancement*. PhD thesis, Department of Electrical and Electronic Engineering, University of Canterbury, 2003.

[Carrera et al., 1995] W. G. Carrera, R. S. Goodman, and R. M. Majewski. *Spotlight synthetic aperture radar: signal processing algorithms*. Artech House, 1995.

[Carter et al., 1973] G. C. Carter, C. H. Knapp, and A. H. Nuttall. Estimation of the magnitude-squared coherence function via overlapped fast fourier transform processing. *IEEE Transactions on Audio Electroacoustics*, 21(4):337–389, August 1973.

[Christoff, 1998] J. T. Christoff. Motion-compensated high frequency synthetic aperture sonar. *Journal of the Acoustics Society of America*, 1:2950, 1998.

[Cumming et al., 1992] I. G. Cumming, F. H. Wong, and R. K. Raney. A SAR processing alogrithm with no interpolation. In *IEEE Transactions on Geoscience and Remote Sensing*, volume 1, pages 376–379. IEEE, 1992.

[Curlander and McDonough, 1996] J. C. Curlander and R. N. McDonough. *Synthetic Aperture Radar: Systems and signal processing*. John Wiley and sons, New York, 1996.

[Dainty, 1975] J. C. Dainty, editor. *Laser Speckle and Related Phenomena*. Springer-Verlag, Berlin, Heidelberg, New York, 1975.

[Douglas and Lee, 1993] B. L. Douglas and H. Lee. Synthetic-aperture sonar imaging with a multiple-element receiver array. In *IEEE International conference on Acoustics, Speech, and Signal Processing*, volume 5, pages 445–448. IEEE, April 1993.

[Dunlop, 1997] J. Dunlop. Statistical modelling of sidescan sonar images. In *Oceans 1997, Marine Technology and Ocean Science Conference*, volume 1, pages 33–38, October 1997.

[Eichel and Jakowatz, 1989] P. H. Eichel and C. V. Jakowatz, Jnr. Phase-gradient algorithm as an optimal estimator of the phase derivative. *Optics Letters*, 14(20):1101–1103, October 1989.

[Eichel et al., 1989] P. H. Eichel, D. C. Ghiglia, and C. V. Jakowatz, Jnr. Speckle processing method for synthetic-aperture-radar phase correction. *Optics Letters*, 14(1):1–3, January 1989.

[Elster and Neumaier, 1995] C. Elster and A. Neumaier. A grid algorithm for bound constrained optimisation of noisy functions. *IMA Journal of Numerical Analysis*, 15:585–608, 1995.

[Fienup and Miller, 2003] J. R. Fienup and J. J. Miller. Aberration correction by maximizing generalized sharpness metrics. *Journal of the Optical Society of America A*, 20(4):609–620, April 2003.

[Fienup, 1989] J. R. Fienup. Phase error correction by shear averaging. In *Signal Recovery and synthesis III*, pages 14–16. Optical Society of America, 1989.

[Fienup, 1997] J. R. Fienup. Invariant error metrics for image reconstruction. *Applied Optics*, 36(32):8352–8357, November 1997.

[Fienup, 2000] J.R. Fienup. Synthetic-aperture radar autofocus by maximizing sharpness. *Optics Letters*, 25(4):221–223, February 2000.

[Finley and Wood, 1985] I. P. Finley and J. W. Wood. An investigation of synthetic aperture radar autofocus. Technical Report Memorandum 3790, Royal Signals and Radar Establishment, April 1985.

[Forsythe et al., 1976] G. E. Forsythe, M. A. Malcolm, and C. B. Moler. *Computer Methods for Mathematical Computations*. Pretice-Hall, 1976.

[Fortune et al., 2001a] S. A. Fortune, P. T. Gough, and M. P. Hayes. Statistical autofocus of synthetic aperture sonar images using image contrast optimisation. In *International Geoscience and Remote Sensing Symposium Proceedings, 2001*. IEEE, 2001.

[Fortune et al., 2001b] S. A. Fortune, M. P. Hayes, and P. T. Gough. Statistical autofocus of synthetic aperture sonar images using image contrast optimisation. In *Oceans 2001, Marine Technology and Ocean Science Conference*, volume 1, pages 163–169, November 2001.

[Fortune et al., 2002] S. A. Fortune, M. P. Hayes, and P. T. Gough. Contrast optimisation of coherent images. In *IVCNZ2002, Image and Vision Computing New Zealand 2002*, pages 299–304. IVCNZ, November 2002.

[Fortune et al., 2003a] S. A. Fortune, M. P. Hayes, and P. T. Gough. Contrast optimisation of coherent images. In *Oceans 2003, Marine Technology and Ocean Science Conference*, pages 2622–2628, 2003.

[Fortune et al., 2003b] S. A. Fortune, M. P. Hayes, and P. T. Gough. Speckle reduction of synthetic aperture sonar images. In *World Congress on Ultrasonics, Paris, September 2003*, pages 669–672, 2003.

[Fortune et al., 2004] S. A. Fortune, M. P. Hayes, and P. T. Gough. Statistics of the contrast of coherent images. *Journal of the Optical Society of America A*, 21(7):1131–1139, July 2004.

[Gensane, 1989] M. Gensane. A statistical study of acoustic signals backscattered from the sea bottom. *IEEE Journal of Oceanic Engineering*, 14(1):84–93, January 1989.

[George and Sinclair, 1976] N. George and D. C. Sinclair. Speckle in optics. *Journal of the Optical Society of America*, 66(11):1143–1316, November 1976.

[Goodman, 1968] J. W. Goodman. *Introduction to Fourier Optics*. McGraw-Hilll, New York, 1968.

[Goodman, 1975] J. W. Goodman. Statistical properties of laser speckle patterns. In J. C. Dainty, editor, *Laser speckle and related phenomena*, pages 9–75. Springer-Verlag, Berlin, 1975.

[Goodman, 1976] J. W. Goodman. Some fundamental properties of speckle. *Journal of the Optical Society of America*, 66(5):1145–1150, 1976.

[Goodman, 1986] J. W. Goodman. A random walk through the field of speckle. *Optical Engineering*, 25(5):610–612, May 1986.

[Gough and Hawkins, 1997] P. T. Gough and D. W. Hawkins. Imaging algorithms for a strip-map synthetic aperture sonar: Minimizing the effects of aperture errors and aperture undersampling. *IEEE Journal of Oceanic Engineering*, 22(1):27–39, January 1997.

[Gough and Hawkins, 1998] P. T. Gough and D. W. Hawkins. A short history of synthetic aperture sonar. In *Geoscience and Remote Sensing Symposium Proceedings, 1998*, volume 3, pages 618–620. IEEE, 1998.

[Gough and Lane, 1998] P. T. Gough and R. G. Lane. Autofocussing SAR and SAS images using a conjugate gradient search alogirithm. In *Geoscience and Remote Sensing Symposium Proceedings, 1998*, volume 2, pages 621–623. IEEE, 1998.

[Gough and Miller, 2004] P. T. Gough and M. A. Miller. Displaced ping imaging autofocus for a multi-hydrophone SAS. *IEE Proceedings on Radar, Sonar, and Navigation*, 151(3):163–170, June 2004.

[Hamaker *et al.*, 1977] J. P. Hamaker, J. D. O'Sullivan, and J. E. Noordam. Image sharpness, fourier optics, and redundant-spacing interferometery. *Journal of the Optical Society of America*, 67(8):1122–1123, August 1977.

[Hansen *et al.*, September 2003] R. E. Hansen, T. O. Saeb, K. Gade, and S. Chapman. Signal processing for AUV based interferometric synthetic aperture sonar. In *Proceedings from Oceans 2003 MTS/IEEE*, San Diego, CA, USA, September 2003.

[Hawkins, 1996] D. W. Hawkins. *Synthetic Aperture Imaging Algorithms: with application to wide bandwidth sonar*. PhD thesis, Department of Electrical and Electronic Engineering, University of Canterbury, October 1996.

[Hayes *et al.*, 2002] M. P. Hayes, H. J. Callow, and P. T. Gough. Strip-map phase gradient autofocus. In David Kenwright, editor, *Image and vision computing New Zealand 2002*, pages 71–76, Auckland, December 2002.

[Henderson and Lewis, 1998] F. M. Henderson and A. J. Lewis, editors. *Principles and Applications of Imaging Radar*, volume 2 of *Manual of Remote Sensing*. John Wiley & Sons, 3rd edition, 1998.

[Hunter, 2005] A. Hunter. *Underwater acoustic modelling for Synthetic Aperture Sonar*. PhD thesis, Department of Electrical and Electronic Engineering, University of Canterbury, 2005.

[Jakeman and Pusey, 1976] E. Jakeman and P. N. Pusey. A model for non-rayleigh sea echo. *IEEE Transactions on Antennas and Propagation*, AP-24(6):806–814, November 1976.

[Jakowatz and Wahl, 1993] C. V. Jakowatz, Jnr. and D. E. Wahl. Eigenvector method for maximum-likelihood estimation of phase errors in synthetic-aperture-radar imagery. *Journal of the Optical Society of America*, 10(12):2539–2546, December 1993.

[Jakowatz et al., 1996] C. V. Jakowatz, Jnr., D. E. Wahl, P. H. Eichel, D. C Ghiglia, and P. A. Thompson. *Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach*. Kluwer Academic Publishers, Boston, 1996.

[Jarvis, 1997] A. R. Jarvis. Function minimisation in the presence of noise using methods that do not require gradient information. Master's thesis, University of Canterbury, 1997.

[Jayant and Noll, 1984] N. Jayant and P. Noll. *Digital Coding of Waveforms*. Prentice Hall, 1984.

[Jenkins and Watts, 1968] G. M. Jenkins and D. G. Watts. *Spectral analysis and its applications*. Holden-Day, San Fransisco, CA, 1968.

[Johnson et al., 1995] K. A. Johnson, M. P. Hayes, and P. T. Gough. Estimating subwavelength sway of sonar towfish. *IEEE Journal of Oceanic Engineering*, 20(4):258–267, October 1995.

[Just and Bamler, 1994] D. Just and R. Bamler. Phase statistics of interferograms with applications to synthetic aperture radar. *Applied Optics*, 33(20):4361–4368, July 1994.

[Knapp and Carter, 1976] C. H. Knapp and G. C. Carter. The generalised correlation method for estimation of time delay. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 24(4):320–327, August 1976.

[Knox and Thompson, 1974] K. T. Knox and B. J. Thompson. Recovery of images from atmospherically degraded short-exposure photographs. *Astrophysical Journal*, 193:L45–L48, Oct. 1974.

[Lagarias et al., 1998] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright. Convergence properties of the nelder-mead simplex method in low dimensions. *SIAM Journal of Optimization*, 9(1):112–147, 1998.

[Lee et al., 1994] J-S. Lee, K. W. Hoppel, S. A. Mango, and A. R. Miller. Intensity and phase statastics of multilook polarimetric and interferometric SAR imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 32(5):1017–1028, September 1994.

[Lee, 1986] J. S. Lee. Speckle suppression and analysis for synthetic aperture radar images. *Optical Engineering*, 25(5):636–643, May 1986.

[Li, 1992] A. Li. Algorithms for the implementation of stolt interpolation in SAR processing. In *International Geoscience and Remote Sensing Symposium*, volume 1, pages 360–362, 1992.

[Lowenthal and Arsenault, 1970] S. Lowenthal and H. Arsenault. Image formation for coherent diffuse objects: Statistical properties. *Journal of the Optical Society of America*, 60(11):1478–1483, November 1970.

[Marron and Morris, 1986] J. Marron and G. M. Morris. Image recognition in the presence of laser speckle. *Journal of the Optical Society of America*, 3:964–971, 1986.

[MathWorks, 1994] MathWorks. Matlab optimization toolbox. The Math Works, Natick, Massachusetts, USA, 1994.

[Morrison and Munson, 2002] R. L. Jr. Morrison and D. C. Jr. Munson. An experimental study of a new entropy-based SAR autofocus technique. In *International Conference on Image Processing 2002, Proceedings.*, volume 2, pages 441–444, 2002.

[Morrison, 2002] R. L. Morrison. Entropy based autofocus for synthetic aperture radar. Master's thesis, University of Illinois, 2002.

[Muller and Buffington, 1974] R. A. Muller and A. Buffington. Real-time correction of atmospherically degraded telescope images through image sharpening. *Journal of the Optical Society of America*, 64(9):1200–1210, September 1974.

[Nelder and Mead, 1965] J. A. Nelder and R. Mead. A simplex method for function minimisation. *Computer Journal*, 7(4):308–313, 1965.

[Otten and van Ginneken, 1989] R.H.J.M. Otten and L.P.P.P. van Ginneken. *The annealing algorithm*. Kluwer Academic Publishers, Boston, 1989.

[Pat, 2000] J. B. Pat. Synthetic aperture sonar image reconstruction using a multiple-receiver towfish. Master's thesis, Department of Electrical and Electronic Engineering, University of Canterbury, March 2000.

[Paxman and Marron, 1988] R. G. Paxman and J. C. Marron. Abberation correction of speckled imagery with an image-sharpness criterion. In G. M. Morris, editor, *Statistical Optics, Proceedings of SPIE*, volume 976, pages 37–47, 1988.

[Quazi, 1981] A. H. Quazi. An overview on the time delay estimate in active and passive systems for target localization. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 29(3):527–533, June 1981.

[Raney and Wessels, 1988] R. K. Raney and G. J. Wessels. Spatial considerations in SAR speckle simulation. *IEEE Transactions on Geoscience and Remote Sensing*, 26(5):666–672, September 1988.

[Raney et al., 1994] R. K. Raney, H. Runge, R. Bamler, I. G. Cumming, and F. H. Wong. Precision SAR processing using chirp scaling. *IEEE Transactions on Geoscience and Remote Sensing*, 32(4):786–799, July 1994.

[Raney, 1983] R. K. Raney. Transfer functions for partially coherent SAR systems. *IEEE Transactions on Aerospace and Electronic Systems*, AES-19(5):740–750, September 1983.

[Reed, 1962] I. S. Reed. On a moment theorem for complex gaussian processes. *IEEE Transactions on Information Theory*, 8(3):194–195, April 1962.

[Reeves and Rowe, 2003] C. R. Reeves and J. E. Rowe. *Genetic algorithms : principles and perspectives : a guide to GA theory*. Kluwer Academic Publishers, Boston, 2003.

[Rihaczek, 1969] A. W. Rihaczek. *Principles of High-Resolution Radar*. McGraw-Hill, 1969.

[Rodriguez and Martin, 1992] E. Rodriguez and J. M. Martin. Theory and design of interferometric synthetic aperure radars. *Proceedings IEE, Part F*, 139(2):147–159, April 1992.

[Runge and Bamler, 1992] H. Runge and R. Bamler. A novel high precision SAR focussing algorithm based on chirp scaling. *IEEE Transactions on Geoscience and Remote Sensing*, 1:372–375, 1992.

[Shippey et al., 1998] G. A. Shippey, P. Ulriksen, and Q. Liu. Quasi-narrowband processing of wideband sonar echoes. In A. Alippi and G. B. Canelli, editors, *Proceedings of the 4th European Conference on Underwater Acoustics*, pages 63–68, 1998.

[Soumekh, 1994] M. Soumekh. *Fourier Array Imaging*. Prentice Hall, Englewood Cliffs, NJ, 1994.

[Soumekh, 1999] M. Soumekh. *Synthetic Aperture Radar Signal Processing with MAT-LAB Algorithms*. John Wiley and sons, 1999.

[Stolt, 1978] R. H. Stolt. Migration by Fourier transform. *Geophysics*, 43(1):23–48, February 1978.

[Sutton et al., 2000] T. J. Sutton, S. A. Chapman, and H. D. Griffiths. Robustness and effectiveness of autofocus algorithms applied to diverse seabed environments. In M. E. Zakharia, editor, *Proceedings of the fifth European Conference on Underwater Acoustics ECUA 2000*, volume 1, pages 407–412. EUCA, July 2000.

[Touzi *et al.*, 1999] R. Touzi, J. Bruniquel A. Lopes, and P. Vachon. Coherence estimation for SAR imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 37(1):135–149, January 1999.

[Tur *et al.*, 1982] M. Tur, K. C. Chin, and J. W. Goodman. When is speckle noise multiplicative? *Applied Optics*, 21(7):1157–1159, April 1982.

[Ulander *et al.*, 2001] L. Ulander, H. Hellsten, and G. Stenstrom. Syntethic-aperture radar processing using fast factoried back projection. submitted to: IEEE Transactions on Aerospace and Electronic Systems, 8 December 2000, 2001.

[Vachon and Raney, 1989] P. W. Vachon and R. W. Raney. Estimation of the SAR system transfer function through processor defocus. *IEEE Transactions on Geoscience and Remote Sensing*, 27(6):702–708, November 1989.

[Wahl *et al.*, 1994a] D. E. Wahl, P. H. Eichel, D. C. Ghiglia, and C. V. Jakowatz, Jnr. Phase gradient autofocus – a robust tool for high resolution SAR phase correction. *IEEE Transactions on Aerospace and Electronic Systems*, 30(3):827–835, July 1994.

[Wahl *et al.*, 1994b] D. E. Wahl, C. V. Jakowatz, Jnr., and P. A. Thompson. New approach to strip-map sar autofocus. In *Sixth IEEE Digital Signal Processing Workshop*, pages 53–56. IEEE, October 1994.

[Warner *et al.*, 2000] D. W. Warner, D. C. Ghiglia, A. FitzGerrell, and J. Beaver. Two-dimensional phase gradient autofocus. *SPIE-Image Reconstruction form Incomplete Data*, 4123:162–173, 2000.

[Wehner, 1987] D. R. Wehner. *High Resolution Radar*. Artech House, Norwood, MA, 1987.

[White, 1984] H. White. *Asymptotic Theory for Econometricians*. Academic Press, 1984.

[Xi *et al.*, 1999] L. Xi, L. Guosui, and J. Ni. Autofocusing of ISAR images based on entropy minimization. *IEEE Transactions on Aerospace and Electronic Systems*, 35(4):1240–1252, October 1999.

[Ye *et al.*, 1999] W. Ye, T. S. Yeo, and Z. Bao. Weighted least-squares estimation of phase errors for SAR/ISAR autofocus. *IEEE Transactions on Geoscience and Remote Sensing*, 37(5):2487–2494, September 1999.

[Zomaya, 2001] A. Y. Zomaya. Natural and simulated annealing. *Computing in Science and Engineering*, 3(6):97–99, November/December 2001.