Signal Processing for Synthetic Aperture Sonar Image Enhancement

Hayden J. Callow

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.

April 2003

Abstract

This thesis contains a description of SAS processing algorithms, offering improvements in Fourier-based reconstruction, motion-compensation, and autofocus.

Fourier-based image reconstruction is reviewed and improvements shown as the result of improved system modelling. A number of new algorithms based on the wavenumber algorithm for correcting second order effects are proposed. In addition, a new framework for describing multiple-receiver reconstruction in terms of the bistatic geometry is presented and is a useful aid to understanding.

Motion-compensation techniques for allowing Fourier-based reconstruction in wide-beam geometries suffering large-motion errors are discussed. A motion-compensation algorithm exploiting multiple receiver geometries is suggested and shown to provide substantial improvement in image quality. New motion compensation techniques for yaw correction using the wavenumber algorithm are discussed.

A common framework for describing phase estimation is presented and techniques from a number of fields are reviewed within this framework. In addition a new proof is provided outlining the relationship between eigenvector-based autofocus phase estimation kernels and the phase-closure techniques used astronomical imaging. Micronavigation techniques are reviewed and extensions to the shear average single-receiver micronavigation technique result in a 3–4 fold performance improvement when operating on high-contrast images.

The stripmap phase gradient autofocus (SPGA) algorithm is developed and extends spotlight SAR PGA to the wide-beam, wide-band stripmap geometries common in SAS imaging. SPGA supersedes traditional PGA-based stripmap autofocus algorithms such as mPGA and PCA—the relationships between SPGA and these algorithms is discussed. SPGA's operation is verified on simulated and field-collected data where it provides significant image improvement. SPGA with phase-curvature based estimation is shown and found to perform poorly compared with phase-gradient techniques. The operation of SPGA on data collected from Sydney Harbour is shown with SPGA able to improve resolution to near the diffraction-limit.

Additional analysis of practical stripmap autofocus operation in presence of under-sampling and space-invariant blurring is presented with significant comment regarding

the difficulties inherent in autofocusing field-collected data. Field-collected data from trials in Sydney Harbour is presented along with associated autofocus results from a number of algorithms.

Acknowledgements

First and foremost I acknowledge the support and guidance of my supervisors Dr. Michael Hayes and Professor Peter Gough; they have provided much needed perspective at times—"Sure it works on simulated data... what about *real* data?". The broad wealth of knowledge these men possess and their skills in intuitive explanation have helped me a great deal. Mike, Peter, it has been a real pleasure working with you.

Many thanks to the other SAS lads: Philip Barclay, Steve Fortune, Edward Pilbrow, Alan Hunter, and Lóic Sibeud for giving me someone to bounce ideas off and to complain to while debugging simulations. Steve, Al and Phil deserve special mention for wading through this and other documents on proof reading duties—let me know when its my turn guys. Thanks also to Mike Cusdin, Peter Lambert and Phil for the work on the KiwiSAS-II hardware and to Dave Wilkinson (a past member of the Acoustics Research Group) for all the help back in the early days (and providing us light relief by occasionally logging onto the linux server with a three-finger-salute). Working in R9 has been memorable; thanks to all those past and present who have made it so—I will particularly miss the cricket matches and the evening-shift "phat beats".

A big thank-you to my family (my parents Clive and Doreen, brothers Kyle and Logan, sister Lauren, and fiancé Tanya) for their love and support—I would not be writing this today without you! Thanks too to all of those who decided to take pity on Tanya and I and invite us to dinner at different times. Paul Ryan and the Ellesmere Concert Band get a mention too for giving me necessary breaks from thinking about thesis work.

Thanks to the Defence, Science and Technology Organisation, Australia for their assistance with data collection in Sydney Harbour and also to the Defence Technology Organisation, New Zealand for their assistance with data collection in Auckland Harbour. Many thanks to the Norwegian Defence Research Establishment (FFI) for allowing me to visit them at Kjeller and to Dr. Roy Hansen for describing their implementation of DPCA/RPC and SAS plans providing me with another perspective on SAS processing. Thanks also to Kongsberg-Simrad for letting me tag along as a "tourist" at the sea-trials of the HUGIN-3000 AUV. For the financial support I received I would like to thank the University of Canterbury for granting me a University

of Canterbury Doctoral Scholarship (1999), and the Foundation for Research, Science, and Technology for a Bright Futures Scholarship (2000-2002).

Last but not least, thanks to my thesis examiners, Prof. Hugh Griffiths, Dr. Reinhard Holtkamp, and Dr. Philip Bones; it is a lot a work to examine a thesis and I am grateful they gave their time to do so. Mike and Peter get mention here too—their in-depth examination, proofing, and the big stick used to repeatedly beat this thesis (the almighty red pen) have made your job easier.

> My pen is at the bottom of a page,
> Which, being finished, here the story ends;
> 'Tis to be wished it had been sooner done,
> But stories somehow lengthen when begun.
> — Byron

Preface

This thesis is written in eleven chapters. Chapters 2–6 describe the deterministic aspects of SAS imaging: governing SAS principles, reconstruction and compensation for known path-deviations. The later chapters (7–10) detail aspects of micronavigation and autofocus needed to overcome random path-deviations and medium fluctuations. Chapter 7 covers fundamentals of autofocus and chapter 8 summarises a number of spotlight autofocus techniques. Techniques for micronavigation and autofocus of stripmap systems are presented in chapters 9 and 10. The last chapter summarises the work in the thesis and presents recommendations for future research.

Papers prepared during the course of this thesis are listed in order of presentation:

P. T. Gough, M. P. Hayes, and H. J. Callow. Strip-map path correction using phase matching autofocus. In *Proceedings of the fifth European Conference on Underwater Acoustics (ECUA) 2000*, pages 412–418, July 2000.

H. J. Callow, P. T. Gough, and M. P. Hayes. Not quite 20,000 leagues under the sea, sonar autofocus in a shallow water reverberant environment. In *Image and Vision Computing New Zealand 2000*, pages 304–309. IVCNZ, November 2000.

H. J. Callow, M. P. Hayes, and P. T. Gough. Autofocus of multi-band, shallow-water, synthetic aperture sonar imagery using shear-averaging. In *Proceedings of International Geoscience And Remote Sensing Symposium (IGARSS) 2001*, pages 1601–1603. IEEE, July 2001a.

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.

H. J. Callow, M. P. Hayes, and P. T. Gough. Noncoherent autofocus of single receiver, broad-band synthetic aperture sonar imagery. In *OCEANS 2001*, volume 1, pages 157–162. IEEE, November 2001b.

M. P. Hayes, P. J. Barclay, P. T. Gough, and H. J. Callow. Test results of a bathymetric synthetic aperture sonar. In *OCEANS 2001*, volume 3, pages 1682–1687. IEEE, November 2001.

H. J. Callow, M. P. Hayes, and P. T. Gough. Advanced wavenumber domain processing for reconstruction of broad-beam multiple-receiver synthetic aperture imagery. In *Image and Vision Computing New Zealand 2001*, pages 51–56. IVCNZ, November 2001c.

H. J. Callow, M. P. Hayes, and P. T. Gough. A wavefront sensing method for synthetic aperture sonar autofocus. In *Proceedings of the sixth European Conference on Underwater Acoustics (ECUA) 2002*, pages 231–234, Gdansk, Poland, 2002a.

H. J. Callow, M. P. Hayes, and P. T. Gough. Wavenumber domain reconstruction of SAR/SAS imagery using single transmitter and multiple-receiver geometry. *Electronics Letters*, 38(7):336–337, March 2002b.

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.

H. J. Callow, M. P. Hayes, and P. T. Gough. Motion compensation improvement for wide-beam, multiple-receiver SAS systems. Accepted for publication in the IEEE Journal of Oceanic Engineering, 2003a.

H. J. Callow, M. P. Hayes, and P. T. Gough. Stripmap phase gradient autofocus. In *Oceans 2003*, San Diego, USA, September 2003b. IEEE/MTS. CDROM.

H. J. Callow, M. P. Hayes, and P. T. Gough. Autofocus of stripmap SAS data using the range variant SPGA algorithm. In *Oceans 2003*, San Diego, USA, September 2003c. IEEE/MTS. CDROM.

Nomenclature

This section defines the notation used in the thesis—a variant on the notation of Gough and Hawkins [1997] and Hawkins [1996]. In particular, the Fourier dual of a variable uses a subscript notation—i.e., k_x is the radial Fourier frequency (wavenumber) corresponding to the along-track variable x ($k_x = 2\pi f_x$). The 2-D Fourier transform for example is defined as

$$E(\omega, k_u) = \iint_{-\infty}^{\infty} e(t, u) \exp\left(-j(\omega t + k_u u)\right) \mathrm{d}t\, \mathrm{d}u. \tag{1}$$

The double letter notation of Gough and Hawkins [1997] does not scale well to 3+ dimensions and has been replaced in favour of a single letter notation. Capitalisation of the single letter denotes a *temporal* Fourier relationship. Note that *spatial* Fourier transforms are not capitalised. An example of this notation is given below, (note $\leftrightarrow_t$ indicates a Fourier transform over the time variable t)

$$e(t, u) \leftrightarrow_t E(\omega, u) \quad e(t, u) \leftrightarrow_u e(t, k_u) \quad \text{(temporal, spatial)}$$

whereas

$$f(x, y) \leftrightarrow_x f(k_x, y), \quad f(x, y) \leftrightarrow_y f(x, k_y) \quad \text{(spatial, spatial)}$$

This notation is also consistent with that commonly used to describe time sequences such as [Haykin 1994],

$$p(t) \leftrightarrow_t P(\omega).$$

Definitions for the remainder of the symbols in the thesis are as follows:

$\mathbf{Arg}\,\{\}$	Angle operator
θ_{3dB}	-3 dB beamwidth
$\theta_{\text{null-to-null}}$	Null-to-null beamwidth
B_c	Signal bandwidth
$\mathbf{Conj}\,\{\}$	Complex conjugate
$\odot_t$	Convolution along time coordinate
$\star_t$	Correlation along time coordinate

c	Wave-speed in the medium
CRLB	Cramér-Rao lower bound
D	Effective combined transmit/receive element along-track length
D_{r}	Receive element along-track length
D_{t}	Transmit element along-track length
Δu	Along-track sample spacing (echo data)
Δy	Along-track sample spacing (image)
Δt	Time sample spacing
Δx	Across-track sample spacing
ϵ	Phase centre approximation error
$e(t, u)$	Raw echo data
$e_{\mathrm{tvg}}(t, u)$	Raw echo data after TVG applied
$\mathcal{F}_t\{\}$	Fourier transform along time axis
$\widetilde{f}(x, y)$	Blurred image
$\mathcal{F}_\omega^{-1}\{\}$	Inverse Fourier transform along frequency axis
$f(x, y)$	Image
$\widehat{f}(x, y)$	Image estimate
$f_n(x, y)$	Small image section
h	Height of the platform
$\mathbf{Im}\{\}$	Imaginary part operator
i_{t}	Transducer aperture function
k	Wavenumber in the medium—is equivalent to ω/c
k_0	Carrier wavenumber
k_b	Baseband wavenumber
λ	Wavelength
∇^2	Laplacian
m	Along-track position of co-located transmitter/receiver pair
η	Elevation angle, also used for complex phasor
$\psi(x, y, z)$	Field measured at (x, y, z)
$p_b(t)$	Baseband transmitted signal
u	Along-track position of the platform; usually used to describe the varying position during recording
$\mathbf{Re}\{\}$	Real part operator
τ_{rep}	Repetition period
δ_r	Range resolution
δ_t	Time resolution

δ_x	Across-track resolution
δ_y	Along-track resolution
ρ_{in}	Range from target to receiver
ρ_{out}	Range from transmitter to target
σ	Along-track position of receiver relative to platform
$s(t,u)$	Pulse compressed data
$s_{\mathrm{rmc}}(t,u)$	Range-migration corrected pulse compressed data
$s_{\mathrm{tvg}}(t,u)$	Pulse compressed data after TVG applied
$\widetilde{s}(t,u)$	Blurred pulse compressed data
τ_c	Signal length
t	Time coordinate
v_s	Along-track velocity of the platform
ω	Temporal frequency (radians / s)
ω_0	Carrier frequency (radians / s)
ω_{max}	Maximum temporal frequency (radians / s)
$\widetilde{\Phi}_{\mathrm{shear},\gamma}$	Tow-fish crab
$\dot{X}(u)$	Sway differential
X_p	Sway (with respect to sonar pulse p)
$X(u)$	Sway
$\Phi_\sigma(u)$	Pulse-by-pulse array yaw
Φ	Global yaw
x	Across-track position in global coordinates
x_s	Slant-range
y	Along-track position in global coordinates
y_{tx}	Along-track position of transmitter
y_{rx}	Along-track position of hydrophone
C_s	Curvature factor
K_{src}	Geometry induced chirp rate
K_s	Measured chirp rate
R_s	Range-migration factor
z	Height in global coordinates
H	Number of receivers
M	Order of eigenvector estimator
N	Number of time samples
P	Number of pulses
$\mathcal{O}\{\}$	Big 'O' notation, computational complexity

2-D PCA	2-D Phase curvature autofocus
DPC	Displaced phase centre
DPCA	Displaced phase centre antenna
FFBP	Fast factorised back projection
FLOS-PGA	Fractional lower order statistics phase gradient autofocus
MD	Map drift autofocus
MOCOMP	Motion compensation
PCA	Phase curvature autofocus
PDA	Phase difference autofocus
PGA	Phase gradient autofocus
PMA	Phase matching autofocus
PPP	Prominent point autofocus
QPGA	Weighted phase gradient autofocus
RPC	Redundant phase centres
SPGA	Stripmap phase gradient autofocus
WPGA	Weighted phase gradient autofocus

Contents

Introduction

1.1 Underwater imaging

The earth is an aquatic planet and, with as much as 80% of the surface covered in water, there is a strong interest in knowing what lies below. Optical wavelengths penetrate poorly into sea-water with common visibilities ranging from 10s of metres down to the order of a metre. This makes the charting and imaging of underwater scenes using lights and cameras difficult. In contrast, acoustic signals and sounds travel well in the ocean, particularly low frequency signals ($< 1\,\mathrm{kHz}$). The result is that acoustic imaging is often the only means of large scale underwater imaging.

1.1.1 Echo detection

Echo detection—the location and detection of objects using sound—is a technique likely to have been in use as long as humans have existed. With the advent of sea-faring craft, echo detection became a way of avoiding running vessels aground. Some evidence exists of Phoenician fishermen (circa 500 B.C) using the echos of ringing bells and the like to detect nearby headlands [Kaharl 2003].

Echo detection, and its extension, echolocation, became popular techniques after the sinking of the Titanic in 1911, with the first patent filed within two weeks of the sinking [Urick 1975]. The echolocation patent was granted for location of icebergs using sound in air. The inventor, J. F. Richardson, was soon awarded another patent for the application of the same technique underwater where the propagation of sound is much better. Underwater echolocation became more important with the threat of submarine warfare in World War I and a range of techniques using steerable arrays of hydrophones came into use. Gradually, these techniques became the beginning of what is now the sonar imaging field.

1.1.2 Sonar

The idea and application of sonar techniques occurred earlier than many realise with Leonardo da Vinci writing late in the 15th century [Burdic 1984]:

> *If you cause your ship to stop, and place the head of a long tube in the*
> *water and place the other extremity to your ear, you will hear ships at a*
> *great distance from you.*

The technique da Vinci outlines is a crude form of passive sonar; more sophisticated passive systems of similar construction were used until late in World War I when electronic devices started to take over [Urick 1975].

The term SONAR came into use late in World War II to describe the *SO*und, *NA*vigation and *R*anging techniques used in anti-submarine warfare at the time and to parallel the then newly coined term RADAR [Urick 1975]. Sonar systems were employed to find the bearing and range of submarine targets often using a mechanically steered hydrophone array and echo-location. Post World War II, sonar techniques found application in civilian activities such as sea-bed imaging, depth sounding, and fish-echolocation [Tucker 1966].

Today, side-scan sonar provides part of the basis for both civilian and military sea-bed imaging. This thesis describes synthetic aperture sonar (SAS) imaging in a side-scanning mode—other sonar imaging geometries also benefit from aperture synthesis although not as greatly as side-scan mode.

1.2   Side-scan imaging geometry

Side-scan (side-looking) sonar systems, both conventional and synthetic aperture, send sound pulses perpendicular to the direction of tow. The returning echo signals are formed into a narrow imaging beam (either by using a large real-aperture in conventional systems or via synthetic aperture processing). The image strips that the beam illuminates are mapped onto a continuous plot as the sonar moves forward—i.e., built into a *stripmap* image or *sonograph*.

Figure 1.1 illustrates a typical side-scan sonar geometry. The imaging platform travels a rectilinear path in the *along-track* or *azimuth* direction u at a fixed height off the mean seafloor H. Sound pulses are transmitted perpendicular to the direction of travel, along the cross-track direction, x, and propagate radially outward. The echos for a particular pulse are recorded as a function of time-delay since transmission t, which is proportional to the two-way range of the scattering target $t = r_{\text{twoway}}/c$, where c is the speed of propagation in the medium.

Images are formed (reconstructed via synthetic aperture processing) on a grid in global (x, y, z) coordinates or possibly in a slant-range coordinate system (x_s, y) (see Section 3.1).

Figure 1.1   Illustration of the typical imaging geometry of side-scanning sonar systems. The sonar image is built up by mosaicking each pulse onto a *stripmap*.

1.2.1   Conventional side-scan sonar

Conventional side-scan sonars operate by transmitting a narrow beam of sound energy, then plotting the echo returns onto a recording chart one pulse at a time. The along-track resolution in the final image is proportional to beam-width—the narrower the beam the better the resolution. Narrowing the beam is possible by either increasing the frequency of the acoustic pulses or increasing the length of the aperture. Thus conventional side-scan sonars typically use high-frequency signals for high resolution imaging.

The disadvantage of using high frequency sonars is that attenuation of the signal is greater at higher frequencies. High-resolution conventional side-scan systems only operate at short range. Another difficulty occurs when conventional systems are used for detection. With sonar beams of 0.1–1°, minor path variations make it easy for the sonar to completely miss imaging a target. This problem cannot be fixed in post-processing—if no energy hits a target it can not give reflections. This has lead to the separate development of low-frequency detection and high-frequency classification sonars.

Regardless of the disadvantages, real-time operation and inherent simplicity make conventional side-scan sonar systems an invaluable mapping tool.

1.2.2   Synthetic aperture sonar (SAS)

Synthetic aperture sonar operates by sending wide-beam acoustic pulses and combining the echoes from many pulses to obtain high-resolution images. The benefit of this approach is that low-frequency signals can generate high-resolution imagery without the need for excessively long arrays. The high along-track resolution is also range-independent—something not possible in conventional systems. However, the disadvantage is that the echoes must retain phase coherency for the length of the synthetic aperture. Without echo coherency the SAS image is severely degraded. The prevention of image degradation requires strict motion and sampling requirements.

The other major disadvantages SAS brings are requirements for heavy computation and increased system complexity. Computing advances and improved reconstruction techniques have only recently made real-time (and thus commercial) SAS feasible.

Spotlight mode

Spotlight-mode imagery is common in synthetic aperture radar (SAR—the radar equivalent of SAS). A spotlight system steers a narrow beam (either mechanically or electronically) onto a small patch of the target area as it passes—see Figure 1.2(b). This allows a large reduction in the data storage and transmission requirements (important in satellite SAR systems).

Spotlight systems are a special case of the stripmap geometry that allow approximations simplifying the image reconstruction process. No published SAS system uses spotlight mode, although with improving navigation accuracy the more relaxed along-track sampling constraints of spotlight mode may be useful. Spotlight mode imaging is discussed in the autofocus chapters of this thesis. Detailed descriptions of spotlight SAR imaging are given by Carrera et al. [1995] and by Jakowatz et al. [1996].

Stripmap mode

Stripmapping is the conventional imaging mode for SAS systems. In stripmap mode the sonar beam always points in the same direction during imaging—see Figure 1.2(a). Stripmap imagery does not allow many simplifying assumptions, making it more difficult than spotlight imagery to reconstruct and autofocus. The majority of the discussion in this thesis relates to stripmap SAS/SAR systems.

1.2.3   Navigation errors

SAS systems require that the recorded echos retain phase coherency for the length of the synthetic aperture. Coherency requires phase accuracies of better than $\lambda/8$ and positioning accuracy of at least $\lambda/16$. Without echo coherency, the resulting SAS

(a)   (b)

Figure 1.2 Stripmap and Spotlight imaging modes. (a) Stripmap mode, the beam points the same direction during imaging. (b) Spotlight mode, the beam is continuously steered onto a small scene patch during imaging. [Hawkins 1996]

Figure 1.3 Sonar system motion errors [Johnson 1992].

images suffer severe degradation[1]. Even on tightly constrained, rail-based systems this type of accuracy is difficult to achieve.

Of the six possible motions shown in Figure 1.3, across-track motion (sway) is the most important to constrain for side-scan SAS operation [Hayes and Gough 1992]. With low grazing angles to the scene, towfish sway[2] accounts for the majority of echo

---

[1]Realistic systems require millimetre accuracy in position over aperture lengths of 10's of metres (due to typical wavelengths of 10 mm–50 mm).

[2]In multiple-receiver systems, yaw motions are similarly important.

de-correlation [Cutrona 1975]. Although sway has the most effect on SAS imagery, it is the combined effect of all six degrees of freedom on slant-range motion that must be controlled. Designing a free-towed or autonomous system able to control all those degrees of freedom to within SAS motion tolerances in general conditions is impractical.

Despite the tight tolerances on motion, diffraction-limited imagery has been shown with free-towed SAS systems, albeit in benign operating conditions [Hawkins 1996; Hayes and Gough 1999]. In conditions where navigation errors would otherwise cause image blurring, inertial navigation system (INS) motion measurements combined with data-driven micronavigation/autofocus offer the potential for producing blur-free imagery. This thesis presents both methods for using INS measurements to correct image-blurring (when using efficient Fourier-based reconstruction techniques), and techniques for data-driven image compensation (autofocus).

1.3   Assumed knowledge

It is expected that the reader has an advanced understanding of the principles of synthetic aperture imaging and a good working knowledge of digital signal processing techniques. In particular, a good knowledge of discrete interpolation techniques is important. Oppenheim and Schafer [1999] and Jakowatz et al. [1996] give detailed descriptions of digital signal processing and discrete interpolation techniques. Fourier transformation is used extensively in the thesis and the reader should have an advanced understanding of Fourier domain concepts. Bracewell [1986] provides a good description of advanced Fourier techniques although a summary of important Fourier properties is provided in Appendix G.

Specialised knowledge of sonar systems design is not required although an understanding of the physical and mathematical constraints are beneficial. The reader should be familiar with aperture synthesis techniques and general imaging procedures. In particular, a basic understanding of SAR imaging ([Carrera et al. 1995; Soumekh 1994, 1999]) and the differences between SAS and SAR (see [Bonifant 1999; Hawkins 1996; Hayes and Gough 1992]) is useful. Wahl et al. [1994a], Jakowatz and Wahl [1993] and Wahl et al. [1994b] provide a good background to PGA and PCA based autofocus for the interested reader.

1.4   Thesis contributions

This thesis presents a number of contributions to the SAS and SAR fields.

Improved SAS system models are derived based on the Helmholtz wave equation and Weyl's identity. These models provide wide-beam, wide-band wavenumber-domain amplitude modelling and allow improvements to reconstruction techniques. A summary

of common reconstruction algorithms in consistent notation and their computational efficiencies is presented in light of the improved modelling. New wavenumber-domain derivations are shown that model Doppler and position shifting effects for continuously moving sonars. An improved reconstruction algorithm making use of this modelling is proposed in Chapter 5.

A number of improvements to the wavenumber algorithm and insights into its use are presented. Phase-centre compensation is improved with a new method for ultra-wide-beam systems proposed. The multiple-receiver reconstruction problem is formulated in terms of a bistatic geometry, providing new insight into the problem. A new multiple-receiver reconstruction algorithm based on the bistatic formulation is proposed and test results presented.

This thesis described wide-beam sway compensation to be used in conjunction with Fourier-based reconstruction algorithms. A summary of single-receiver algorithms and a new multiple-receiver algorithm is presented. Simulation results from the new algorithm are shown to provide image improvement. Another new algorithm for the compensation of sway movement occurring during the pulse and additional wide-beam yaw compensation techniques provide further motion compensation improvement.

Chapter 7 provides a comprehensive summary of the techniques and ideas behind autofocus. New motion constraints are derived and show that sway gradient errors provide closer correspondence to image blurring than absolute sway errors. Time-delay and phase estimation techniques from a number of fields are summarised in a common framework. The Cramér-Rao lower bounds of each of the techniques are provided to give insight into algorithm accuracy. PGA's eigenvector phase estimator is shown to implement phase-closure for order $M \geq 3$ and also shown to be equivalent to the bispectrum and trispectrum techniques used in astronomical imaging.

Chapter 8 provide a summary of spotlight autofocus and various PGA algorithms is presented in consistent notation. The parametric PDA method is shown to use a phase estimator closely related to PGA's ML estimator.

Chapter 9 provides a summary of micronavigation techniques and the limitations of image correlation—it is effectively limited by the same constraints as traditional phase gradient autofocus. This thesis gives a description of the problems facing shear average autofocus on field-collected data. Extensions are suggested to alleviate those problems. The autofocus improvement the extensions provide is demonstrated both on simulated data and data collected in the field at Sydney Harbour.

Chapter 10 provides an extensive investigation into stripmap autofocus using phase gradient techniques. A new wide-band, wide-beam stripmap blurring model is derived allowing improvements to autofocus algorithms. The blurring model is extended for the case of large sways and the implications of narrow-band, narrow-beam approximations are discussed. A consistent and in-depth summary of traditional autofocus techniques,

the blurring models they employ, and their (stripmap) autofocus limitations is also given in Chapter 10.

This thesis presents the new SPGA algorithm. The algorithm is based on the wide-beam, wide-band blurring model and the new insight that Doppler spectrum shifting can be used to aid autofocus. In-depth discussion of SPGA's technical aspects and implementation issues provide the necessary information for successful operation. Extensions to the algorithm for operating where gross sway is present are presented. SPGA supersedes traditional algorithms and the implementation of traditional algorithms using the SPGA framework is discussed. The autofocus problems caused by aperture undersampling and space-invariant blurring are also presented leading to discussions on the mitigation of those effects. Results from in-depth quantitative testing of SPGA on simulated data are shown and provide a guide to the selection of various autofocus parameters.

Real data testing on field collected data from Sydney Harbour is demonstrated with SPGA offering significant image improvement. Analysis of the failure of 2-D PCA (SPGA with a phase curvature kernel) on the same data is presented, demonstrating the benefit of phase-gradient processing. Realistic quantitative autofocus testing is difficult and a summary of testing difficulties shows the areas where improvement is necessary.

1.5   Thesis outline

Chapter 2 covers the fundamental principles of SAS imaging including a section on complex-valued data collection and sampling. Acoustic propagation is reviewed leading to the methods used for SAS system modelling. A discussion on array theory and the resolution constraints on a typical side-scan SAS system are also presented.

Chapter 3 covers the system models chosen for SAS systems and provides spatial and wavenumber domain model derivations. New wavenumber domain derivations for multiple-receiver SAS systems and second-order effects models are presented.

Chapter 4 reviews the common SAS reconstruction algorithms and presents new derivations with amplitude terms to appropriately compensate spreading loss.

Chapter 5 goes into the detail of wavenumber algorithm based reconstruction and implementation issues, providing a discussion on the trade-offs between reconstruction accuracy and efficiency. In addition, a new bistatic path derivation for multiple-receiver reconstruction and improved phase-centre approximation compensation are presented. New inversions for moving sonar compensation and

temporal Doppler effects are provided. An additional 3-D wavenumber reconstruction is presented to aid the understanding of interferometric processing.

Chapter 6 reviews current techniques for motion compensation and suggests improvements for multiple-receiver wide-beam sonar systems. The approximations made in current motion compensation are considered and more accurate descriptions and corrections of yaw and sway effects are provided.

Chapter 7 provides introductory material for the following chapters and discusses the issues underlying SAS autofocus. The autofocus problem is presented and the differences between stripmap and spotlight autofocus considered. A summary of wide-band echo-correlation techniques with their theoretical accuracies is also given. Phase estimation techniques from a number of difference fields are discussed and compared with phase-only time-delay estimation. An overview of current SAS autofocus techniques is presented to introduce the following chapters.

Chapter 8 provides a detailed investigation into phase gradient based autofocus for stripmap systems. Traditional algorithms are summarised and related to a new technique based on the use of the wavenumber transform. Autofocus results for simulated and field-collected data are presented. Discussions on the difficulties of stripmap autofocus, the failure of traditional and wavenumber transform based blurring models, simulation deficiencies with regard to autofocus, and aperture undersampling and its effect on autofocus are included.

Chapter 9 provides a review of the echo-correlation based autofocus algorithms. The redundant phase centre (RPC) algorithm is analysed and considered relative to image correlation and to the single-receiver shear average algorithm. Improvements to the shear average algorithm for operation in scenes with high contrast and for bulk motion estimation are presented. Real-world autofocus results on both simulated and field collected data are demonstrated.

Chapter 10 presents the novel wide-beam, wide-band stripmap phase gradient autofocus (SPGA) algorithm. The derivation of an appropriate stripmap blurring model and the failure of traditional blurring models are presented. Traditional SAS autofocus algorithms are investigated in relation to autofocus implementation and related to the new SPGA algorithm. A number of the major steps of the SPGA algorithm are described and those common to spotlight autofocus discussed from a stripmap SAS viewpoint. Aperture undersampling and the problems caused by space-invariant blurring are also discussed. Autofocus testing on both simulated and field-collected data is presented and conclusions regarding SPGA's operational parameters drawn.

Chapter 11 presents the conclusions drawn from this work and provides recommendations for future research.

Synthetic Aperture Sonar (SAS) fundamentals

This chapter describes the notation and fundamental theory needed to describe SAS imaging. Complex and band-pass signals are covered in Section 2.1 and the sampling constraints on band-pass signals in Section 2.2. A description of pulse-compression techniques in Section 2.3 is followed by acoustic wave propagation theory required for accurate SAS system modelling. The chapter finishes with an overview of sidescan SAS resolution and sampling constraints.

2.1 Complex signals

The waveforms used in echo imaging systems may be expressed in the form

$$p(t) = a(t) \cos(\Phi(t)), \qquad (2.1)$$

where $a(t)$ is a time-varying, amplitude-only function (often called the signal envelope) and $\Phi(t)$ is a phase modulation term. Usually, the signal, $p(t)$, is a band-pass signal and (2.1) may be written as

$$p(t) = a(t) \cos(\phi(t) + \omega_0 t), \qquad (2.2)$$

where the phase function $\Phi(t) = \phi(t) + \omega_0 t$ is composed of a time-varying phase function $\phi(t)$ and an arbitrary carrier frequency ω_0. Notation is often simplified for band-pass signals by rewriting (2.2) in terms of a low-pass complex envelope signal, $p_b(t)$, and exponential carrier

$$p(t) = \mathbf{Re}\{p_b(t) \exp(j\omega_0 t)\}, \qquad (2.3)$$

where $p_+(t) = p_b(t) \exp(j\omega_0 t)$ is the pre-envelope and the $\mathbf{Re}\{\}$ operator takes the real part of the signal [Haykin 1994]. The baseband complex envelope, $p_b(t)$, may be obtained from the pre-envelope, $p_+(t)$, by demodulating by the carrier frequency[1]

$$p_b(t) = p_+(t) \exp(-j\omega_0 t). \qquad (2.4)$$

---

[1]In practice, the baseband signal is usually created using quadrature (I-Q) sampling [Haykin 1994].

The baseband signal $p_b(t) = a(t) \exp(j\phi(t))$ is sufficient to describe the band-pass signal, completely independent of the carrier frequency [Haykin 1994].

The use of baseband signals makes for simplified computer processing and reduces the amount of storage space required (see Section 2.2).

2.2 Sampling of complex baseband signals

If the sonar signal is a bandpass signal, as described above in Section 2.1, traditional sampling at the Nyquist rate [Haykin 1994] often results in more samples being taken than required to fully characterise the signal. The low-pass baseband signal, $p_b(t)$, contains all of the information of the real signal, $p(t)$, thus only the band-width of $p_b(t)$ need be Nyquist sampled[2]. For typical signals in pulse-echo imaging systems, sampling the baseband signal can result in a large reduction in data storage requirements. Nyquist sampling results in a sample spacing of

$$\Delta t = \frac{4\pi}{\omega_{\max} + \omega_0},\qquad(2.5)$$

where $\omega_{\max}$ is the maximum frequency component contained in $p_b(t)$ and ω_0 is the carrier frequency. Complex-baseband sampling by contrast results in a sample spacing of [Haykin 1994]

$$\Delta t = \frac{2\pi}{\omega_{\max}},\qquad(2.6)$$

where each sample is now complex. For bandpass signals the storage saving is approximately equal to the bandwidth to carrier ratio (Q). When sampled baseband signals are used, high-Q sonar systems gain a significant saving in the storage space needed for the echo data.

2.3 Pulse-compression

The range of a sonar system, and the dynamic range of the final imagery are proportional to the transmitted signal's energy. High power impulse-like pulses of extremely short duration are unable to be used due to constraints on the maximum signal amplitude in water (owing to non-linearity and cavitation effects [Urick 1975])[3]. The range resolution of a waveform is proportional to its pulse duration (for impulse-like signals[4]) but the energy able to be transmitted when using impulse-like signals is severely constrained. This trades resolution against energy; high-energy, high-resolution imaging is

---

[2]Only the real and imaginary parts of $p_b(t)$ (the I-Q channels) need be sampled at $2\pi/\omega_{\max}$ where $\omega_{\max}$ is the maximum frequency component contained in $p_b(t)$.

[3]Radar has similar constraints on peak amplitude, these are constraints on the transmitter electronics rather than the medium.

[4]Shorter duration pulses result in higher resolution because they have a larger bandwidth (see Section 2.6.2).

not possible. Lowering the signal energy decreases the SNR and causes a corresponding decrease in system dynamic range.

The constraint on signal amplitude may be avoided by transmitting long-duration phase-modulated signals. The use of such signals allows more energy to be put into the water and gives an increase in system SNR proportional to the time-bandwidth product of the signal $(\tau_c B_c)$ [Cook and Bernfeld 1967; Curlander and McDonough 1996; Hawkins 1996; Hayes 1989]. Moreover, resolution (which is proportional to the signal bandwidth) is also improved. To improve the resolution of the system, the received echo signal is *pulse-compressed* [Hovanessian 1980, page 130]. This operation is usually performed as a *matched-filtering* with the transmitted signal and gives a final resolution that is proportional to the signal bandwidth. The classical matched-filtering used is the same as complex correlation methods outlined in Section 7.6.1 and in Haykin [1994] and Camp [1970].

In the notation used in this thesis, the pulse-compressed echo signal, $s(t, u)$, is calculated from the raw echo signal, $e(t, u)$ via correlation

$$s(t, u) = \int e(t', u) p^*(t' - t) \, \mathrm{d}t' \tag{2.7}$$

$$= e(t, u) \odot_t p^*(-t) \tag{2.8}$$

$$= e(t, u) \star_t p(t), \tag{2.9}$$

where $p(t)$ is the transmitted signal and $\odot_t$ and $\star_t$ denote convolution and correlation respectively in the time axis. The operation of (2.7) may also be represented in the temporal-frequency domain as

$$S(\omega, u) = E(\omega, u) P^*(\omega), \tag{2.10}$$

which is how the calculation is normally performed.

Pulse-compression, in conjunction with transmitted signals with large time-bandwidth products, allows for the use of high-resolution, high-energy signals in sonar imaging and avoids the constraint imposed by the cavitation limit of the medium [Hayes and Gough 1992].

2.4   Acoustic wave propagation

Sonar systems "paint" images using acoustic energy. The round-trip range and acoustic scattering strength of the sea-floor are combined to create images of the sea-floor. Due to the sonar's reliance on sound propagation, it is important that the effect is modelled properly. This section describes the linear wave equation that approximately models the sound pulses of most sidescan sonars. With many SAS inversion techniques relying

on wavenumber representations of the signal, acoustic propagation modelling in the wavenumber domain is also covered.

The propagation of small-signal acoustic waves in the ocean is approximately governed by the Helmholtz wave equation for a scalar field, ψ, in an unbounded homogeneous medium [Morse and Feshbach 1953][5],

$$\left(\nabla^2 + k^2\right)\psi(x, y, z) = 0, \tag{2.11}$$

where ∇^2 is the Laplacian operator,

$$\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}, \tag{2.12}$$

$k = \omega/c$, is the wavenumber for a given acoustic frequency ω, with c the field propagation speed in the medium. If an acoustic field $f(x, y, z)$ exists then (2.11) becomes the forced Helmholtz equation,

$$\left(\nabla^2 + k^2\right)\psi(x, y, z) = -f(x, y, z). \tag{2.13}$$

Taking the 3-D spatial Fourier transform of both sides the frequency response is represented as

$$\left(-k_x^2 - k_y^2 - k_z^2 + k^2\right)\psi(k_x, k_y, k_z) = -f(k_x, k_y, k_z), \tag{2.14}$$

$$\psi(k_x, k_y, k_z) = \frac{-f(k_x, k_y, k_z)}{\left(k^2 - k_x^2 - k_y^2 - k_z^2\right)}. \tag{2.15}$$

Substituting in an impulse (point) source at (x_0, y_0, z_0) as the forcing field into (2.15)

$$-f(k_x, k_y, k_z) = \exp\left(-jk_x x_0 - jk_y y_0 - jk_z z_0\right),$$

$$\psi(k_x, k_y, k_z) = \frac{\exp\left(-jk_x x_0 - jk_y y_0 - jk_z z_0\right)}{\left(k^2 - k_x^2 - k_y^2 - k_z^2\right)}. \tag{2.16}$$

and taking the inverse Fourier transform of (2.16), we get the spatial response,

$$\psi(x, y, z) = \frac{\exp\left(-jk\sqrt{(x - x_0)^2 + (y - y_0)^2 + (z - z_0)^2}\right)}{4\pi\sqrt{(x - x_0)^2 + (y - y_0)^2 + (z - z_0)^2}}. \tag{2.17}$$

The unconstrained (3-D) spatial response given by (2.17) is the result of convolving a

---

[5][Ziomek 1995] provides another derivation of (2.11) beginning with the basic physical properties of the medium.

point scatterer at (x_0, y_0, z_0) with the system *Green's function*,

$$g(x, y, z) = \frac{\exp\left(-jk\sqrt{x^2 + y^2 + z^2}\right)}{4\pi\sqrt{x^2 + y^2 + z^2}}. \qquad (2.18)$$

Superposition of the responses from many point-sources can be used to derive an expression for the measured field $\psi(x, y, z)$ for an arbitrary distribution of sources.

Assuming a model where the spreading of energy is contained in 2-D, similar derivations can be undertaken [Chew 1995], starting with the 2-D version of (2.15) and (2.18),

$$g(k_x, k_y) = \frac{-1}{k^2 - k_x^2 - k_y^2}. \qquad (2.19)$$

Inverse Fourier transforming (2.19) gives the 2-D spatial impulse response (Green's function) as,

$$g(x, y) = -\frac{j}{4} H_0^{(2)}(k\rho), \qquad (2.20)$$

where $\rho = \sqrt{x^2 + y^2}$ and $H_0^{(2)}$ is a Hankel function of the 2nd kind. When $\rho \gg \lambda = 2\pi/k$, i.e., the range is much greater than the wavelength, (2.20) may be approximated by an asymptotic expansion [Morse and Feshbach 1953],

$$g(x, y) \approx \frac{-1}{\sqrt{j8\pi k\rho}} \exp\left(-jk\rho\right). \qquad (2.21)$$

Using these models for the propagation of sound in water, the data received by the sonar may be synthesised and reconstruction algorithms derived. This modelling is covered in more detail in Chapter 3.

2.5 Array theory

This section develops the theory required to understand basic imaging with real aperture arrays—such as that used in conventional sidescan imaging sonars. The extension to synthetic aperture arrays follows with a summary of resolution and sampling constraints.

2.5.1 Transducers and beam-patterns

Sonar transducers are often rectangular and the beam-patterning problem is able to be separated into two 1-D problems. The following analysis may be extended to non-separable aperture functions in 2-D.

Given a 1-D aperture illumination function, $i_t(y)$, the field incident on a target at (x, y), for a transducer centred on $(0, 0)$, is given by (using the 2-D Green's func-

Figure 2.1 Ground-plane geometry for beam-pattern calculations.

tion (2.20))

$$\psi_i(x,y) = -j\frac{1}{4}\int i_t(u)H_0^{(2)}\left(k\sqrt{x^2+(y-u)^2}\right)\mathrm{d}u. \tag{2.22}$$

Taking a 1-D Fourier transform of (2.22) in y and employing Weyl's identity (Appendix B) gives

$$\psi_i(x,k_y) = -\int i_t(u)\frac{\exp\left(-j\,|x|\,\sqrt{k^2-k_y^2}-k_yu\right)}{\sqrt{k^2-k_y^2}}\,\mathrm{d}u. \tag{2.23}$$

Recognising the Fourier transform relation in (2.23) allows the incident field to be written

$$\psi_i(x,k_y) = -I_t(k_y)\frac{\exp\left(-j\,|x|\,\sqrt{k^2-k_y^2}\right)}{\sqrt{k^2-k_y^2}}, \tag{2.24}$$

where $I_t(k_y)$ is the Fourier transform of $i_t(y)$. Using the stationary phase method [Cook and Bernfeld 1967; Gough and Hawkins 1997; Soumekh 1994] to perform an Inverse Fourier transform over k_y, the spatial equivalent of (2.24) may be written

$$\psi_i(x,y) \approx \frac{-j}{4}I_t(k\sin\theta)\,H_0^{(2)}\left(k\sqrt{x^2+y^2}\right), \tag{2.25}$$

where $\theta = \sin^{-1}\left(y/\sqrt{x^2+y^2}\right)$. Equation (2.25) is a weighted version of (2.20). This result shows a transducer may be treated as a point-transducer (i.e., as having no spatial-extent) with the physical spatial-extent of the transducer causing a filtering of the spatial-frequency bandwidth. This filtering is usually known as the *beam-pattern* of the transducer.

Figure 2.2 Aperture illumination for a rectangular aperture, $i_\mathrm{t}(y)$.

Rectangular transducer

Given a transducer with an aperture illumination function

$$i_\mathrm{t}(y) = \operatorname{rect}\left(\frac{y}{D}\right), \tag{2.26}$$

(shown in Figure 2.2) its beam-pattern is given by

$$I_\mathrm{t}(k_y) = D \operatorname{sinc}\left(\frac{k_y D}{2\pi}\right). \tag{2.27}$$

The beam-pattern, $I_\mathrm{t}(k_y)$, is shown in Figure 2.3. With the nulls of (2.27) being given by $k_y = 2\pi n/D$ where $n \in \pm1, \pm2, \pm3, \ldots$, and using the relation $k_y = k\sin\theta$, the null-to-null width of the transducer's main-lobe may be calculated as

$$
\begin{aligned}
\theta_\text{null-to-null} &= \sin^{-1}\left(\frac{4\pi}{kD}\right) \\
&= \sin^{-1}\left(\frac{2\lambda}{D}\right) \\
&\approx \frac{2\lambda}{D},
\end{aligned}
\tag{2.28}
$$

Figure 2.3   Aperture illumination for a rectangular aperture, $I_{\mathrm{t}}(k_y)$.

and the -3 dB[6] beam-width, θ_{3dB}, is

$$\theta_{\mathrm{3dB}} \approx \frac{\lambda}{D}, \tag{2.29}$$

where the result applies in the *far-field* of the transducer. This far-field (or *Fraunhofer*) region starts at a range of [Goodman 1968]

$$\frac{2D^2}{\lambda}, \tag{2.30}$$

which is called the *Rayleigh range*. The far-field region is the distance from the transducer where the wave fronts arriving from a single point source may be treated as plane-waves (i.e., where the wavefront curvature may be ignored). For distances closer than this region, the *near-field* or *Fresnel region*, the effective beam-pattern is much more complicated[7].

The resolution possible from a rectangular transducer is fixed by the effective length

---

[6]λ/D is actually the -3.9 dB bandwidth; for the purposes of the discussion in this thesis the approximation is sufficient.

[7]The ranges of interest are usually in the far field of an individual transducer due to its small D. When considering the beam-pattern of an array of transducers, the effective D may be large causing near-field effects at the ranges of interest.

of the transducer

$$D = {}_{\max}\{D_{\mathrm{r}}, D_{\mathrm{t}}\}, \tag{2.31}$$

where D_{t} and D_{r} are the transmit and receive transducer lengths respectively. Thus the along-track resolution, δ_y, can by expressed in terms of the transducer length D as

$$\begin{aligned}\delta_y &= R\theta_{\mathrm{3dB}}, \\ &= \frac{R\lambda}{D}.\end{aligned} \tag{2.32}$$

This resolution is fixed immediately on signal transmission and again on reception by the physical extent of the transducer.

2.5.2 Real arrays

A real array is formed when the signals from several transducers are combined. The coherent summation of signals from the individual transducers produces a signal equivalent to that from a single transducer the same shape as the array[8].

The benefit of recording the transducers separately is that the array may be *focused* or *beam-steered* [Nielsen 1991; Urick 1975] using computer processing. This extra flexibility is possible because the echo data are recorded. Receive beam-steering and focusing may be performed for the entire image in post-recording processing (also called *dynamic focusing*).

The along-track resolution of a focused linear array is given by

$$\delta_y = \frac{R\lambda}{L}, \tag{2.33}$$

where R is the range to target and L is the length of the array—either the transmitter or receiver array, whichever is longer. This resolution is fixed for a given length L at the time the signal is focused. The motivation for the use of synthetic aperture arrays is to increase L and gain an associated increase in resolution.

2.5.3 Synthetic arrays

Synthetic arrays are conceptually similar to real arrays; a synthetic array samples the transducer locations of a real array using the constant forward motion of the imaging platform. If the scene does not vary with time (and the path of the platform is predictable) then an equivalent aperture of arbitrary length may be synthesised. Resolution is given by $\delta_y = R\lambda/(2L)$, where L is the aperture length and R the range to the target, thus increasing the length of the effective aperture improves resolution.

---

[8]This is possible to show by partitioning (2.22) for multiple-receivers and summing the resulting signals.

The resolution of a synthetic aperture system is improved by a factor of 2 when compared to a real aperture system[9]. Taking this into account, the resolution of a synthetic aperture system is given by

$$\delta_y = \frac{R\,\lambda}{2L_{\text{synthetic}}}, \tag{2.34}$$

which gives the resolution for a spotlight system. In practice, the effective aperture length of a spotlight systems is constrained by its beam steering angle limit.

Stripmap systems have their maximum possible synthetic array length $L_{\text{synthetic}}$ limited by beam-patterning effects. In practice, the length of aperture able to be synthesised for a given range is limited by the transducer beam-width,

$$\begin{aligned} L &= R\,\theta_{\text{3dB}}, \\ &= \frac{R\,\lambda}{D}, \end{aligned} \tag{2.35}$$

where D is the larger of transmit, D_{t}, and receive, D_{r}, aperture extents, and the -3 dB beam-width is given by (2.29). Thus the resolution of a stripmap synthetic aperture system is given by

$$\delta_y = \frac{R\,\lambda D}{2R\,\lambda}, \tag{2.36}$$

$$= \frac{D}{2}, \tag{2.37}$$

where D, as before, is the effective transducer extent. Note that this resolution is independent of range and imaging frequency, something not possible with a real aperture system. This allows long-range, low-frequency, high-resolution imagery without excessively long arrays.

Factor of two resolution improvement

Synthetic aperture arrays have a resolution that is half of that a real array of the same length achieves. This factor of two improvement comes about because of the phase-doubling effect caused by two-way propagation over multiple transmitted pulses.

For illustration, a hypothetical real aperture system that has an omni-directional transmitter and a receiver array of length L_{real} is considered and compared with a monostatic synthetic aperture system obeying the stop-and-hop approximation (see Chapter 3 for a more thorough discussion of these models). The extent of spatial frequencies (and thus) resolution of the real aperture system is fixed by the angular

---

[9]See later in the section.

extent at reception. The spatial frequencies at reception are described by

$$k_u = k \sin \theta, \tag{2.38}$$

and resolution is directly calculated from beam-width. In the synthetic system both the transmitter and receiver move causing a phase-doubling (because of the two-way propagation). The phase doubling is sometimes treated as if the wave-speed were halved (which is equivalent to wave-number doubling—see Chapter 3) and the spatial frequencies are described by

$$k_u = 2k \sin \theta. \tag{2.39}$$

Comparing the spatial frequencies of the real and synthetic systems (2.38) and (2.39), it is apparent that for the same beam-width the synthetic aperture system has twice the spatial frequency extent. This leads to the factor of two resolution improvement of synthetic systems (and also altered sampling constraints).

Another way of considering the problem is to treat the real array using the phase-centre approximation (see Chapter 3). Under this approximation (a transmitter/receiver pair is modelled as a co-located transducer midway between), a real array can be seen to be equivalent to a synthetic array of half the length (and thus resolution is poorer by a factor of two). Bistatic modelling, such as used in Chapter 3, or spatial frequency analysis [Hawkins 1996; Soumekh 1994] can be used to arrive at the same conclusion.

2.6   Sidescan SAS imaging

The sidescan imaging mode refers to the scenario where the sonar system images to the side at a low grazing-angle. This mode is commonly used for mapping, producing images roughly similar to aerial photographs in appearance. This section summarises some of the important features and constraints of the particular imaging mode as it relates to SAS imaging.

2.6.1   Collection geometry

As described above, the imaging sonar looks to the side and effectively maps many range-image strips. To simplify understanding, the geometry is regarded as being 2-D but the extension of the ideas into 3-D space is straightforward. Chapter 3 delves into the 3-D geometry in more detail.

Figure 2.4 illustrates the commonly used 2-D ground-plane (flat) geometry. The sonar system records the response of the targets in range within the beam as a function of time-delay, t, for many different along-track positions, u. The scene is then reconstructed into the image coordinates x, y.

Figure 2.4 Ground-plane geometry for a sidescan SAS system. Image coordinates (x, y), sonar coordinates (u, t), scene offset (x_0, y_0).

2.6.2 Range resolution

The range resolution of a sidescan SAS system is the same as for any pulse-echo ranging system. The resolution of a pulse-echo system is given by

$$
\begin{aligned}
\delta_r &= \frac{c}{2}\delta_t, \\
\delta_r &= \frac{c}{2\,B_c},
\end{aligned}
\tag{2.40}
$$

where B_c is the transmitted bandwidth and c the wave speed in the medium.

2.6.3 Range constraints

The repeated pulse transmission of side-scan systems causes targets outside a given range to appear in the echo from the subsequent pulse. This results in range ambiguity.

The unambiguous range of the system is limited by the pulse repetition frequency (PRF). A higher PRF implies that the maximum unambiguous range is reduced. The maximum range is given by the distance that the transmitted sound pulse travels out and back before the next pulse is transmitted,

$$
x_{\text{unambiguous}} = \frac{c\,\tau_{\text{rep}}}{2},
\tag{2.41}
$$

where τ_{rep} is the repetition period of the transmitting system.

By way of example, for the KiwiSAS-III travelling at $4\,\text{knots}$ $(2\,\text{ms}^{-1})$ and using

$D/4$ sampling[10] τ_{rep} must be kept below $0.04\,\mathrm{sec}$ (a sample rate of close to $25\,\mathrm{Hz}$) giving a maximum unambiguous range of only $30\,\mathrm{m}$.

If a larger unambiguous range is desired, either the system pulse repetition frequency (PRF) must be decreased or multiple orthogonal signals must be sent. Decreasing the system PRF tends to cause along-track undersampling requiring slower system speeds or additional transducers. Alternating two orthogonal pulses is roughly equivalent to doubling the system PRF. Using multiple orthogonal pulses in flight simultaneously is a technique employed in space-borne SAR imaging systems [Curlander and McDonough 1996]. However, a system employing orthogonal chirp transmission suffers a degradation in SNR [Axelsson 2001]. When the orthogonal signals are cross-correlated some energy still occurs as no finite-time signal can be truly orthogonal with another finite-time signal. This undesired cross-correlation energy causes an increase in the noise floor of the image.

2.6.4 Along-track resolution

The along-track resolution of a synthetic aperture system is half of the transducer extent, i.e.,

$$\delta_y = \frac{D}{2}, \tag{2.42}$$

where D is the larger of either the transmitting aperture, D_{t}, or the receiving aperture, D_{r}. The resolution limit is due to the limiting of the synthetic aperture length by the transducers. The resolution is also half that of a real aperture system of the same length as the synthetic aperture. This is due to the transmitter and receiver both moving in a synthetic system causing phase (and along-track spatial frequency) doubling.

2.6.5 Along-track sampling constraint

The appropriate along-track sampling constraint is a contentious issue in the SAS community. The true along-track sampling constraint of any synthetic aperture system[11] requires sampling finer than $\lambda/4$. Practically, this constraint is much tighter than necessary and a figure of $D/2$, where D is calculated as in Sections 2.5.2 and 2.6.4, is often given as a sampling constraint [Douglas and Lee 1992; Hayes and Gough 1992; Rolt and Schmidt 1992; Tomiyasu 1978][12].

The sampling constraint chosen for a synthetic aperture system is dependent on the along-track ambiguity to signal ratio (AASR) and peak to grating lobe peak ratio (PGLR) desired in the final image. PGLR provides an estimate of the dynamic range and improves with increasing system bandwidth due to grating lobe smearing.

---

[10]See Section 2.6.5.

[11]This is equivalent to having $180°$ beam-width—i.e., point transducers.

[12]This constraint is arrived at by considering the windowing (bandwidth limiting) effects the finite aperture has on the spatial frequency spectrum.

AASR does not improve with increasing bandwidth[13] and grating lobe smearing is finite [Châtillon 2000; Gough and Hawkins 1997]. A better along-track sampling requirement which avoids some of the adverse effects is

$$\Delta u = \frac{D}{4}. \tag{2.43}$$

Sampling at a rate of $D/4$ prevents the main-lobe of the beam-pattern being aliased [Hawkins 1996]. When the main-lobe does not alias, the alias target energy is much lower than the energy of the true target; however, the side lobes of the beam-pattern do still alias resulting in some finite energy being aliased. For example the AASR for a $D/2$ sampled system is $\approx$ -8 dB whereas is $\approx$ -21 dB for a $D/4$ sampled system.

Multiple-receiver systems (vernier arrays)

Multiple-receiver arrays provide a method of overcoming the along-track sampling constraint. Systems using multiple-receiver geometry gain an increase in the sampling rate equal to the number of receivers used [de Heering 1982; Gilmour 1978]. Thus the sampling constraint (2.43) becomes

$$\Delta u = \frac{N_h D}{4}. \tag{2.44}$$

Noting that $N_h D$ is the length of the receiver array, D_{array}, the sampling constraint may instead be written

$$\Delta u = \frac{D_{\text{array}}}{4}. \tag{2.45}$$

Systems employing many receivers on relatively short arrays are thus able to provide high-resolution imagery at realistic mapping rates.

Vernier-array systems are now common in SAS imaging and their use is likely to continue. Image reconstruction algorithms have to be adapted slightly to account for the altered collection geometry (see Section 5.2).

2.7   Summary

A good grasp of the SAS fundamentals helps the understanding of its more difficult problems. This chapter summarises some of the fundamental principles required for designing and processing data from SAS systems. The relationship between recorded sonar echoes and the complex baseband signals has been summarised. Complex baseband representations of sonar signals offer storage benefits as well as simplified algorithm implementation. Complex baseband signals and their pre-envelopes are used throughout this thesis. Pulse-compression, which allows improved resolution without

---

[13] Increasing the system bandwidth only causes increased smearing of the grating-lobe. Since it does not have an influence on the grating lobe energy the AASR is unchanged.

sacrificing range, has been summarised. Pulse-compressed data is the preferred input for the majority of the SAS reconstruction techniques. Improved wave propagation modelling for SAS imaging has been presented. This modelling is used in the following chapter to obtain improved SAS system models. Beam-patterning effects are derived using the improved modelling and correspond to those used previously. SAS along-track sampling requirements and other important design parameters have been summarised with regard to the aperture extent and swath-width requirements. An along-track sampling rate of $D/3$ to $D/4$ is recommended based on AASR considerations.

System modelling

To be able to reconstruct images of a sea-floor scene an understanding of the imaging process is necessary. An accurate system model is a good step in achieving that understanding, particularly if the approximations made in deriving the model are well understood.

This chapter outlines the basic SAS system modelling needed to design and implement reconstruction and autofocus techniques. To do this, the *bistatic* scattering geometry for multiple-receiver sonars is considered and compared to the more usual monostatic (single receiver sonar) system modelling. This modelling is performed as if in the ground-plane geometry—Section 3.1 describes this particular geometry. (A system model in 3-D is described in Section 3.4.) It is shown in Section 3.3 that multiple-receiver sonars can be treated as a monostatic equivalent sonar using the phase-centre approximation. Errors resulting from the phase-centre approximation are discussed with the intent of correction during reconstruction. A summary of some of the (often neglected) second-order effects in SAS modelling (such as the stop-and-hop approximation and temporal Doppler shifting) is also included.

Modelling of the complicated scattering from objects on the sea-floor and ray-bending (common in underwater imaging) are only discussed in passing in this thesis, the interested reader is referred to Ziomek [1995] for information on these topics. Scene variant ray-bending (medium fluctuation) has the potential to cause image blurring and is discussed briefly with regard to autofocus in Chapter 7.

3.1   Ground plane geometry

The imaging sonar travels a nominally straight flight-path at a given height above the sea-floor. In side-scanning mode, the sonar images the sea-floor at a low grazing-angle. This is illustrated in Figure 3.1

To simplify the mathematics, the height dimension to the problem is removed [Soumekh 1992]. This is done by replacing the across-track and height dimensions with a single dimension of slant-range. Figure 3.2 illustrates this transformation. The new

Figure 3.1   General 3-D geometry for side-scan SAS imaging. The sonar travels above the floor at a fixed height and images down and to the side. Axis variables are as shown.

Figure 3.2   Slant-range mapping. The across-track distance variable x and the height variables h, z are replaced by a single parameter x_s.

slant-range variable is x_s, which is given by the non-linear (polar) mapping

$$x_s = \sqrt{x^2 + (h - z)^2}, \tag{3.1}$$

$$\eta = \tan^{-1}\left(\frac{h - z}{x}\right), \tag{3.2}$$

where z is the sea-floor height and h the height of the sonar. If the image is recon-
structed in the slant-range (normal procedure for 2-D reconstruction algorithms), the
new slant-plane image, $f_s(x_s, y)$, is given by

$$f_s(x_s, y) = \int_{-\pi}^{\pi} f(x_s, y, \eta) x_s \, \mathrm{d}\eta, \tag{3.3}$$

where $f(x, y, z)$ and $f(x_s, y, \eta)$ represent the true 3-D reflectivity function on cartesian
and slant-range coordinates.

2-D image reconstruction techniques reconstruct the slant-plane image, $f_s(x_s, y)$;
i.e., the image coordinates are in x_s, y. Reconstruction without correcting for the
ground-plane mapping causes some range variant image compression and shifting. Re-
construction onto the true ground map x, z is possible if the elevation map is known
before reconstruction [Soumekh 1992], or estimated using some form of bathymetry
[Banks et al. 2001].

The system models derived in the following sections assume a ground-plane where
the imaging height $h - z = 0$; this makes x and x_s equivalent. A 3-D system model is
presented in Section 3.4 for dealing with situations where $h - z \neq 0$.

Aside from making the mathematics easier, another reason for making the slant-
plane transformation is that the sea-floor is usually the only object of interest in a
sonar image. Reconstructing a full 3-D image is not required when most of the image
is empty space.

3.1.1 Multiple-receiver geometry

Current SAS systems employ multiple-element receiver-arrays to mitigate the sampling
constraints described in Section 2.6.5. The constraints are now based on the length of
the receiver-array rather than the length of an individual receiver. The benefit of
multiple-receiver geometry is that resolution is determined by the individual receiver
length and sampling determined by the array length. A SAS with N_h receivers is thus
able to travel N_h times faster than a single-receiver SAS with the same resolution.

Sonar system models must account for the extra complexity caused by multiple-
receiver collection. Multiple-receiver sonars transmit from a single transmitter. The
reflected wave propagates back to the sonar and is received at the hydrophones of the
receiver array.

In the ground-plane (discussed in the previous section), the geometry of a multiple-
receiver system is as shown in Figure 3.3 where u is the along-track transmitter position,
σ is the distance from transmitter to hydrophone, and x, y are the across and along
track positions of the scene.

Figure 3.3 Bistatic collection geometry of a multiple-receiver sonar. The received echo is a function of both the projector position, u, and the hydrophone position relative to the projector, σ, as well as the target location.

3.2   Multiple-receiver system model

Synthetic aperture sonars perform undersea-imaging by interrogating the sea-floor with acoustic energy and measuring the back-scattered field. The outgoing acoustic pulse hits a region of differing density or refractive index, causing some energy to be scattered back to the sonar receiver.

To model this interaction accurately, both the outgoing and incoming acoustic paths should be regarded separately. Modelling both paths is known as the *bistatic* imaging model. This is important in modelling multiple-receiver SAS systems. The simpler alternative, *monostatic* modelling (described in Section 3.3), is to treat the sonar as if both transmission and reception were made on the same element and the incoming and outgoing paths are identical.

The multiple-receiver SAS model starts by approximating the reflection caused by the propagating incident acoustic pulse as a set of re-radiating secondary sources (following Huygen's principle [Goodman 1968]). The energy from the secondary sources propagates back to the sonar and is measured by the receiver hydrophones. During the time the pulse is in the water column, it is assumed that the imaging sonar is stationary[1]. This allows the slow-time dimension (travel along the u axis) to be decoupled from fast-time (the time variable t) and simplifies the mathematics (see Section 3.5.1).

Bistatic system modelling is common in many fields related to Fourier imaging. Soumekh [1991, 1994, 1999] outlines system models for both a true bistatic SAR arrangement of separate transmit and receive imaging platforms and the multiple receive

---

[1]This approximation is called the *stop-and-hop approximation* [Douglas 1993; Hawkins 1996; Hayes 1989] and is dealt with in more detail in Section 3.5.1.

element systems considered here. Bistatic imaging models are also in common use in the seismic imaging fields [Soumekh 1994], medical ultrasound [Lerner and Waag 1988], and ground penetrating radar [Leuschen and Plumb 2001].

Using the 2-D Green's function (2.20) for both outgoing and incoming propagation paths, the response from a general scene $f(x, y)$ may be given as

$$e(\omega, \sigma, u) = \frac{1}{16} P(\omega) \iint f(x, y) H_0^{(2)}(k\rho_{\text{in}}) H_0^{(2)}(k\rho_{\text{out}}) \, \mathrm{d}x \, \mathrm{d}y, \tag{3.4}$$

where u is the position of the transmitter, σ is the position of the hydrophone relative to the transmitter position, $\rho_{\text{out}} = \sqrt{x^2 + (y - u)^2}$, and $\rho_{\text{in}} = \sqrt{x^2 + (y - u - \sigma)^2}$. If the asymptotic approximation of Hankel functions is used [Chew 1995; Morse and Feshbach 1953], (3.4) may be approximated as

$$e(\omega, \sigma, u) \approx \frac{1}{8\pi jk} P(\omega) \iint \frac{f(x, y)}{\sqrt{\rho_{\text{in}}\rho_{\text{out}}}} \exp\left(-jk(\rho_{\text{in}} + \rho_{\text{out}})\right) \mathrm{d}x \, \mathrm{d}y, \tag{3.5}$$

which is valid for

$$\rho_{\text{in}}, \ \rho_{\text{out}} \gg \lambda \equiv \frac{2\pi c}{\omega} \equiv \frac{2\pi}{k}.$$

For designing inversion/reconstruction methods it is useful to have wavenumber domain representations of the above system models. Starting with (3.4) and taking 2-D spatial Fourier transforms in similar manner to that employed in Section 2.4, gives the wavenumber domain system model as [Callow et al. 2001a, 2002b; Soumekh 1991]

$$E(\omega, k_\sigma, k_u) = P(\omega) \frac{f(k_x, k_y)}{\sqrt{k^2 - k_\sigma^2}\sqrt{k^2 - (k_u - k_\sigma)^2}}, \tag{3.6}$$

where the Fourier variables are related via [Callow et al. 2001a]

$$k_x = \sqrt{k^2 - k_\sigma^2} + \sqrt{k^2 - (k_u - k_\sigma)^2}, \tag{3.7}$$

$$k_y = k_u. \tag{3.8}$$

3.2.1   Offset variables

The models presented above are presented for the bandpass signals (pre-envelope signals). The signals in the sonar/radar fields are usually complex baseband with a known offset frequency. Spatial variables are usually measured relative to an offset where, for example, a small swath is taken about a large range offset. In addition, reconstruction algorithms perform better if mapping between baseband variables as this prevents interpolation error caused by high-frequency phase functions [Gough and Hawkins 1997].

In his thesis, Hawkins [1996] takes particular care in discussing offset frequency

measurements (baseband signals) when considering reconstruction algorithms. System models for offset measurements are derived by substituting offset measurement parameters and remembering that a spatial-offset causes phase-shifting in the Fourier domain.

Example—across-track offset:

Often the collected image is centred about a particular range offset. Making the substitution $x' \equiv x - x_0$ and $t' \equiv t - t_0$, where x_0 is the scene offset and t_0 is the time offset of the echo signal, (3.4) becomes

$$
\begin{aligned}
e_{t'}(\omega, \sigma, u) = \frac{1}{16} P(\omega) \exp\left(j\omega t_0\right) & \\
\iint f_{x'}(x', y) H_0^{(2)}\left(k\sqrt{(x' + x_0)^2 + (y - u)^2}\right) & \\
H_0^{(2)}\left(k\sqrt{(x' + x_0)^2 + (y - u - \sigma)^2}\right) \mathrm{d}x' \, \mathrm{d}y, &
\end{aligned} \tag{3.9}
$$

where $e_{t'}(\omega, \sigma, u)$ is the equivalent of $e(\omega, \sigma, u)$ referenced to t_0.

The equivalent of (3.6) follows as

$$
E_{t'}(\omega, k_\sigma, k_u) = P(\omega) \exp\left(j\omega t_0\right) \frac{f_{x'}(k_x, k_y) \exp\left(-j k_x x_0\right)}{\sqrt{k^2 - k_\sigma^2}\sqrt{k^2 - (k_u - k_\sigma)^2}}, \tag{3.10}
$$

where the Fourier phases of $f_{x'}(k_x, k_y)$ are relative to the offset variable x'—i.e., $f_{x'}(k_x, k_y)$ is the offset equivalent of $f(k_x, k_y)$. In the usual imaging scenario, the echo signal's time offset is related to the scene offset by

$$
t_0 = \frac{2x_0}{c}, \tag{3.11}
$$

and (3.10) becomes

$$
E_{t'}(\omega, k_\sigma, k_u) = P(\omega) \exp\left(j 2k \, x_0\right) \frac{f_{x'}(k_x, k_y) \exp\left(-j k_x x_0\right)}{\sqrt{k^2 - k_\sigma^2}\sqrt{k^2 - (k_u - k_\sigma)^2}}, \tag{3.12}
$$

expanding for k_x

$$
\begin{aligned}
E_{t'}(\omega, k_\sigma, k_u) = \\
P(\omega) \frac{f_{x'}(k_x, k_y) \exp\left(-j(\sqrt{k^2 - k_\sigma^2} + \sqrt{k^2 - (k_u - k_\sigma)^2} - 2k)x_0\right)}{\sqrt{k^2 - k_\sigma^2}\sqrt{k^2 - (k_u - k_\sigma)^2}}.
\end{aligned} \tag{3.13}
$$

Comparable derivations for offsets in frequency [Gough and Hawkins 1997] (also Section 5.1.1) and along-track position [Soumekh 1994] are calculated in the same way.
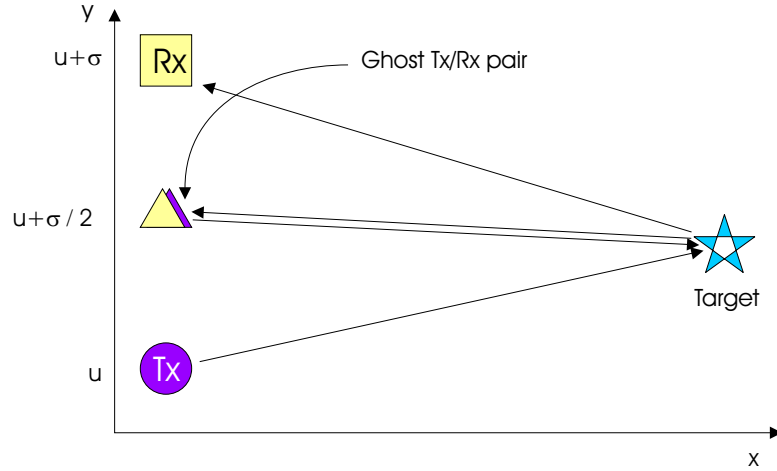
Figure 3.4 Phase-center geometry, the bistatic transmitter/receiver pair is modelled using an equivalent co-located transducer midway between the two.

3.2.2 Phase-centre approximation

One approximate method of modelling a true bistatic sonar is to treat it as a monostatic sonar measuring many individual samples. This modelling method is often chosen to show the possibility of using fast single-receiver reconstruction techniques on multiple-receiver data. The phase-centre approximation [Bellettini and Pinto 2002; Bonifant 1999; Wilkinson 2001] allows this modelling. The basis of the approximation is to treat a bistatic transmitter/receiver pair as if it were a single co-located transducer located midway between the two. By doing this, a vernier-array sonar may be modelled as a single-receiver sonar taking samples at the positions of the phase centres.

It is possible to show how this approximation affects the wavenumber domain representations of the system model (3.4) [Callow et al. 2001a; Soumekh 1991]. Noting that the phase-centre sonar Fourier data only exists on the line where [Callow et al. 2001a][2]

$$2k_\sigma = k_u,$$

then (3.6) becomes

$$E(\omega, k_\sigma, k_u) = P(\omega)\frac{f(k_x, k_y)}{|k^2 - k_\sigma^2|}, \tag{3.14}$$

where the coordinate transform of (3.7) becomes

$$k_x = \sqrt{k^2 - k_\sigma^2} + \sqrt{k^2 - k_\sigma^2}, \tag{3.15}$$

$$= \sqrt{4k^2 - k_u^2}, \tag{3.16}$$

$$k_y = k_u. \tag{3.17}$$

---

[2]This assumption is equivalent to assuming that the angle of transmission and reception is identical.

The error caused in making the phase-centre approximation is the difference between the two-way bistatic path and the two-way equivalent monostatic path. Writing out the approximation error, ϵ, gives

$$\epsilon = \sqrt{x^2 + (u - y)^2} + \sqrt{x^2 + (u + \sigma - y)^2} - 2\sqrt{x^2 + (u + \sigma/2 - y)^2}, \qquad (3.18)$$

where σ is the position of the hydrophone. The approximation error (3.18) has been derived previously [Bellettini and Pinto 2002; Bonifant 1999; Wilkinson 2001]. Taking a Taylor expansion about $\sigma = 0$ to linearise the problem gives the error as [Pinto et al. 2002; Soumekh 1991]

$$\epsilon \approx \frac{\sigma^2}{4r}\cos^2\theta + \frac{\sigma^4}{64r^3}\cos^2\theta\left(4 - 5\cos^2\theta\right) + \dots, \qquad (3.19)$$

where θ and r are the angle and range to the target from the co-located transducer.

The phase-centre approximation treats the in-scene targets as if they were in the far-field of the physical receiver array [Banks and Griffiths 2002] (but not that of the synthetic array). Reconstruction techniques using the approximation must take the error of phase-centre approximation into account when the targets are in the near-field of the receiver array.

3.3   Single-receiver model

A monostatic sonar transmits and receives using the same physical element. Similar data are obtained using a bistatic sonar where a single-receiver element is used that is physically located close to the transmitter element; these sonars are usually treated as monostatic [Gough and Hawkins 1997]. Most early SAS and the majority of SAR systems may be modelled in this fashion [Gough and Hawkins 1998; Hayes and Gough 1992; Jakowatz and Wahl 1993].

To model a monostatic sonar, some approximations regarding the physical model are made to make the mathematics tractable. Instead of treating the scene as a collection of scatterers and modelling the two-way path, the scene is modelled as a collection of self radiating, coherent sources where the acoustic wave travels at half the true propagation speed. This physical approximation models the phase terms correctly but the amplitude terms are approximate since only one-way spreading losses are considered. Making this physical approximation, called the *exploding sources* model [Bonifant 1999], it is possible to directly apply the wave equation results from Section 2.4.

Using the exploding sources model, the spatial impulse response for a realistic monostatic sonar is given by (assuming $\rho \gg \lambda$, see Section 2.4)

$$g(x, y) = \frac{-1}{\sqrt{j16\pi k\rho}} \exp\left(-j2k\rho\right), \qquad (3.20)$$

(cf. with (2.21), in particular note the additional factor of 2). Rewriting this in more familiar notation by expanding ρ gives the response, $E(\omega, u)$, from a scene $f(x, y)$ as

$$E(\omega, u) \approx P(\omega) \iint \frac{-f(x, y)}{\sqrt{j16\pi k \sqrt{x^2 + (u - y)^2}}} \exp\left(-j2k\sqrt{x^2 + (u - y)^2}\right) \mathrm{d}x\,\mathrm{d}y.$$
(3.21)

If the amplitude terms are neglected (because of incorrect spreading losses inherent in the previous sonar model), this is the same model seen in previous sonar literature [Hawkins 1996; Hayes and Gough 1992],

$$E(\omega, u) \approx P(\omega) \iint \frac{f(x, y)}{8\pi \sqrt{x^2 + (u - y)^2}} \exp\left(-j2k\sqrt{x^2 + (u - y)^2}\right) \mathrm{d}x\,\mathrm{d}y.$$
(3.22)

The time-domain equivalent of (3.22) is given by

$$e(t, u) \approx \iint \frac{f(x, y)p\left(t - \frac{2}{c}\sqrt{x^2 + (u - y)^2}\right)}{8\pi \sqrt{x^2 + (u - y)^2}} \mathrm{d}x\,\mathrm{d}y,$$
(3.23)

where again the amplitude terms are approximated as a one-way spreading loss.

Similar derivations to those in previous sections (using Weyl's identity and avoiding approximation) allow the wavenumber domain representation of (3.21) to be written as

$$E(\omega, k_u) = P(\omega) \frac{f(k_x, k_y)}{\sqrt{4k^2 - k_u^2}},$$
(3.24)

where the Fourier coordinate transform (the Stolt mapping [Bamler 1992; Hawkins 1996]) is given by

$$k_x = \sqrt{4k^2 - k_u^2},$$
(3.25)

$$k_y = k_u.$$
(3.26)

Note that (3.24), (3.25), and (3.26) are equivalent to (3.14), (3.16), and (3.17). Thus a monostatic sonar collects the *same data* as a phase-centre sonar sampling at the same positions. Aside from the amplitude terms, a hypothetical phase centre sonar is equivalent to a monostatic sonar. Using the phase-centre approximation, single-receiver reconstruction techniques are able to be employed on multiple-receiver data sets.

3.4   3-D model

Bathymetry (height measurement) is possible using a system with multiple, vertically-displaced receivers. Interferometric techniques for performing this estimate are of par-

ticular interest to the sonar community. This section presents a model for a single-receiver SAS system in an unbounded 3-D medium.

Following similar arguments to those of Sections 3.2 and 3.3, and using the free space Green's function for 3-D instead of 2-D, the wavenumber domain representation of the echo signal may be calculated as[3]

$$E(\omega, k_u, k_h) = \frac{f(k_x, k_y, k_z)P(\omega)}{\sqrt{4k^2 - k_u^2 - k_h^2}}. \tag{3.27}$$

The Fourier domain coordinate transform in (3.27) is given by

$$k_x = \sqrt{4k^2 - k_u^2 - k_h^2}, \tag{3.28}$$

$$k_y = k_u, \tag{3.29}$$

$$k_z = k_h, \tag{3.30}$$

and $(z, h) \leftrightarrow (k_z, k_h)$ are the Fourier pairs for the global and sonar height respectively.

Another representation of (3.27) is obtained by inverse Fourier transformation over k_x, k_z and k_h. Thus

$$E(\omega, k_u, h) = -j\frac{1}{4}P(\omega) \iint f(x, k_y, z)H_0^{(2)}\left(\sqrt{x^2 + (h-z)^2}\sqrt{4k^2 - k_u^2}\right) dx\, dz. \tag{3.31}$$

Using the slant-range mapping of Section 3.1 this may be written

$$E(\omega, k_u, h) = -j\frac{1}{4}P(\omega) \iint f(x_s, k_y, \eta)H_0^{(2)}\left(x_s\sqrt{4k^2 - k_u^2}\right)x_s\, dx_s\, d\eta, \tag{3.32}$$

or more compactly as

$$E(\omega, k_u, h) = \frac{P(\omega)f_s(k_{xs}, k_y)}{\sqrt{4k^2 - k_u^2}}, \tag{3.33}$$

where

$$k_{xs} = \sqrt{4k^2 - k_u^2}. \tag{3.34}$$

Modelling the 3-D propagation problem in this way shows that the beam-forming in azimuth and elevation are separable. Each direction can be reconstructed separately. 2-D reconstruction simply obtains various slant-plane images which may be combined to achieve 3-D imagery. Alternately, full 3-D reconstruction via (3.27) could be used avoiding the slant-range representation. These techniques have a possible application in simplifying interferometry by avoiding the iterative estimation otherwise used (see [Banks et al. 2001] for an example of iterative interferometric techniques).

---

[3]Again using 1-way propagation with a halved wave-speed.

3.5 Second order effects

It is impractical to include the effects of all phenomena in the system model when describing complex physical phenomena. Approximations and assumptions that apply for some sonar systems do not hold for others—a good example is the phase-centre approximation made in the previous section. This section describes some of the physical effects that have not been modelled above. These *second order* effects become important and must be accurately modelled where improvements in the mapping rates and resolution of sonar systems are sought.

3.5.1 Stop-and-hop approximation

In the derivation of the models above, it was assumed that the sonar transmits at a particular position and waits for all of the echoes to return before moving instantaneously to the next transmit position; the so-called *stop-and-hop* model. A real sonar moves continuously along the aperture during data collection. The time and along-track position parameters that are treated as decoupled by the stop-and-hop approximation are loosely coupled. This section considers modelling the continuous sonar movement between pulses and the related coupling of the along-track position parameter.

The along-track coupling causes two major effects: a temporal Doppler shifting of the echo pulse signal (due to movement during the sonar ping) and a difference between receive and transmit positions. The Doppler shifting is modelled in Section 3.5.1 and has only a minor influence on most SAS images (if Doppler tolerant waveforms are employed [Hawkins 1996]).

Temporal Doppler

Temporal Doppler, the scaling of the echo spectra for targets forward and aft of the beam caused by sonar motion, is also neglected using the stop-and-hop assumption. The spectrum of returns received forward of the sonar are scaled in frequency (scaling factor $\alpha > 1$). Similarly for targets aft of the sonar ($\alpha < 1$).

Doppler shifting/spectral scaling of the echo returns in the fashion described, causes a small phase error in reconstructed imagery if the effect is not taken into account. This effect is not a problem for slow moving sonars with a narrow beam-width [Hawkins 1996; Hayes 1989]. However, as higher and higher resolutions are demanded, beam-widths become wider and the Doppler shifting becomes more pronounced. This becomes particularly apparent for the faster moving SAS systems likely to be used in the near future. The resulting imagery suffers a slight geometrical error and some minor blurring if the effect is not modelled. Temporal Doppler shifting is modelled starting with

the expression for acoustic Doppler shifts[4] [Hayes 1989]

$$k_r = k\left(\frac{c + v_s}{c - v_s}\right), \tag{3.35}$$

where k_r is the received frequency, k is the transmitted frequency, c is the wave-speed in the medium, and v_s is the velocity of the sonar. Equation (3.35) applies for a target directly in front of the sonar; generalising for target angle gives

$$k_r = k\left(\frac{c + v_s \sin\theta_\sigma}{c - v_s \sin\theta_u}\right), \tag{3.36}$$

where θ_σ is the angle to the target at signal reception and θ_u the angle at transmission. Since

$$k_u \equiv k \sin\theta_u, \tag{3.37}$$

$$k_\sigma \equiv k_r \sin\theta_\sigma, \tag{3.38}$$

k_r may be solved for using the quadratic equation. For notational simplicity let,

$$k_r = q(k, k_\sigma, k_u), \tag{3.39}$$

Equation (3.39) is useful for including within-pulse Doppler in the wavenumber domain representation of the system model. Thus the new bistatic system model (cf. (3.6)) is

$$E(\omega, k_\sigma, k_u) = P(\omega)\frac{f(k_x, k_y)}{\sqrt{k_r^2 - k_\sigma^2}\sqrt{k_r^2 - (k_u - k_\sigma)^2}}, \tag{3.40}$$

where the Fourier variables are related via [Callow et al. 2001a]

$$k_x = \sqrt{k_r^2 - k_\sigma^2} + \sqrt{k_r^2 - (k_u - k_\sigma)^2}, \tag{3.41}$$

$$k_y = k_u. \tag{3.42}$$

Thus the Doppler inclusive system model is almost identical to that normal bistatic model with the predictable Doppler shifting encapsulated within k_r.

Image Skew

Airborne SAR systems do not normally consider the position shift between transmission and reception. When the medium propagation speed is the speed of light, the imaging platform does not move very far between transmission and reception causing very little error. Even with current SAS systems, where tow velocity can be of the order of 1% of

---

[4]This assumes a fixed medium, a moving sonar, and stationary targets. Relativistic effects have also been neglected.

the medium wave-speed and wide swath widths are used [Keeter 2001], the monostatic model starts to break down because of the position shift.

Image skew is due to range dependent position shift. Bonifant [1999] notes this problem in his thesis where a phase correction is suggested during reconstruction to account for the approximation error. Soumekh [1991] undertakes a detailed analysis and derives a model for a system receiving the signal at a different position to transmission. Both of these derivations neglect temporal Doppler during the pulse transmission[5] and treat only the along-track shift between transmission and reception. In addition, both models are derived for a fixed range. However, the fixed range derivation (centre of patch approximation) is not likely to be a hindrance to correction. The error is small compared to a resolution cell (for most SAS systems) and narrow-band correction may be applied to each range independently.

The effect of image skew may be modelled starting with the monostatic system model (3.23) and including the time varying position, u',

$$e_{\text{moving}}(t, u) \approx \iint \frac{f(x, y) p\Big(t - \frac{2}{c}\sqrt{x^2 + (u' - y)^2}\Big)}{2\sqrt{x^2 + (u' - y)^2}} \, \mathrm{d}x \, \mathrm{d}y, \qquad (3.43)$$

where under a phase-centre approximation

$$u' = u + \frac{v_s t}{2} \qquad (3.44)$$

for a given sonar velocity v_s. Using Fourier skew properties (see Appendix G) the wavenumber domain representation of (3.43) is given by

$$E_{\text{moving}}(k, k_u) = E_{\text{hop}}\Big(k - \frac{v_s}{2c}\, k_u, k_u\Big), \qquad (3.45)$$

$$E_{\text{moving}}(\omega, k_u) = E_{\text{hop}}(\omega - v_s/2\, k_u, k_u), \qquad (3.46)$$

which is a straightforward Doppler shifting of the received echo signal.

Image skew removal is straightforward with minor alterations to the reconstruction algorithms (e.g., in the wavenumber algorithm a preliminary interpolation in along-track). The principle behind skew removal is identical to that used for yaw compensation (see Chapter 6).

Avoiding stop-and-hop modelling

It is possible to calculate the overall continuous movement model and include both effects discussed in the previous sections. The instantaneous time-delay of the transmitted signal for the round-trip propagation may be derived considering the movement

---

[5]Temporal Doppler has only a minor influence [Hawkins 1996]. Despite this, it is possible to correct during reconstruction (see Section 3.5.1).

during the sonar pulse [Hayes 1989]. This allows a better system model than the earlier models. The model for the received signal from a *continuously moving* SAS is given by

$$e_{\text{moving}}(t, u) \approx$$
$$\iint \frac{f(x,y)p\Big(t - \frac{2}{(1-a^2)c}\sqrt{x^2 + (v_s t + u - y)^2} - \frac{2a}{c(1-a^2)}(v_s t + u - y)\Big)}{2\sqrt{x^2 + (v_s t + u - y)^2}}\, \mathrm{d}x\,\mathrm{d}y, \quad (3.47)$$

where $a = v_s/c$. Using Fourier scaling and shifting properties, the wavenumber domain echo signal for a continuously moving sonar can be calculated from (3.47) as

$$E_{\text{moving}}(\omega, k_u) = P(\omega)\frac{f(k_x, k_y)}{k_x}, \quad (3.48)$$

where (k_x, k_y) are given by

$$k_x = \sqrt{\left(\frac{2(k - ak_u)}{(1 - a^2)}\right)^2 - \left(k_u + \frac{2a(k - ak_u)}{-(1 - a^2)}\right)^2}, \quad (3.49)$$

$$k_y = k_u. \quad (3.50)$$

When the continuous sonar movement between pulses is ignored—by setting the inter-pulse sonar velocity v_s to zero—(3.49) collapses to (3.25), that used in the familiar stop-and-hop approximation model of (3.24).

Modelling a continuously moving sonar using (3.48) allows reconstruction algorithms to compensate these effects. Whilst the compensation is not particularly important for the imagery itself, autofocus techniques derive benefit due to their reliance on accurate phase information.

3.5.2   Motion effects

Platform motion, both measured and unmeasured, can cause severe degradation in the resulting SAS imagery [Cutrona 1975; Johnson et al. 1995]. The positions of the transmitter and hydrophones must be known to better than $\lambda/16$ to prevent this[6]. Accuracy in positioning of that order, when the wavelength is typically 10-20 mm, is extremely difficult (to the point of being practically impossible for free-towed or autonomous systems). The topic is dealt with in detail in later chapters.

3.5.3   Medium coherence

The coherence of the medium, both temporal and spatial, is an important parameter in SAS operation. If the medium is not homogeneous and causes variations in the

---

[6]It is even tighter than this for rapidly varying position variations, see Section 7.3.

wave-speed, models derived using the free-space Green's functions fail to be accurate [Jensen et al. 1993]. When the wave-speed variations cause phase fluctuation of the order of $\pi/4$, distortion begins to occur in SAS imagery.

Medium turbulence was initially thought to make SAS unworkable. Several studies in the early stages of SAS research investigated the temporal phase coherence of the medium and found it to be better than expected, with Williams [1976] suggesting that the phase coherence was sufficient for synthetic aperture processing.

Studies by Christoff et al. [1982] and Hayes and Gough [Gough and Hayes 1989a, b; Hayes 1989], showed temporal phase stability to within the $\pi/4$ limit over the period of a minute. Although these results were presented for short range coherence tests (50 m and 130 m respectively), they show that Synthetic Aperture processing was possible underwater. Since those early times, many SAS images have been produced at short range and a smaller number at long range [Châtillon et al. 1999, 1992; Marx et al. 2000].

Imaging platform stability represents the limit on short range SAS imagery with long range imagery limited by medium coherence. The long range limit is attributable to the dual effects of the longer integration times required for targets at long range, and the medium's spatial coherence properties that limit the integration length [Chang and Tinkle 2001]. Sufficiently advanced autofocus techniques are able to ameliorate both medium phase variation and platform movement (see Chapter 7).

3.6   Summary

Accurate modelling of the SAS process is required for accurate reconstruction using the algorithms discussed in Chapters 4 and 5. System models for both multiple-receiver and single-receiver SAS systems in a ground-plane geometry have been presented. These models have amplitude modelling improvements over the standard models due to the use of Weyl's identity (see Appendix B). The relationship between multiple-receiver and single-receiver SAS under the phase-centre approximation has been explored and a new wavenumber domain interpretation of the phase-centre approximation has been presented. An extension of the single-receiver SAS model to 3-D space was discussed to aid in interferometric reconstruction; this was shown to be equivalent to the ground-plane model in the case of ground-plane geometry. Common modelling approximations and errors were discussed and improved models that account for the errors have been derived. A wavenumber-domain model representing the full effect of both the stop-and-hop approximation and image skew has been presented and does not suffer the swath-width limitations of previous derivations.

Chapter 4

Image reconstruction techniques

The reconstruction techniques presented in this chapter attempt to estimate the sea-floor scatterer distribution from the data measured by an active sonar. The reconstruction process gives the resolution improvement that synthetic aperture systems are capable of. The goal is to invert the sonar system model (discussed in Chapter 3) and to produce an image to the desired accuracy as efficiently as possible.

Originally, SAS systems used the time-delay and sum technique, known as *time-domain correlation* (Section 4.2) or *exact matched-filtering* [Gough and Hawkins 1997], common in standard sonar beam-forming [Gough and Hawkins 1997; Hawkins 1996; Hayes 1989; Nielsen 1991; Urick 1975]. The technique allows reconstruction of general, arbitrary geometry imaging problems and is widely used. The major disadvantage of time domain correlation is a large computational load [Gough and Hawkins 1997].

Spatial-frequency domain methods such as *fast correlation* (Section 4.4), *range-Doppler* (Section 4.5), and the *wavenumber algorithm* (Chapter 5), provide algorithmic efficiency and reduce the computational load compared with time-domain correlations. These methods make more stringent assumptions about the collection geometry but have the large computational savings of the FFT algorithm. Chapter 6 covers motion compensation methods for extending the scenarios where frequency/wavenumber domain methods are applicable.

This chapter discusses monostatic (single-receiver) SAS reconstruction algorithms. Spatial domain algorithms (such as time-domain correlation) allow straightforward reconstruction of multiple-receiver sonar data. Algorithms requiring an along-track Fourier transform (such as range-Doppler and wavenumber), need modification for reconstruction of multiple-receiver SAS data (bistatic collection geometry); these modifications are discussed in Section 5.2.2).

4.1   System model

The algorithms summarised in this chapter use a simplified system model. The model
is derived starting with the 3-D system model (see (3.27))

$$E(\omega, u, h) \approx \iiint \frac{f(x, y, z)P(\omega)}{(4\pi)^2 \left|x^2 + (u - y)^2 + (h - z)^2\right|} \\ \exp\left(-j2k\sqrt{x^2 + (u - y)^2 + (h - z)^2}\right) \mathrm{d}z\, \mathrm{d}x\, \mathrm{d}y. \quad (4.1)$$

The system model used is derived by taking the ground plane of the 3-D system model
at $h = z = 0$ (common in the SAS community). Under these conditions the model
becomes

$$E(\omega, u) \approx \iint \frac{f(x, y)P(\omega)}{(4\pi)^2 \left|x^2 + (u - y)^2\right|} \exp\left(-j2k\sqrt{x^2 + (u - y)^2}\right) \mathrm{d}x\, \mathrm{d}y, \quad (4.2)$$

and the image is equivalent to the ground plane image—i.e., $x_s = x$, $f(x, y) = f_s(x_s, y)$.
In other sections of this thesis the 4π constant scale factor is neglected for simplicity.

4.1.1   Fourier algorithm system model

The Fourier domain algorithms (block algorithms) described later in this chapter and in
Chapter 5 are derived using a different system model to that given above in (4.2). This
is done so that the Fourier algorithms bear resemblance to those published previously
in [Hawkins 1996] and in [Soumekh 1994]. The model used is the exploding sources
model discussed in Section 3.3 and is summarised below

$$E(\omega, u) \approx P(\omega) \iint \frac{-f(x, y)}{\sqrt{j16\pi k\sqrt{x^2 + (u - y)^2}}} \exp\left(-j2k\sqrt{x^2 + (u - y)^2}\right) \mathrm{d}x\, \mathrm{d}y. \quad (4.3)$$

The amplitude terms of (4.3) are different from those of (4.2), these differences and the
TVG terms necessary to allow amplitude compensation for Fourier domain reconstruc-
tion as detailed in Appendix C.

4.1.2   Amplitude compensation / Time varying gain (TVG)

In the discussion of the following algorithms, amplitude correction terms accounting
for the spreading loss of the system have been derived. For these to be appropriate,
a range varying gain (RVG) needs to be applied to the imagery before processing to
account for the range spreading loss.

   Imagery without amplitude correction shows a $1/r^2$ decrease in signal amplitude

with range[1]. Spreading losses are typically corrected in the raw data with time or range varying gain (TVG/RVG) being applied either in post-collection processing or in the preamplifier hardware.

In typical side-scan sonar systems RVG is applied to the raw echo image

$$e_{\text{tvg}}(t, u) = e(t, u)r^2, \tag{4.4}$$

where

$$r = \sqrt{x^2 + (u-y)^2} \approx \frac{c\,t}{2}, \tag{4.5}$$

the approximation holding for short-duration transmitted pulses. For historical reasons (and due to the approximate equivalence of range and time-delay (4.5)) the RVG is often termed *Time Varying Gain* (TVG) instead[2].

For the operation of SAS reconstruction algorithms an applied TVG of

$$s_{\text{tvg}}(t, u) = s(t, u)2\pi r, \tag{4.6}$$

is assumed, where $s(t, u)$ is the pulse-compressed echo signal. After the application of TVG, the system model described by (4.2) is

$$S_{\text{tvg}}(\omega, u) \approx \iint \frac{f(x, y)\,|P(\omega)|^2}{8\pi\sqrt{x^2 + (u-y)^2}} \exp\left(-j2k\sqrt{x^2 + (u-y)^2}\right) \mathrm{d}x\,\mathrm{d}y. \tag{4.7}$$

which is the signal used in the time-domain algorithms described in the remainder of the chapter[3]. The reconstruction itself provides some additional gain, removing the rest of the r^2 amplitude dependence.

Using TVG also allows improved amplitude accuracy in Fourier-based reconstruction algorithms (following the method of Appendix C). The TVG operation is optional and depends on the degree of inverse filtering desired. The obvious disadvantage of implementing TVG is the noise-power in the final image estimate becomes range variant.

4.2 Time-domain correlation

Time-domain correlation is a reconstruction algorithm that has been in use since the early SAS experiments and is the same algorithm used in towed array beam-former systems [Nielsen 1991]. The algorithm performs a matched-filtering for the SAS system

---

[1]True in deep water. In shallow water the spreading loss is approximately $1/r$ due to the waveguide properties of the medium. The initial return from a target decreases as $1/r^2$ but subsequent multi-path echoes contain additional energy. The clutter background appears to decline as $1/r$.

[2]Traditionally TVG was implemented in the receiver analogue electronics—a true time varying gain. Time varying gains also alter the Frequency distribution of linear FM chirp signals so should be employed after pulse-compression (pulse-compression is needed by the short-time pulse assumption).

[3]TVG for Fourier-based algorithms is the same but additional post-reconstruction steps need to be taken, this is discussed in Appendix C.

model (3.21) [Bonifant 1999; Pat 2000]. The inverse temporal Fourier transform of the raw echo data gives the time domain representation of the echo signal as

$$e(t, u) \approx \iint \frac{f(x, y)}{8\pi\sqrt{x^2 + (u-y)^2}} \, p\left(t - \frac{2}{c}\sqrt{x^2 + (u-y)^2}\right) \mathrm{d}x \, \mathrm{d}y, \qquad (4.8)$$

where $p(t)$ is the transmitted signal, and $e(t, u)$ the echo signal, neglecting the diffraction limiting effects of beam-patterning. The image estimate, $\widehat{f}(x, y)$, reconstructed using the matched-filter for the echo signal, $e(t, u)$, is then given by [Cook et al. 2001]

$$\widehat{f}(x, y) = \iint \frac{e(t, u)}{8\pi\sqrt{x^2 + (u-y)^2}} \, p^*\left(t - \frac{2}{c}\sqrt{x^2 + (u-y)^2}\right) \mathrm{d}u \, \mathrm{d}t, \qquad (4.9)$$

which may also be written as

$$\widehat{f}(x, y) = \iint \frac{E(\omega, u)}{8\pi\sqrt{x^2 + (u-y)^2}} \, P^*(\omega) \exp\left(j2k\sqrt{x^2 + (u-y)^2}\right) \mathrm{d}u \, \mathrm{d}\omega, \qquad (4.10)$$

using simple Fourier transform properties. Often the amplitude term,

$$\frac{1}{8\pi\sqrt{x^2 + (u-y)^2}}, \qquad (4.11)$$

is neglected, leading to a *phase-only* correlation. An *inverse filter* results if the amplitude term

$$8\pi\sqrt{x^2 + (u-y)^2}, \qquad (4.12)$$

is used instead of (4.11). The remainder of the algorithms discussed in this chapter are presented as inverse filters for range effects. (However, the effect of beam-patterning is neglected.)

The reconstruction described by (4.9) can be computationally very intensive; full integrals over all time samples and sonar pulses[4] need to be computed to get a single output image pixel. Thus the algorithm has a computational efficiency of $\mathcal{O}\left(P^2 N^2\right)$ ($\mathcal{O}\left(H^2 P^2 N^2\right)$ for multiple-receiver systems). The benefit of using time-domain correlation is that it is flexible and is easily adapted to account for platform motion and Doppler effects by using an appropriate system model.

Previous SAS systems, such as SAMI [Adams et al. 1993, 1996], the original KiwiSAS [Hayes 1989; Hayes and Gough 1992], and the sonar used by Douglas and Lee [Douglas and Lee 1992, 1993a, b; Silkaitis et al. 1994], have been research instruments where computational effort was not a major concern. Thus the flexibility of time-domain correlation makes it a good algorithm choice for research groups. Commercial SAS systems[5] require a reconstruction with low computational burden, strongly sug-

---

[4] The integral over u is limited by the number of pulses that see the particular output pixel.

[5] Autofocus or other processing using iterative algorithms provide another strong motivation for Fourier-based reconstruction algorithms.

gesting the use of the faster, Fourier-based, algorithms.

4.3 Back projection

Back projection is a reconstruction technique used extensively in the medical imaging fields [Cho et al. 1993; Liang and Lauterbur 2000] for the similar reconstruction problems encountered in computed tomography. Back projection algorithms take the received signal for a given pulse and *back-project* that signal over a spherical arc to all the possible contributing image points. Once back projection is performed on the remainder of the pulses and accumulated, an estimate of the image is obtained. A minor modification to the algorithm, called *Filtered back projection* (sometimes known as *modified back projection* [Enright 1992]), is often employed to produce images more suitable for human interpretation [Liang and Lauterbur 2000]. Several differing filters may be employed but are not in common use in the SAS community.

In the SAS scenario, a back projection reconstruction may be summarised by

$$\widehat{f}(x,y) = \int 8\pi \sqrt{x^2 + (u-y)^2}\, s\left(\frac{2}{c}\sqrt{x^2 + (u-y)^2}, u\right) \mathrm{d}u, \qquad (4.13)$$

where

$$s(t,u) = e(t,u) \odot_t p^*(-t) \equiv e(t,u) \star_t p(t), \qquad (4.14)$$

is the pulse-compressed image data. The comments of Section 4.2 regarding the amplitude term $8\pi \sqrt{x^2 + (u-y)^2}$ apply.

It is clear that since t is a discrete variable, obtaining $s\left(2/c\,\sqrt{x^2 + (u-y)^2}, u\right)$ from $s(t,u)$ requires interpolation. If Fourier-based interpolation is used (4.13) is the same as the time-domain correlation inversion of (4.10). Rewriting (4.13) using Fourier-based interpolation gives

$$\widehat{f}(x,y) = \int 8\pi \sqrt{x^2 + (u-y)^2}\left(s\left(\frac{2}{c}\sqrt{x^2 + (u-y)^2}, u\right)\right) \mathrm{d}u, \qquad (4.15)$$

$$= \int 8\pi \sqrt{x^2 + (u-y)^2}\left(\int S(\omega, u)\exp\left(j2k\sqrt{x^2 + (u-y)^2}\right) \mathrm{d}\omega\right) \mathrm{d}u, \quad (4.16)$$

$$= \int 8\pi \sqrt{x^2 + (u-y)^2}\left(\int E(\omega, u)P^*(\omega)\exp\left(j2k\sqrt{x^2 + (u-y)^2}\right) \mathrm{d}\omega\right) \mathrm{d}u, \qquad (4.17)$$

which is identical to (4.10). Back projection using Fourier interpolation has no processing efficiency gain over time-domain correlation and the algorithms are equivalent.

Other, less computationally-intensive interpolation methods can be used with a resulting loss of accuracy. Soumekh [1999, pp 214–215] notes that inaccuracies in the interpolation lead to a loss of high resolution information over the entire image.

This contrasts with spatial-frequency interpolation algorithms, such as the wavenumber algorithm, where information loss due to interpolation inaccuracies occurs at the edges of the reconstructed image.

Motion compensation and multiple-receiver geometries are easily incorporated into the back projection algorithm in the same way as in the equivalent time-domain correlation algorithm. A straightforward alteration of the range term in (4.13) can account for any spatial variation encountered.

Fourier back projection [Lockwood et al. 2001] is closely related to the fast correlation algorithm described later in this chapter. The details of the algorithm are not described in this thesis.

Back projection provides a reconstruction technique with very similar properties to those of time-domain correlation. Computational savings over a traditional time-domain correlation are possible depending on the accuracy of the interpolation used—in general it is not necessary to use a full Fourier interpolation (e.g., the interpolation described by Shippey et al. [2001], or the spatial-frequency beamformer described by Groen et al. [2001]). The computational complexity of the back projection algorithm is $\mathcal{O}(P^2N)$ although this depends on the type of interpolator used ($\mathcal{O}(H^2P^2N)$ for a multiple-hydrophone system) [Ulander et al. 2001]. Back projection allows the flexibility of time-domain-correlation and, with suitable interpolation, provides significant computational savings.

4.3.1   Fast factorised back projection (FFBP)

Recent modifications have been made to the back projection based algorithms to enable reconstruction approaching speeds of the Fourier based schemes described later in the chapter. The most recent algorithm (known as *fast factorised back projection*), partitions both echo data and image space recursively and allows redundancy in the image reconstruction processing to be exploited [Ulander et al. 2001]. The benefit of the partitioning is an algorithm with tunable performance, with a best performance of the order[6] $\mathcal{O}(P \log PN)$, where P is the number of pulses and N the number of time samples (as opposed to the $\mathcal{O}(P^2N)$ performance of direct back projection) [Ulander et al. 2000, 2001; Xiao et al. 2000].

Any processing gain over traditional back projection is obtained by the use of depth-of-focus approximations in the recursive partitioning of the reconstruction problem. Because of this, a so-called exact reconstruction has the same $\mathcal{O}(P^2N)$ performance [Banks and Griffiths 2002]. Clear performance gains over traditional back projection have been obtained on SAS reconstruction tasks where performance gains

---

[6]The order is between $\mathcal{O}(H^2P \log PN)$ and $\mathcal{O}(HP \log(HP)N)$ for multiple-receiver systems depending on the phase centre interpolation scheme used.

of approximately two orders of magnitude have been reported for realistic SAS scenes [Banks and Griffiths 2002; Shippey et al. 2001].

Multiple-receiver data reconstruction is straightforward but sacrifices some computational efficiency. To give computational savings, the multiple-receiver correction is only applied at the first recursive stage; thereafter, the system is treated using the phase centre approximation (see Section 3.2.2) [Banks and Griffiths 2002]. The approximation results in phase centre correction for only the centre of the scene and causes minor blurring in the other parts of the image.

Fast factorised back projection is a promising algorithm for SAS image reconstruction. Potentially, the algorithm offers the flexibility of traditional back projection/time-domain correlation with the benefits of improved computational efficiency. It remains to be seen how the computational efficiency compares with that of Fourier-based reconstruction for real-world imaging situations. The recursive partitioning used in FFBP can be applied to any of the SAS reconstruction techniques including Fourier-based algorithms such as the wavenumber algorithm. The recursive partitioning is somewhat similar to the depth-of-focus partitioning used with fast correlation (see following section).

4.4 Fast correlation

Fast correlation is a method of implementing the correlation of an equivalent time-domain correlation reconstruction (see Section 4.2) in the wavenumber (spatial-frequency) domain. The algorithm uses Fourier domain matched-filtering to implement the convolution described by (4.9) [Gough and Hawkins 1997]. *Convolution beamforming* [Groen and Sabel 2002; Groen et al. 2001], and *range-stacking* [Soumekh 1999, pp 206–212], are other names used for implementations of the algorithm.

The algorithm is straightforward to derive using Weyl's identity (see Appendix B) to take an along-track Fourier transform of the system model and calculate the correlation of (4.10) in the wavenumber domain. Taking an along-track Fourier transform of (4.10) gives the fast correlation implementation (neglecting some of the amplitude effects) as

$$\widehat{f}(x,y) = \iint \sqrt{4k^2 - k_u^2}\, S(\omega, k_u) \exp\left(j\,|x_0|\,\sqrt{4k^2 - k_u^2} + jk_u y\right) \mathrm{d}k_u\, \mathrm{d}\omega. \qquad (4.18)$$

Fourier Domain filtering is only locally correct around the centre range of the scene, x_0 because of the space-variant nature of the SAS imaging problem [Hayes 1989; Hayes and Gough 1992]. This leads to depth-of-focus issues, where the reconstruction is only valid for small sub-regions of the reconstructed image [Groen and Sabel 2002; Hawkins 1996; Hayes 1989; Hayes and Gough 1992]. Images in such cases are partitioned into

regions smaller than the depth-of-focus; these images are processed separately then used in a mosaic.

The processing cost for the algorithm is $\mathcal{O}(P \log PN \log N)$ if the entire image is inside the depth-of-focus [Groen and Sabel 2002]. As most of the scene is outside the depth-of-focus, the efficiency is somewhere between $\mathcal{O}(P^2 N \log N)$ and $\mathcal{O}(P \log PN \log N)$ depending on what the depth-of-focus is for the particular system.

The reconstruction may be summarised as

$$\widehat{f}(x, y) = \mathcal{F}^{-1}_{\omega, k_u} \left\{ \sqrt{4k^2 - k_u^2}\ S(\omega, k_u) \exp\left(j\,|x_0|\,\sqrt{4k^2 - k_u^2}\right) \right\}, \qquad (4.19)$$

where $S(\omega, k_u)$ is the 2-D Fourier transform of the pulse compressed echo image $s(t, u)$,

$$S(\omega, k_u) \equiv E(\omega, k_u) P^*(\omega), \qquad (4.20)$$

x_0 is the focus depth, and

$$k_y \equiv k_u. \qquad (4.21)$$

The reconstruction algorithms described in subsequent sections (such as the range-Doppler and wavenumber algorithms) implement forms of fast correlation with differing methods for extending the depth-of-focus.

4.5   Range-Doppler algorithm

The range-Doppler algorithm is a popular reconstruction algorithm in the satellite imaging field [Bamler 1992; Carrera et al. 1995; Curlander and McDonough 1996], and has been used extensively since its invention in 1982. The range-Doppler algorithm performs fast correlation in the along-track direction (Doppler matched-filtering) and utilises time-domain interpolation to extend the depth-of-focus.

The range-Doppler algorithm starts by taking a 1-D Fourier transform of the pulse-compressed echo signal, $s(t, u)$, in the along-track dimension u to give the along-track wavenumber k_u. This gives the range-Doppler domain representation of the pulse-compressed echo signal. Performing the Fourier transform on the modelled echo signal (3.21) gives the range-Doppler representation of the pulse-compressed echo signal (ignoring diffraction limiting effects) [Bamler 1992; Hawkins 1996; Raney 1992]

$$s(t, k_u)|_{x_0} \approx \frac{f(x_0, k_u)}{\sqrt{4k_0^2 - k_u^2}}\ \delta\left(t - \frac{2}{c}\Delta R_s(x_0, k_u)\right) \exp\left(j\,|x_0|\,\sqrt{4k_0^2 - k_u^2}\right), \qquad (4.22)$$

where $k_u \equiv k_y$ (see Chapter 3), k_0 is the sonar centre frequency, the *range-migration* for a target at x is given by

$$\Delta R_s(x_0, k_u) = x_0 C_s(k_u), \qquad (4.23)$$

and the *curvature factor*, $C_s(k_u)$, is

$$C_s(k_u) = \frac{1}{\sqrt{1 - (\frac{k_u}{2k_0})^2}} - 1. \tag{4.24}$$

The first step of the range-Doppler algorithm is to perform a coordinate transformation decoupling the range and across track variables:

$$s_{\text{rmc}}(x, k_y)|_{x_0} = \mathcal{T}\{s(t, k_u)|_{x_0}\}, \tag{4.25}$$

where the coordinate transform $\mathcal{T}\{\}$ is given by,

$$x \equiv \frac{c}{2}(x_0 - \Delta R_s(x_0, k_u)), \tag{4.26}$$

$$k_y \equiv k_u. \tag{4.27}$$

This has the effect of straightening the range-migration of (4.23) and is implemented as a range-Doppler domain interpolation.

Once the range-migration-correcting coordinate-transform has been applied for all received ranges and for all k_u, the resulting image is matched-filtered with a narrow-band propagation filter

$$\widehat{f}(x, k_y) = s_{\text{rmc}}(x, k_y)q(x, k_y), \tag{4.28}$$

where

$$q(x, k_y) = \sqrt{4k_0^2 - k_y^2} \, \exp\left(j \, |x| \, \sqrt{4k_0^2 - k_y^2}\right). \tag{4.29}$$

After phase filtering the image estimate is obtained by inverse Fourier transformation. the range-Doppler reconstruction may be summarised by

$$\widehat{f}(x, y) = \mathcal{F}_{k_y}^{-1}\{q(x, k_y)\mathcal{T}\{s(t, k_u)\}\}. \tag{4.30}$$

The range-Doppler algorithm is the standard reconstruction for a large number of SAR systems [Carrera et al. 1995]. Although in common use, it is computationally more expensive than the chirp-scaling (see Section 4.6) for similar performance and requires SRC to match the accuracy of the wavenumber algorithm (which is also less computationally expensive). In SAS imaging either chirp-scaling or the wavenumber algorithm provide better solutions to the reconstruction problem.

4.5.1 Secondary range compression (SRC)

The use of the range-Doppler algorithm produces image degradation in wide-beam systems. This results from a narrow band assumption ($k \approx k_0$) made in the derivation of the algorithm. The degradation is particularly noticeable in systems that employ large

Image Data

(t,u)

Azimuth FFT

(t,ku)

$\otimes$ —— φ1
Chirp scaling

Range FFT

(ω,ku)

$\otimes$ —— φ2
Bulk range correction,
pulse compression

Range IFFT

(t,ku)

$\otimes$ —— φ3
Azimuth focus,
phase compensation

Azimuth IFFT

(t,u)

Focused data
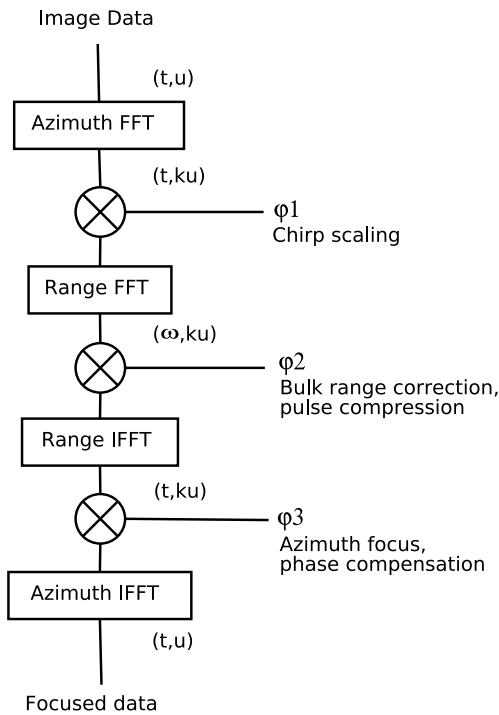
Figure 4.1   Chirp-scaling algorithm operation.

beam-widths or have low Q-factors (most SAS systems). Secondary range compression (SRC) was developed to account for these issues [Curlander and McDonough 1996].

The narrow-band approximation neglects spreading in the range direction of the range-migration corrected data. This spreading is noticeable for non-zero k_u values. SRC provides a k_u dependent compression in the range direction to give improved imagery.

SRC is likely to be needed on any range-Doppler based image reconstruction for SAS systems since most systems under development have spatial bandwidths larger than those of the 1996 KiwiSAS (which required SRC [Hawkins 1996]).

4.6   Chirp-scaling algorithm

The chirp-scaling algorithm is a Fourier domain algorithm similar in derivation to the range-Doppler algorithm. In 1992 two groups independently presented the algorithm as an attempt to remove the computationally intensive interpolation of the range-Doppler algorithm [Cumming et al. 1992; Runge and Bamler 1992]. The chirp-scaling algorithm avoids the interpolation of the range Doppler algorithm by using only the first three[7] terms in the Taylor expansion of the Stolt transformation [Raney et al. 1994] (used in wavenumber reconstruction see Chapter 5). With only the first three terms used,

[7]Gimeno and Lopez-Sanchez [2001] outlines a method for allowing higher order Taylor series terms to be used.

the time-shifting and scaling properties of linear-FM chirps are exploited to replace the interpolation. An in-depth analysis of the chirp-scaling algorithm is presented in Hawkins [1996] and Hawkins and Gough [1997a]—the following description is a summary of those works in the notation of this thesis.

The chirp-scaling algorithm starts by taking an along-track Fourier transform of the *raw* echo data—i.e., the data is not pulse compressed unlike the range-Doppler algorithm. Note that the transmitted signal must be a linear-FM chirp of the form

$$p(t) = \exp\left(j\pi K t^2\right). \tag{4.31}$$

After the along-track Fourier transform, the data may be modelled as

$$e(t, k_u)|_x \approx \frac{f(x, k_u)}{\sqrt{4k_0^2 - k_u^2}} \operatorname{rect}\left(\frac{t - \frac{2}{c}\Delta R_s(x, k_u)}{\tau_c}\right)$$

$$\exp\left(jK_s\left(t - \frac{2}{c}\Delta R_s(x, k_u)\right)^2\right) \exp\left(j\,|x|\,\sqrt{4k_0^2 - k_u^2}\right), \tag{4.32}$$

where

$$K_s(k_u)|_x = \frac{1}{1/K - K_{\mathrm{src}}(k_u)} \tag{4.33}$$

is the received chirp rate and

$$K_{\mathrm{src}}(k_u)|_x = \frac{8\pi x}{c} \frac{k_u^2}{\left(4k_0^2 - ku^2\right)^{3/2}} \tag{4.34}$$

the chirp rate due to imaging geometry. $\Delta R_s(x, k_u)$ and $C_s(k_u)$ are as for the range-Doppler algorithm.

The *chirp-scaling* phase multiply is applied to the range Doppler data via

$$m(t, k_u)|_x = e(t, k_u)\varphi_1(t, k_u)|_x, \tag{4.35}$$

where the phase multiply is given by

$$\varphi_1(t, k_u)|_x = \exp\left(j\pi K_s(k_u)C_s(k_u)(t - t_0(k_u))^2\right) \tag{4.36}$$

and the time shift to the reference locus

$$t_0(k_u) = \frac{2}{c}x_0(1 + C_s(k_u)). \tag{4.37}$$

This step removes the range dependence of the *phase* of the range-Doppler signal (but not the range dependence of the signal envelope). Afterwards, targets at all ranges have the same phase signal as a target at the reference range x_0.

A range Fourier transform of the range-Doppler image is taken to allow bulk cur-

vature compensation (matched-filtering for the reference range).

$$M(\omega, k_u) = \mathcal{F}_t\{m(t, k_u)\}. \tag{4.38}$$

Pulse-compression, SRC, and bulk curvature compensation are all applied with another phase multiply via[8]

$$N(\omega, k_u) = \sqrt{4k^2 - k_u^2}M(\omega, k_u)\varphi_2(\omega, k_u), \tag{4.39}$$

where

$$\varphi_2(\omega, k_u) = \exp\left(j\left(\frac{c^2k^2}{4\pi K_s(k_u)(1 + C_s(k_u))}\right)\right)\exp\left(j2kx_0C_s(k_u)\right). \tag{4.40}$$

After this multiply, the signal azimuth and range coordinates have been decoupled and

$$N(k_x, k_y) \equiv \mathcal{C}\{N(\omega, k_u)\}, \tag{4.41}$$

where the coordinate transform $\mathcal{C}\{\}$ is given by the trivial mapping

$$k_x = 2k, \tag{4.42}$$

$$k_y = k_u. \tag{4.43}$$

An inverse Fourier transform is applied to obtain the decoupled range Doppler (image Doppler instead of signal Doppler) image. Once in this domain, the azimuth compression used in the range-Doppler algorithm is applied along with additional phase-compensation to obtain the focused range-Doppler image. This is performed via

$$\widehat{f}(x, k_y) = N(x, k_y)\varphi_3(x, k_y), \tag{4.44}$$

where the final phase multiply is given by

$$\varphi_3(x, k_y) = q(x, k_y)\exp\left(-j\frac{4\pi}{c^2}K_s(k_y)C_s(k_y)(1 + C_s(k_y))(x - x_0)^2\right), \tag{4.45}$$

where $q(x, k_y)$ is as for the range-Doppler algorithm, (4.29).

The final image estimate is now obtained with an azimuth inverse Fourier transform of $\widehat{f}(x, k_y)$.

Chirp-scaling, whilst mathematically complicated, is simple to implement—only phase multiplies and Fourier transforms are required. The intentional lack of an interpolator that allows this simplicity also improves reconstruction efficiency. In typical imaging scenarios, chirp-scaling has been found to be about 2 times faster than the

---

[8]Note the altered frequency dependent amplitude term used here ($\sqrt{4k^2 - k_u^2}$). This is a result of the improved modelling of Chapter 3

wavenumber algorithm for equivalent results [Gimeno and Lopez-Sanchez 2001]. However, Bonifant [1999, page 153] raises concerns about the phase performance of the algorithm implying care must be taken with regard to autofocus scenarios. In light of these concerns, the wavenumber algorithm (see Chapter 5) has been chosen over chirp-scaling for the reconstructions used in this thesis.

4.6.1   Accelerated chirp-scaling

Accelerated chirp-scaling is an extension of chirp-scaling that allows computational savings to be made when reconstructing a patch of small extent compared to the chirp duration [Hawkins 1996; Hawkins and Gough 1997b]. By noting that the extent of the linear-FM chirp signal only needs to be as long as the maximum range-migration, ΔR_s, the amount of data needed for the chirp-scaling may be reduced [Hawkins 1996, pp 92–93].

   The data reduction is obtained by reducing the length of the linear-FM chirp in a preliminary pulse-compression-like operation. Once the chirp is reduced in length to that required for the chirp-scaling step, the reconstruction proceeds normally.

   An additional benefit of the technique is that pulse-compressed data may be *re-chirped* (i.e., remodulated to make it appear as if a linear FM signal had been used) and subsequently processed using chirp scaling [Hawkins 1996, page 92];[Gimeno and Lopez-Sanchez 2001]. This allows the extension of the chirp-scaling to arbitrary waveforms. If the waveform is pulse-compressed normally then re-chirped with a linear-FM chirp, it is usable in a chirp-scaling reconstruction processor.

4.7   Post-processing

Reconstructed sonar images are often difficult for human observers to interpret and post-processing is sometimes employed. This section covers some of the commonly-used post-processing techniques used on SAS images.

4.7.1   Speckle-reduction / multi-look imagery

SAS intensity images generally have a grainy appearance caused by speckle-noise. The multiplicative nature of speckle-noise makes interpretation difficult[9] for both automated and human observers.

   A SAS speckle-image (or for that matter any other coherent image of a surface that is rough compared with the imaging wavelength) has a negative-exponential intensity distribution [Goodman 1976]. The highly probably outcome is that the image intensity

---

[9]Speckle-noise appears as a particularly strong multiplicative noise where the variance of a point is identical to its mean.

is zero for any given pixel[10]. To reduce the variance, several overlapping images are sometimes summed on an intensity basis to lower the noise and make the statistics more Gaussian (which improves the ease of image interpretation). The approach trades image resolution for increased accuracy in image amplitude. This is usually termed *multi-look* processing or speckle-reduction processing [Moreira 1991].

Multi-look processing is performed by taking separate sub-bands of the image and summing them on an intensity basis [Carrera et al. 1995; Hawkins 1996; Jakowatz et al. 1996]. In spotlight SAR imagery the multi-look images are often obtained by processing different portions of the along track aperture (called sub-aperture processing). Due to the equivalence of aperture position and spatial frequency (explored later in the thesis— see Chapter 8) sub-aperture processing and sub-band processing are identical. With SAS systems, range sub-bands can be used instead [Chanussot et al. 2002; Hawkins 1996; Hayes and Gough 1992]—it is usually preferable to sacrifice range resolution in SAS systems rather than azimuth resolution. This is unusual compared to airborne SAR systems—they have much higher azimuth resolution than range resolution and so tend to use multi-look in azimuth.

Trahey and Smith [1988] studied the spatial resolution vs. speckle reduction trade-off and concluded that for detecting lesions in ultrasound images (similar to detecting the shadow behind proud targets) full resolution imaging was preferable. Non-linear speckle-reduction techniques [Moreira 1991] hint at the possibility of retaining resolution while reducing the effects of speckle. At this stage the benefits in applying speckle-reduction techniques with SAS data remain unclear. The field-collected data shown elsewhere in this thesis has had no speckle-reduction applied.

4.7.2 Non-coherent processing

Non-coherent processing, a variant of multi-look processing, is sometimes used to reduce the sensitivity of the synthetic aperture processing to motion errors and turbulence. The processing is equivalent to the extreme case of the multi-look reconstructions employed in SAR [de Heering 1982; Hayes 1989; Jakowatz and Wahl 1993].

Non-coherent processing reconstructs a series of images from sub-sections of the available aperture, takes intensities (discarding the phase information), and sums those images. Depending on the length of sub-aperture used, either full non-coherent processing is possible [Douglas 1993] or various orders of multi-look images [Hawkins 1996; Jakowatz and Wahl 1993].

Multi-look processing still requires normal synthetic aperture reconstruction algorithms to be used. It is regarded as an auxiliary post-processing technique. Time-domain correlation is well suited to reconstruction if a full non-coherent image is required. Hayes [1989] covers the topic in more detail.

---

[10]Due to the variance of any given image pixel equalling the intensity of the underlying scene.

4.8 Summary

Image reconstruction is required to realise the resolution improvement SAS systems offer. This chapter has presented summaries of all of the common reconstruction algorithms (except wavenumber reconstruction which is described separately in Chapter 5) within a common framework and with consistent notation. In addition, the reconstruction algorithms have been amended to include the improved modelling in Chapter 3. This allows images to be reconstructed free from amplitude term approximations.

Fourier domain reconstruction offers large computational savings over most time-domain methods. This improvement comes at the expense of less flexibility—Fourier domain methods require separate motion compensation schemes (discussed in Chapter 6). Of the Fourier domain methods, accelerated chirp-scaling with SRC or the wavenumber algorithm should be used as they provide improved reconstruction with a lower computational cost than the other algorithms. Some concern over the phase performance of chirp-scaling make the wavenumber algorithm a better choice where autofocus or interferometry are required. Second order effects with accurate wavenumber domain models (such as those discussed in Section 3.5) are easily compensated during reconstruction using Fourier-based methods.

Of the time-domain methods, fast factorised back projection (FFBP) should be used as it offers improved efficiency and implements traditional back projection as a special case. FFBP promises flexible processing at speed rivalling Fourier based algorithms (the recursive partitioning used is reminiscent of the recursive techniques used in FFT type algorithms). Further research is required to determine how promising FFBP is in relation to combined Fourier reconstruction and motion compensation, particularly with regard to the highly optimised Fourier transform routines found on common computing hardware.

Wavenumber domain processing

Wavenumber reconstruction entered synthetic aperture imaging from the seismic imaging field [Stolt 1978] with the technique first being used in the open SAR literature in the early 1990s [Cafforio et al. 1991a, b; Milman 1993]. The wavenumber algorithm relies on inverting the effect of the imaging system by the use of a coordinate transformation in the spatial-frequency domain (wavenumber domain) [Soumekh 1994]. The technique is often termed wavenumber interpolation because the coordinate transform is implemented using wavenumber domain interpolation.

Reconstruction via Fourier space interpolation is also common practice in other imaging fields and is an accepted inversion technique [Poularikas 1996]. Some usage examples include: radio astronomy, magnetic resonance imaging (MRI), computer aided tomography (CAT) (although the Fourier data collection is markedly different) [Liang and Lauterbur 2000; Poularikas 1996; Stark and Woods 1994], spotlight SAR (polar reformat algorithm) [Jakowatz et al. 1996], and seismic imaging [Soumekh 1994]. Once the positions of the recorded samples in Fourier space are known, the algorithms interpolate those samples onto a regular grid and inverse Fourier transform to reconstruct the image.

The major benefit of using Fourier interpolation methods, aside from computational efficiency, is that second order effects (see Section 3.5) can be compensated with an almost insignificant increase in computation. Accurate compensation of second order effects only requires better wavenumber domain modelling. The majority of wavenumber algorithm computation is in Fourier transforms/interpolation, and only a small amount in calculating the wavenumber domain system model. Jakowatz et al. [1996] use a similar approach where many effects (such as yaw and slant-range collection) are compensated in the polar reformat algorithm (a wavenumber domain interpolation algorithm). This is in direct contrast with algorithms such as time-domain correlation where the kernel calculation weighs heavily on the computation—any extra burden in the kernel causes a large increase in computational cost.

5.1   Wavenumber algorithm for single-receiver systems

The wavenumber algorithm itself has been developed independently by more than one group and has several names including: the seismic migration algorithm [Cafforio et al. 1991a, b], the Stolt mapping algorithm, the range-migration algorithm (RMA) [Curlander and McDonough 1996], and the (ω, k_u) algorithm [Milman 1993].

The traditional wavenumber algorithm operates by 2-D Fourier transforming the pulse-compressed image, $s(t, u)$, into the wavenumber domain (along-track wavenumber/temporal frequency-domain, $S(\omega, k_u)$). This is followed by matched-filtering for a target at the reference range, x_0 (to ease the Fourier interpolation requirements) and a nonlinear coordinate transformation (Stolt mapping) to give

$$\widehat{f}(k_x, k_y) = \mathcal{S}^{-1}\Big\{ \sqrt{4k^2 - k_u^2}\, S(\omega, k_u) \exp\Big(j\,|x_0|\,\sqrt{4k^2 - k_u^2}\Big)\Big\}. \tag{5.1}$$

The coordinate transform that the Stolt mapping, $\mathcal{S}^{-1}\{\}$, describes is given by [Hawkins 1996]

$$k_x = \sqrt{4k^2 - k_u^2}, \tag{5.2}$$

$$k_y = k_u, \tag{5.3}$$

and the inverse Stolt mapping, $\mathcal{S}\{\}$,

$$k_u = k_y, \tag{5.4}$$

$$\omega = \frac{c}{2}\sqrt{k_x^2 + k_y^2}. \tag{5.5}$$

The Stolt coordinate transform is implemented by interpolating input samples, located on spheres (arcs in 2-D) of constant temporal-frequency (ω, k_u) centred around the wavenumber origin, onto a regular grid, (k_x, k_y). Figure 5.1 illustrates the baseband version of the coordinate mapping (discussed in the next section—the major difference in implementation is to shift the wavenumber origin to DC prior to Fourier domain interpolation). A 2-D Inverse Fourier transform from the interpolated wavenumber domain data $\widehat{f}(k_x, k_y)$ is all that is required to get an estimate of the slant-range image $\widehat{f}(x, y)$.

5.1.1   Implementation details

As in all algorithms, the mathematics belies some of the difficulties in actually implementing reconstruction. This section details the major implementation issues.

Figure 5.1 Wavenumber interpolation detail. The wavenumber algorithm takes the input samples samples (black dots), measured on circles of constant temporal-frequency, and interpolates to get a regular grid of samples for output. The interpolation shown is the baseband Stolt mapping described by (5.8). A rectangular region is normally extracted to ensure a consistent point spread function in the reconstructed image [Hawkins 1996].

Baseband mapping

The nonlinear Stolt coordinate transformation of (5.2) and (5.3) requires an accurate interpolation to ensure errors are not injected into the image. The algorithm uses baseband data both for input and output and to remove the rapidly varying phase functions. In addition, the algorithm is formulated for operating on data about a range offset r_0 and returning data offset about x_0 (see Section 3.2.1—the output image is a function of x' where $x' = x - x_0$). With the modulations and offsets taken into account, the reconstruction becomes

$$\widehat{f}(k_{xb}, k_y) = \mathcal{S}_b^{-1}\left\{ \sqrt{4(k_b + k_0)^2 - k_u^2} \exp\left(j(k_{xb} + k_{x0})x_0 - j2(k_b + k_0)r_0\right)S(\omega_b, k_u)\right\}, \tag{5.6}$$

with the baseband Stolt mapping, $\mathcal{S}_b^{-1}\{\}$, given by

$$\begin{aligned} k_{xb} &= \sqrt{4(k_b + k_0)^2 - k_u^2} - k_{x0}, \\ k_y &= k_u, \end{aligned} \tag{5.7}$$

where k_b is the baseband temporal-frequency around a carrier k_0, k_{xb} is the baseband across-track spatial-frequency around a carrier k_{x0}, r_0 is the distance to the input image centre, and x_0 is the distance to the output image centre. In normal operation $x_0 \equiv r_0$,

and $k_{x0} \equiv 2k_0$, thus the baseband mapping (5.6) simplifies to

$$\widehat{f}(k_{xb}, k_y) = \mathcal{S}_b^{-1}\left\{ \sqrt{4(k_b + k_0)^2 - k_u^2} \exp\left(j(k_{xb} - 2k_b)x_0\right)S(\omega_b, k_u)\right\}. \qquad (5.8)$$

To produce a non-baseband image (as the algorithms covered in Chapter 4 produce), the reconstruction is modified to be

$$\widehat{f}(x, y) =$$
$$\mathcal{F}_{k_{xb},\ k_y}^{-1}\left\{ \mathcal{S}_b^{-1}\left\{ \sqrt{4(k_b + k_0)^2 - k_u^2} \exp\left(j(k_{xb} - 2k_b)x_0)\right)S(\omega_b, k_u)\right\}\right\} \exp\left(jk_{x0}x\right). \quad (5.9)$$

This step has no impact if only the magnitude of the image, $\left|\widehat{f}(x, y)\right|$, is required[1].

For notational simplicity, the implementation details described by (5.6) and (5.8) are neglected for the remainder of the thesis.

Reconstruction accuracy versus efficiency

Wavenumber reconstruction with ideal interpolators would provide error-free imagery. Given the constraints of operational hardware and a desire to perform the reconstruction as fast as possible, low-order interpolators are often chosen. The order of the interpolator must be traded against the accuracy desired in the final image.

Whatever the order of interpolator chosen, a reconstruction using the wavenumber algorithm is always valid at the reference range, x_0. Both wavenumber reconstruction techniques (via (5.1)) and fast correlation (via (4.19)) use a matched-filter for the reference range. The interpolation in the wavenumber algorithm *extends the focus depth* of fast correlation. Without interpolation, the wavenumber algorithm is *identical* to the fast correlation algorithm (see Section 4.4) [Li 1992]. Thus the image is always well focused at the reference range (even for large spatial bandwidths) and image distortion from inaccurate interpolation only appears at the edges of the image.

Reconstructing the image in sub-swaths by reducing the distance from edge of scene to the centre lessens interpolation constraints. Computational efficiency is gained by splitting a wide swath-width image before reconstruction. Near-range imaging is a harder problem than far-range imaging (non-linear and space-variance) and so requires higher-order interpolation. By splitting the image, high order interpolators only need be applied to the close-range swath (at heavy computational expense) and low-order interpolators can be used for the remainder of the image[2].

Efficient derivatives of the wavenumber algorithm exist if the swath width is small

---

[1]It is preferred if the images are reconstructed as baseband images and the results of the previous chapter's algorithms are often base-banded by multiplying by $\exp\left(-jk_{x0}x\right)$ [Soumekh 1999].

[2]The disadvantage of splitting the image into range sub-swaths as discussed is that the efficiency gained from the use of the FFT to calculate $S(\omega, k_u)$ is reduced—$\mathcal{O}\{N \log N\}$ versus $\mathcal{O}\{N^2\}$ for splitting into N sub-swaths.

compared to the average range and the system is narrow-band (i.e., if wavenumber domain coordinate mapping shifts the samples only a small amount). For these derivative algorithms, the interpolation process in the wavenumber domain is replaced by phase multiplication in the range-Doppler domain in a similar manner to chirp-scaling [Lanari 1994, 1995; Li 1992]. This technique avoids interpolation at the expense of two additional 1-D Fourier transforms per along-track spatial-frequency sample.

A useful technique for determining the required interpolation accuracy is to run the wavenumber algorithm backward[3] on a reconstructed image. The slight differences between the initial pulse-compressed image and the reconstructed image after being run through the backward wavenumber algorithm show the effect of finite precision interpolation. Alternately, reconstruction of simulated data sets with investigation of peak to grating lobe ratios and phase distortion can provide further interpolation accuracy checks.

Quick-look imagery

In recent times there has been a desire to provide low-resolution images[4] with only a minimum of processing. This is one of the proposed benefits of FFBP [Banks and Griffiths 2002] (see Section 4.3.1).

Processing of low-resolution images via the wavenumber algorithm is straightforward. In all synthetic aperture imaging systems, the final image resolution in azimuth is constrained by the physical element size (see Section 2.6.4). Simply summing elements together prior to reconstruction (effectively low-pass filtering and decimating the input image) provides a low-resolution output with significantly reduced computation. An identical procedure in the range direction also provides a reduced cross-track resolution at lower computational effort. This processing method for obtaining quick-look low-resolution imagery is possible using any of the reconstruction algorithms outlined in Chapter 4.

Low-resolution images can also be obtained at various stages in wavenumber algorithm processing by performing the wavenumber domain interpolation in sub-bands. As output along-track spatial frequencies become available, a 2-D inverse Fourier transform will give images progressively increasing in resolution. The obvious cost of requiring progressive low-resolution images is a 2-D inverse Fourier transform for each image required. The requirement is computationally demanding although the Stolt interpolation is often the most computationally expensive part of the wavenumber algorithm.

The application of quick-look imagery as described is limited, with operator display

---

[3]Running the wavenumber algorithm backward is equivalent to simulating the echo given a scene and is straightforward to derive and implement.

[4]The intention is that low-resolution images may be processed quickly for determining areas of interest to later reconstruct at full resolution. Application also exists for some interferometric processing steps.

being the current primary use. For this purpose fixed focus at x_0 or even unfocused imagery may suffice. The long-term prospect is much more interesting with possible real-time, multi-scale, autofocus techniques or region-of-interest processing likely to be attempted on SAS data. Any of the common reconstruction techniques can be modified to perform quick-look imagery.

5.2   Wavenumber reconstruction for multiple-receiver systems

To overcome the along-track sampling constraints that SAS imposes, most current SAS systems employ multiple-receiver hydrophone arrays (vernier-arrays). Reconstruction of data collected with these systems is more challenging than data from a highly sampled single-receiver equivalent. For most time-domain methods, little additional complication is introduced. This is not the case for Fourier-based methods. The following sections cover the two most prominent methods of reconstructing multiple-receiver SAS imagery using the wavenumber algorithm. Other Fourier-based imaging algorithms also require the use of the methods.

Many SAS systems employ a single transmitter and multiple-receiver configuration. This type of system is by definition *multistatic*, containing many bistatic transmit/receive pairs. Reconstruction of vernier-array data is often performed by massaging it into a form approximating the monostatic case (the phase-centre approximation) and using the standard wavenumber algorithm [Gough et al. 2000b; Hayes and Gough 1999] (see Section 5.2.1). Generalised reconstruction then follows from the simple monostatic (single-receiver) wavenumber reconstruction outlined in Section 5.2.2.

5.2.1   Phase-centre approximation inversion

The simplest way of reconstructing multiple-receiver sonar data is to treat the data as if it were collected using a monostatic sonar system. Usually, the return at each of the separate hydrophones is treated as if coming from a co-located transmit/receive transducer midway between the actual transmitter and receiver (via the phase-centre approximation). In this manner, each hydrophone provides a sample of an equivalent monostatic sonar.

Treating the data collection as monostatic allows the use of monostatic reconstruction algorithms if the along-track phase-centre samples are collected at the same positions as a monostatic sonar would collect them [Pat 2000]. However, the equivalent monostatic samples are usually taken with non-uniform spacing because it is very difficult to keep the imaging platform within the narrow velocity range where ideal sampling occurs [Pat 2000].

   The first suggestion in the open literature of using the wavenumber algorithm for multiple-receiver reconstruction was made soon after the algorithm's use in SAR reconstruction [Sheriff 1992]. The same paper describes the use of an along-track interpolation prior to reconstruction so that the monostatic equivalent samples are uniformly spaced. Both Hayes and Gough [1999] and Wilkinson [2001] describe an interpolation based on the DFT that calculates the monostatic wavenumber spectrum directly. Other possibilities include FFT algorithms designed for using non-uniformly spaced data [Lui et al. 1998]. An additional phase correction is also needed if the scene is not in the far field of the physical hydrophone array [Banks and Griffiths 2002; Bonifant 1999; Sheriff 1992; Wilkinson 2001]. In making the phase-centre approximation, the near-field phase terms are neglected (see Section 3.2.2 equation (3.18)). The suggested corrections are only valid for a given range and broadside to the sonar.

   Full time-shift compensation is valid only for the centre of the image with the phase-centre correction valid only for a given range. Phase modulation[5] of the pulse-compressed data set can provide approximate compensation for ranges other than the centre of the image. The angular-dependent error is much more difficult to correct and causes an additional lobe to appear in the reconstructed image (apparent in the conventionally compensated images in [Bonifant 1999, page 42]). Currently, the extra lobe is not considered a problem; as ultra-wide-beam sonars start to appear this may not remain the case indefinitely. A method for mitigating the error is presented in the next section.

Phase-centre compensation

This correction only applies for multiple receive-hydrophone sonars when using the reconstruction method described in Section 5.2.1. The purpose of the correction is to remove any distortion inducing phase errors caused by making the phase-centre approximation.

   The approximation error for a target at (x, y) is given by (3.18) and its 2-term Taylor series expansion is given by [Bonifant 1999]

$$\epsilon \approx \frac{\sigma^2}{4x} + \frac{\left(\frac{3y\sigma^2}{8} + \frac{3yu\sigma^2}{4} - \frac{3y^2\sigma^2}{8} - \frac{3u^2\sigma^2}{8} - \frac{3u\sigma^3}{8} - \frac{7\sigma^4}{64} \right)}{x^3}, \qquad (5.10)$$

where u and σ are the transmitter and hydrophone along-track positions respectively. The terms that vary in y and u are difficult to compensate, so dealing with only the terms involving x for the centre of the swath, x_0, then

$$\epsilon \approx \frac{1}{4}\frac{\sigma^2}{x_0} - \frac{7}{64}\frac{\sigma^4}{x_0^3} \qquad (5.11)$$

---

[5]A narrow-band approximation to time shifting data. The approximation error is minor if the shifts are small compared to the range-resolution.

is the broadside approximation error. Correcting the bulk error by compensating (5.11) using time-shifts allows the error for other ranges in the image to also be corrected using phase-only compensation[6].

Similar compensations are described in the literature concerning other multiple-receiver SAS system reconstruction techniques [Sheriff 1992; Wilkinson 2001; Yamaguchi 1999; Yamaguchi and Kato 1998].

Along-track interpolation

Once multiple-receiver, pulse-compressed sonar data is obtained, it needs to be coerced into a form suitable for single receiver reconstruction techniques. Starting with the multiple-receiver pulse-compressed data set $s(\omega, \sigma, u)$ the wavenumber representation of the monostatic equivalent sonar is given by [Hayes and Gough 1999]

$$S_m(\omega, k_m) = \iint s(\omega, \sigma, u) \exp\left(-jk_m m\right) \mathrm{d}u\, \mathrm{d}\sigma, \tag{5.12}$$

where m, the position of the co-located transducer, is given by

$$m = u + \sigma/2. \tag{5.13}$$

Expanding and separating (5.12) gives

$$S_m(\omega, k_m) = \int \left(\int s(\omega, \sigma, u) \exp\left(-jk_m u\right) \mathrm{d}u \right) \exp\left(-j\tfrac{1}{2}k_m \sigma\right) \mathrm{d}\sigma, \tag{5.14}$$

$$= \int S(\omega, \sigma, k_m) \exp\left(-j\tfrac{1}{2}k_m \sigma\right) \mathrm{d}\sigma, \tag{5.15}$$

$$= S(\omega, 2k_m, k_m), \tag{5.16}$$

which is a change of variables from $S(\omega, k_\sigma, k_u)$ to $S(\omega, 2k_m, k_m)$ where

$$k_\sigma = 2k_m, \tag{5.17}$$

$$k_u = k_m. \tag{5.18}$$

The DFT-based interpolation schemes simply calculate the Fourier integral (5.15) directly (see [Hayes and Gough 1999] and Wilkinson [2001] for more detail on the discrete implementations).

The major disadvantage of the technique is that an FFT over HP samples becomes an FFT over P pulses and a DFT over H hydrophones [Hayes and Gough 1999]. This reduces the algorithmic efficiency of the wavenumber algorithm from $\mathcal{O}\{HP \log(HP) N \log N\}$ to $\mathcal{O}\{H^2 P \log P N \log N\}$ (where N is the number of time

---

[6]The actual errors are likely to be small compared with a resolution cell once the bulk error has been compensated. This allows for accurate phase compensation of the remainder.

samples in the image). While for the low numbers of hydrophones in the current kiwi-SAS implementations the loss of efficiency is negligible [Hayes and Gough 1999], systems employing a large number of hydrophones will see a significant loss in efficiency.

Improved phase-centre compensation

Assuming that the interpolation to $S(\omega, k_m)$ is implemented via the DFT (5.15) improvements to the phase-centre correction are possible. With both hydrophone position, σ, and angle from the co-located transducer to target, θ_m, available in the DFT interpolation, the angular error caused by the phase-centres approximation (3.19) can be removed[7].

The corrected monostatic equivalent signal, $S_{\text{corr}}(\omega, k_m)$, is given by

$$S_{\text{corr}}(\omega, k_m) = \int S(\omega, \sigma, k_m) \exp\left(-j\tfrac{1}{2}k_m\sigma\right) \exp\left(-j\phi_{\text{corr}}\right) \mathrm{d}\sigma, \qquad (5.19)$$

where ϕ_{corr} is calculated based on the the phase-centre approximation error ϵ (from (3.18))

$$\phi_{\text{corr}} = \frac{\omega}{c}\epsilon, \qquad (5.20)$$

$$= \frac{\omega}{c}\left(\sqrt{x^2 + (u-y)^2} + \sqrt{x^2 + (u+\sigma-y)^2} - 2\sqrt{x^2 + (u+\sigma/2-y)^2}\right), \quad (5.21)$$

$$\approx \frac{\omega}{c}\left(\sqrt{x_0^2 + (x_0\tan\theta_m - \sigma/2)^2} + \sqrt{x_0^2 + (x_0\tan\theta_m + \sigma/2)^2} - \frac{2x_0}{\cos\theta_m}\right), \quad (5.22)$$

where the approximation is exact for the centre range of the image ($x = x_0$), and θ_m is calculated via

$$\theta_m = \cos^{-1}\frac{\sqrt{4k^2 - k_m^2}}{2k}. \qquad (5.23)$$

The penalty of the improved correction is additional computation; this penalty is minor if the correction is performed during reconstruction because the $\sqrt{4k^2 - k_m^2}$ term is calculated anyway. It should be acknowledged that the improved phase-centre correction offers little benefit to current generation SAS systems as the narrow-beam approximation error is small. Future ultra-wide-beam SAS systems (those with a 3 dB beam-width greater than 40°) will suffer blurring if the improved correction is not implemented.

Imaging example

Figure 5.2 shows an imaging example with and without phase centre correction on a simulated dataset from a hypothetical SAS. The hypothetical sonar has 32 receive elements spaced equally in along-track making up a 2 m array and travels 1 m between

---

[7]Wide-beam compensation of across-track array distortions is also possible although is of little benefit due to their small magnitude.

Figure 5.2   Comparison of along-track main-lobe responses for point target at 25.0 m range. (a) Standard wavenumber reconstruction technique, without phase-centre correction. (b) Standard wavenumber reconstruction with narrow-beam phase-centre correction applied for 25.0 m.

transmit pulses (corresponding to an (inadequately) $D/2$ sampled aperture [Hawkins and Gough 1997a]). The transmitted signal has a 20 kHz bandwidth at a centre frequency of 100 kHz. This configuration approximately matches the specification of the US Navy SAS under construction [Keeter 2001].

5.2.2   Bistatic inversion

The alternative to invoking the phase-centre approximation and treating a vernier-array as a number of monostatic samples, is to use a bistatic derivation of the wavenumber algorithm. Such a derivation considers the true *bistatic* path from transmitter to object and back to an individual receiver element.

Starting with the shallow water sonar response (3.4), taking the spatial Fourier transform, and using Weyl's identity (see Appendix B, [Chew 1995; Morse and Feshbach 1953]), gives the wavenumber-domain spatial-impulse response,

$$H(k, k_\sigma, k_u) = \frac{\exp\left(-j(\sqrt{k^2 - k_\sigma^2} + \sqrt{k^2 - (k_u - k_\sigma)^2})\,|x|\right)}{\sqrt{k^2 - k_\sigma^2}\sqrt{k^2 - (k_u - k_\sigma)^2}}. \tag{5.24}$$

Rewriting the received field using (5.24) gives,

$$E(\omega, k_\sigma, k_u) = \frac{1}{\sqrt{k^2 - k_\sigma^2}\sqrt{k^2 - (k_u - k_\sigma)^2}}$$
$$\times \iint f(x, y) \exp\left(-j\sqrt{k^2 - k_\sigma^2}\,|x| - j\sqrt{k^2 - (k_u - k_\sigma)^2}\,|x| - jk_\sigma y - jk_u y\right) \mathrm{d}x\,\mathrm{d}y. \tag{5.25}$$

Recognising the last term of (5.25) as a 2-D Fourier transform gives,

$$E(\omega, k_\sigma, k_u) = \frac{f\left(\left(\sqrt{k^2 - k_\sigma^2} + \sqrt{k^2 - (k_u - k_\sigma)^2}\right), k_u\right)}{\sqrt{k^2 - k_\sigma^2}\sqrt{k^2 - (k_u - k_\sigma)^2}}, \qquad (5.26)$$

where the Fourier domain change of variables is given by,

$$k_x = \sqrt{k^2 - k_\sigma^2} + \sqrt{k^2 - (k_u - k_\sigma)^2}, \qquad (5.27)$$

$$k_y = k_u. \qquad (5.28)$$

The basis of the wavenumber algorithm is to convert the measured data into the wavenumber domain and perform a coordinate transform[8]. An inverse Fourier transform now provides an estimate of the sea-floor scatterers. This may be summarised as,

$$\begin{aligned} \widehat{f}(k_x, k_y) = \mathcal{S}^{-1}\Big\{ & E(\omega, k_\sigma, k_u) \\ & \times \sqrt{k^2 - k_\sigma^2}\sqrt{k^2 - (k_u - k_\sigma)^2} \\ & \times \exp\left(j\left(\sqrt{k^2 - k_\sigma^2} + \sqrt{k^2 - (k_u - k_\sigma)^2}\right)x_0\right)\Big\}, \end{aligned} \qquad (5.29)$$

where $\mathcal{S}^{-1}\{\}$ is the Stolt coordinate transform of (5.27) and (5.28) and is performed via a frequency-domain interpolation. The image estimate, $\widehat{f}(x, y)$, is then obtained by inverse Fourier transforming $\widehat{f}(k_x, k_y)$.

Interpretation

A helpful notation in aiding interpretation is to measure both along-track variables relative to $y = 0$ (instead of σ relative to transmit position). In this notation, with the projector position given by y_{tx}, and the hydrophone position by y_{rx}, the Stolt coordinate transform is now given by

$$\begin{aligned} \widehat{f}(k_x, k_y) = \mathcal{S}^{-1}\Big\{ & \sqrt{k^2 - k_{rx}^2}\sqrt{k^2 - k_{tx}^2}\, E(\omega, k_{rx}, k_{tx} + k_{rx}) \\ & \exp\left(j\left(\sqrt{k^2 - k_{rx}^2} + \sqrt{k^2 - k_{tx}^2}\right)x_0\right)\Big\}, \end{aligned} \qquad (5.30)$$

where [Soumekh 1994]

$$k_x = \sqrt{k^2 - k_{tx}^2} + \sqrt{k^2 - k_{rx}^2}, \qquad (5.31)$$

$$k_y = k_{tx} + k_{rx}, \qquad (5.32)$$

---

[8]This coordinate transformation is usually implemented via a frequency-domain interpolation.

Figure 5.3   Stolt interpolation of multiple-receiver data using the bistatic wavenumber algorithm. Input samples on curved ellipses of constant temporal frequency. The transmitter and receiver define the focii of the ellipses.

Figure 5.4   Wavenumber spectral coverage detail. The points in 2-D wavenumber space that samples exist on are governed by the two-way propagation and the collection geometry. The wavenumber that interrogates the target is the vector sum of $\bar{k}_{\text{tx}}$ and $\bar{k}_{\text{rx}}$.

Figure 5.5   Notation from medical ultrasound for the scattering wavenumbers [Lerner and Waag 1988]. $\bar{k}_{\text{tx}}$ is the incident wavenumber (k_i), $\bar{k}_{\text{rx}}$ is the scattering wavenumber (k_s) and $\bar{k}$ is the imaging wavenumber. k_{tx} and k_{rx} are the individual components in the y direction of their respective vector equivalents $\bar{k}_{\text{tx}}$ and $\bar{k}_{\text{rx}}$.

and $k_{tx} = k \sin \theta_{tx}$ and $k_{rx} = k \sin \theta_{rx}$ relate to the angles to target seen at the respective transducers. Comparing with the previous derivation it is possible to show $k_{tx} = k_u - k_\sigma$ and that $k_{rx} = k_\sigma$. Having a monostatic arrangement means that $k_{tx} = k_{rx}$, i.e., the angle from transmitter to target is the same as from receiver to target.

An alternative way of representing the reconstruction described by (5.29) is also useful for interpretation. Using a few geometrical properties (5.29) may be written [Li et al. 1993; Soumekh 1991]

$$\widehat{f}(k_x, k_y) = \mathcal{S}^{-1} \left\{ \sqrt{k^2 - k_{rx}^2} \sqrt{k^2 - k_{tx}^2} \; E(\omega, k_{rx}, k_{tx} + k_{rx}) \right.$$
$$\left. \exp\left(j\left(\sqrt{4k^2 \cos^2 \beta - k_m^2} \right) x_0 \right) \right\}, \quad (5.33)$$

where 2β is the bistatic angle given by

$$\beta = \sin^{-1}\left(\frac{k_{rx} - k_{tx}}{2k} \right) = \sin^{-1}\left(\frac{2k_\sigma - k_u}{2k} \right),$$
$$\approx \frac{\theta_{rx} - \theta_{tx}}{2}, \quad (5.34)$$

with the approximation holding for $\beta \ll 1$ and

$$k_m = k_{tx} + k_{rx} = k_u = k_y. \quad (5.35)$$

This particular representation is sometimes seen in bistatic SAR literature [Jakowatz et al. 1996]. Equation (5.33) demonstrates that the bistatic collection geometry scales the imaging wavenumber by $\cos \beta$. When $\beta = 0$ the monostatic inversion and the bistatic inversion are identical.

Imaging sonars have a small bistatic angle, 2β, thus $k_u \approx 2k_\sigma$. The small bistatic angle is also why the phase-centre approximation ($k_u \equiv 2k_\sigma$, see Section 3.2.2) is usually valid. The additional processing performed by the bistatic formulation is to correct the error caused by the phase-centre approximation (for the entire image at once).

The only trick in implementing an improved bistatic wavenumber algorithm is in knowing the actual positions of the samples in 2-D wavenumber space. This should be faster than the phase-centre based algorithm above because it does not need to do the $\mathcal{O}\{H^2\}$ fractional interpolation (this is replaced with a standard FFT). The along-track frequency sample locations are geometry dependent and difficult to calculate. The bistatic wavenumber algorithm implementation presented in [Callow et al. 2002b] avoids the calculation but requires extra interpolations to account for this (H times more). The algorithm operates by interpolating onto a monostatic grid and averaging over all possible bistatic angles. This requires more computation than the optimum solution as the sonar samples only a limited set of bistatic angles.

(a)   (b)

Figure 5.6 Comparison of along-track main-lobe responses for point target at 25.0 m range reconstructed with centre of the swath at 150 m range. (a) Standard wavenumber reconstruction technique, applying phase-centre correction for 150 m. (b) Bistatic wavenumber reconstruction technique.

Imaging example

Figure 5.6 demonstrates the benefit obtained in using the new bistatic wavenumber reconstruction technique presented in [Callow et al. 2002b] on a simulated data set from the same sonar modelled previously. There is an obvious improvement in the main-lobe response; the side-lobes (not shown) have increased to levels slightly higher than expected, probably due to the effects of inadequate sampling (D/2 sampling).

5.2.3 Bathymetric wavenumber reconstruction

Many of the problems currently faced in interferometric SAS systems are caused by geometrical considerations [Banks and Griffiths 2002]. This is evident where images from vertically displaced receivers are reconstructed separately and the phases compared for the purpose of interferometry. If the terrain varies considerably and the path length is significantly different between the images (image-warping/footprint-shift), the image cross-correlation is severely reduced [Banks et al. 2001][9] Using images from the vertically displaced receivers and performing focusing in elevation as well as in azimuth, can alleviate some of these difficulties, particularly as the footprint-shift is removed. Reconstruction in elevation with a number of vertical hydrophones also allows layover effects [Gatelli et al. 1994] to be removed.

The algorithm presented is a generalisation of the monostatic wavenumber algorithm to 3-D space. Similar derivations for a 2-D synthetic aperture (used for creating 3-D SAR images [Busse 1992; Lopez-Sanchez and Fortuny-Guasch 2000]) are based

---

[9]A different effect causes the well known baseline de-correlation effect (which can be partly removed with appropriate pre-filtering [Gatelli et al. 1994]). In typical single-pass SAS interferometry this has little effect and foot-print shift is much more prevalent.

on the same techniques. The extension to vernier-array systems is a straightforward application of the techniques outlined earlier in the chapter.

Starting with the 3-D monostatic system model, (3.27), (summarised below)

$$E(\omega, k_u, k_h) =$$

$$P(\omega) \iiint \widehat{f}(x, y, z) \frac{\exp\left(-j\sqrt{4k^2 - k_u^2 - k_h^2}\,|x| - jk_u y - jk_h z\right)}{\sqrt{4k^2 - k_u^2 - k_h^2}}\,\mathrm{d}x\,\mathrm{d}y\,\mathrm{d}z, \quad (5.36)$$

where v and its Fourier pair k_h measure the vertical position of the monostatic hydrophone (position of the equivalent co-located transducer under the phase-centre approximation). The reconstruction is a Fourier domain change of variables,

$$k_x = \sqrt{4k^2 - k_u^2 - k_h^2}, \qquad (5.37)$$

$$k_y = k_u, \qquad (5.38)$$

$$k_z = k_h. \qquad (5.39)$$

Thus summarising the reconstruction

$$\widehat{f}(k_x, k_y, k_z) = \mathcal{S}^{-1}\Bigg\{ E(\omega, k_u, k_h)$$

$$\times \sqrt{4k^2 - k_u^2 - k_h^2} \qquad (5.40)$$

$$\times \exp\left(j\sqrt{4k^2 - k_u^2 - k_h^2}\,|x_0|\right)\Bigg\},$$

where $\mathcal{S}^{-1}\{\}$ performs the coordinate transform of equations (5.37), (5.38) and (5.39).

The result of the reconstruction is a 3-D volume map, $\widehat{f}(x, y, z)$. With all of the propagation effects corrected inside the reconstruction algorithm, phase-based interferometry is able to be applied by phase differencing slices at varying heights without additional footprint-shift correction or iteration. In addition, discrete effects in the 3-D inversion do not cause significant problems. The inversion amounts to a vertical beamformer and traditional array theory can be used to predict results. Vertical resolution is poor with a short vertical array but interferometric accuracy is unaffected. Even in the worst case—a single element—the inversion (5.40) becomes the same as a 2-D reconstruction, i.e., height is unresolvable.

5.3   Correction of second order effects in wavenumber processing

This section briefly outlines correction of some of the second order effects when using wavenumber or similar frequency-domain reconstruction methods. Most of the effects are discussed in Section 3.5.

5.3.1   Moving sonar compensation (avoiding the stop-and-hop approximation)

When imaging wide swaths, the receiving elements move some distance in along-track between transmission and reception (see Section 3.5.1). This is not generally a problem for monostatic sonars—the along-track velocity must be small in order to obey the sampling constraints and the distance moved is small. The effect can be substantial for multiple-receiver SAS systems that travel a significant distance while the sonar ping is in the water. The basic result is a skew in the reconstructed image (see Section 3.5).

In addition to the image skew, temporal Doppler-shifting has a minor effect on SAS images and is often neglected in systems using tolerant waveforms [Hawkins 1996]. The overall outcome is a slight loss in along-track resolution. Two effects contribute to the resolution loss: a shift in the along-track wavenumber, i.e., targets forward of the beam reconstruct at a different range to those backward of the beam when using linear-FM chirps, and a signal de-correlation caused by the Doppler shift moving the echo spectrum before pulse compression.

The within-pulse effects described are predictable in SAS systems (see Section 3.5.1) and can be compensated in straightforward fashion. For example, when using the time-domain correlation reconstruction method, both the position of the target and the angle from sonar to target are known at all times and the matched filter easily incorporates the effect[10].

Due to the ease of modelling temporal Doppler and image skewing effects in the wavenumber domain, the wavenumber algorithm is particularly suited to compensation processing. From (3.48), and (3.49), the wavenumber algorithm may be re-written to include moving sonar compensation. Thus

$$\widehat{f}(k_x, k_y) = \mathcal{S}_{\text{Doppler}}^{-1}\{E(\omega, k_u)P(\omega)k_x \exp\left(jk_x\left|x_0\right|\right)\}, \tag{5.41}$$

---

[10]The obvious downside to this reconstruction technique is large computational cost.

where the modified Stolt transform[11] for Doppler compensation, $\mathcal{S}_{\text{Doppler}}^{-1}\{\}$, is given by

$$k_x = \sqrt{\left(\frac{2(k - ak_u)}{(1 - a^2)}\right)^2 - \left(k_u + \frac{2a(k - ak_u)}{(1 - a^2)}\right)^2}, \tag{5.42}$$

$$k_y = k_u. \tag{5.43}$$

The modified Stolt transform may also be written

$$k_x = \sqrt{k_r^2 - (k_u + ak_r)^2}, \tag{5.44}$$

$$k_y = k_u, \tag{5.45}$$

where

$$k_r = \frac{2(k - ak_u)}{(1 - a^2)}, \tag{5.46}$$

and $a = v_s/c$, describe the effect of the sonar movement on illumination wavenumbers used.

It is worth noting that if the movement within the pulse is neglected, $a = 0$ and (5.46) and (5.44) become the same as used in the traditional monostatic wavenumber algorithm. Reconstruction using the modified wavenumber algorithm can thus conditionally provide within-pulse compensation, reverting to the traditional wavenumber algorithm when desired.

Compensation in the modified Stolt mapping (5.42), (5.43) is achieved by calculating the true illumination wavenumber, k_r (based on the measured wavenumbers k, k_u). Once the true illumination wavenumber is calculated, calculation of the appropriate image frequency-components in a mapping identical to the standard Stolt mapping is possible. The interpolation step of the wavenumber algorithm performs both the frequency-shift and scaling of k_x normally required in Doppler compensation. This remapping of the frequency spectra in the modified Stolt mapping causes only a marginally increased computational load—a significant savings compared with traditional Doppler compensation techniques.

Within-pulse movement and Doppler effects are accurately modelled (and compensated for) using the wavenumber algorithm modifications outlined above. Similar derivations for multiple-receiver reconstruction are achieved in straight-forward manner (phase-centre based wavenumber reconstruction follows directly from the modifications above). The modifications to the wavenumber algorithm presented here should be used in any wavenumber reconstruction suite—particularly as the traditional wavenumber algorithm is a special case of the modified version.

---

[11]The across-track spatial frequency k_x used for demodulation prior to the Stolt mapping is also given by (5.42).

5.4   Summary

This chapter has presented an investigation of the commonly-used wavenumber algorithm. Enhancements have been presented in light of the modelling improvements of Chapter 3 to obtain more accurate reconstructions with regard to image amplitude. A number of wavenumber algorithm implementation issues were discussed. The baseband wavenumber algorithm is recommended and high-order interpolators preferred over image subdivision. Interestingly, the interpolation used in the chirp-scaling algorithm was found to be possible to use for wavenumber interpolation suggesting some similarity in the algorithms.

Current methods for reconstructing multiple-receiver SAS data using the phase-centre approximation and along-track interpolation were described. Phase-centre approximation compensation was discussed and improvements for extending the size of the swath and using wide-beam modelling were suggested. The swath width improvement is useful for any sonar system where the imaged ranges are in the near field of the array. However, the wide-beam improvements only offer phase-centre reconstruction improvements for SAS's with beam-widths of over 45°.

A novel approach to multiple-receiver reconstruction using bistatic system modelling was suggested. The method aids understanding of the SAS imaging problem and provides a wavenumber domain interpretation of the phase-centre approximation. The new method hints at the possibility of improving reconstruction efficiency by avoiding the along-track interpolation used in traditional multiple-receiver reconstruction. More research is required to realise the potential of the algorithm.

A new derivation of the wavenumber algorithm for imaging in an unbounded 3-D medium including appropriate amplitude terms has been presented following similar derivations in the SAR field. This algorithm is suitable for bathymetric reconstruction and should assist interferometric processing. Other straightforward modifications allow improved compensation of the within-pulse second-order effects discussed in Section 3.5. Wavenumber domain algorithms allow compensation of such effects with only a minor increase in processing cost.

Chapter 6

Motion compensation for known path errors

Uncompensated motion errors have a severe effect on synthetic aperture sonar (SAS) imagery. Time-domain correlation reconstruction is able to compensate arbitrary track-errors but the more efficient frequency-domain reconstruction algorithms, such as the range-Doppler, chirp-scaling, and wavenumber algorithms do not inherently allow for straightforward compensation (especially for systems with wide beam-widths). Data processed via these block algorithms is usually compensated prior to reconstruction in a computationally inexpensive preprocessing step [Bonifant 1999; Hawkins 1996; Wilkinson 2001]. This compensation assumes a narrow-beam geometry, leading to blurring in wide-beam sonar images.

This chapter discuses techniques for compensation of known motion errors in wide-beam SAS data. The most promising technique relies on the multiple-receive element configuration typical in high-resolution SAS systems. The correction requires little extra processing over standard, narrow-beam correction. Admittedly, the benefit for systems with beam-widths of less than $10°$ is limited although becomes substantial for short-range imaging. The technique requires an extra FFT and inverse FFT along the receiver dimension compared with conventional motion compensation and so has an increased computational load.

6.1 Overview

Compensation for image distortion is possible given accurate estimates of the imaging platform's deviation from the nominal path. This type of compensation is useful given precise micronavigation (see Chapter 9, [Wang et al. 2001]) or precise inertial navigation system (INS). Compensation for known motion in a time-domain beam-former is straightforward; however the algorithm itself is computationally expensive [Banks and Griffiths 2002; Gough and Hawkins 1997; Hawkins 1996]. To improve computational efficiency it is better to use more efficient block reconstruction algorithms, [Bamler 1992]). Block algorithms for SAS motion compensation currently treat the error caused by motion as a timing-error [Hawkins 1996; Wilkinson 2001]. The timing-error approximation breaks down for wide-beam sonars. Section Section 6.2 gives a

more detailed analysis of this problem.

The image blurring caused by the timing-error assumption also causes difficulties in SAR systems. Typically, SAR systems form images using the efficient range-Doppler, wavenumber, or chirp-scaling based reconstruction algorithms and so wide-beam radar images [Madsen 2001; Soumekh 1999] also suffer some blurring under conventional motion compensation schemes. Moreover, SAR systems are often used for repeat-pass interferometry [Reigber 2001] and the loss of phase signature is extremely destructive in this application. Research is being undertaken in the radar community to mitigate these negative effects of conventional motion-compensation and retain the benefits of efficient reconstruction techniques [Madsen 2001; Potsis et al. 2001; Reigber 2001; Soumekh 1999; Ulander et al. 2000, 2001]. Post-reconstruction algorithms for compensation are based on space-variant filtering [Soumekh 1999] and short-term Fourier transformation [Potsis et al. 2001]. Another algorithm attempts the space-variant filtering by using techniques reminiscent of chirp-scaling [Madsen 2001]. Other research groups have been improving the back-propagation reconstruction technique, well suited to arbitrary-path reconstruction [Ulander et al. 2000, 2001]. This type of compensation is an active area of research in the SAR community and further developments are expected in the near future.

6.2   Motion compensation geometry

For a sonar with the 2-D multiple-receiver collection geometry shown in Figure 3.3, the recorded sonar echo as a function of time, t, along-track transmitter position, u, and along-track receiver offset, σ, may be expressed as (see Chapter 3, [Callow et al. 2002b]),

$$e(t, u, \sigma) \approx \iint \frac{f(x,y)\; p\Big(t - \frac{1}{c}\Big(\sqrt{x^2 + (y-u)^2} + \sqrt{x^2 + (y-u-\sigma)^2}\Big)\Big)}{(4\pi)^2 \sqrt{x^2 + (u-y)^2}\sqrt{x^2 + (u+\sigma - y)^2}}\; \mathrm{d}x\,\mathrm{d}y,$$

(6.1)

where $f(x,y)$ is the illuminated scene and $p(t)$ is the transmitted signal[1]. To make the following analysis more tractable, the phase centres approximation (see Section 3.2.2, [Callow et al. 2001a, 2002b; Wang et al. 2001]) can be used to give a revised system model

$$e(t, u, \sigma) \approx \iint \frac{f(x,y)\; p\Big(t - \frac{2}{c}\sqrt{x^2 + (u+\sigma/2 - y)^2}\Big)}{(4\pi)^2(x^2 + (u+\sigma/2 - y)^2)}\; \mathrm{d}x\,\mathrm{d}y.$$

(6.2)

This approximation treats a transmitter-receiver pair as virtual transmit and receive transducers co-located mid-way between the actual transducers. After introducing a platform position error given by $X(u)$ into the system model (6.2)) the (distorted) echo

---

[1] Beam-patterning and the vertical dimension are ignored in this representation.

signal may be represented by[2]

$$\widetilde{e}(t, u, \sigma) = \iint \frac{f(x,y) \, p\left(t - \frac{2}{c}\sqrt{(x - X(u))^2 + (u + \sigma/2 - y)^2}\right)}{(4\pi)^2((x - X(u))^2 + (u + \sigma/2 - y)^2)} \, \mathrm{d}x \, \mathrm{d}y. \qquad (6.3)$$

Simplifying the distorted echo signal (6.3) by using the ideal signal (6.2), assuming that the amplitude variation is negligible, that $X(u) \ll x$, and that $xX(u) \ll r^2$, gives

$$\widetilde{e}(t, u, \sigma) \approx e\left(t - \frac{2}{c} \, X(u) \cos \theta, u, \sigma\right). \qquad (6.4)$$

The consequence of modelling wide-beam motion errors is that the amount of motion-caused distortion is dependent on the angle to the target. Wide-beam motion compensation techniques must therefore account for this in their processing.

6.2.1 Narrowbeam approximation

Provided that the system is narrow-beam and $|u + \sigma| \ll |x|$, then the distorted echo model (6.3) may be approximated,

$$\widetilde{e}(t, u, \sigma) \approx e\left(t - \frac{2}{c} \, X(u), u, \sigma\right). \qquad (6.5)$$

Comparing narrow-beam compensation using (6.5) with wide-beam using (6.4), it becomes evident that as the beam-width becomes narrower, the wide-beam compensation approaches the narrow-beam (timing-error) compensation. This is a consequence of approximating $\cos \theta \approx 1$ for small θ in the wide-beam system model (6.4). Figure 6.1 demonstrates the motion-compensation geometry of a wide-beam system and shows the difference in wide and narrow-beam compensation.

6.3 Timing-error based sway compensation

When reconstructing SAS/SAR imagery using block (Fourier based) algorithms, known motion-errors are normally treated as timing-errors. This assumes that a sideways displacement from the straight flight-path can be treated as an equivalent timing-error in the raw data (via (6.5)). These *pseudo* timing-errors are then removed, prior to reconstruction, by time-shifting the received signal on a pulse-by-pulse basis. The time-shift is often applied using a linear phase-shift in the temporal-frequency domain to avoid having to perform sub-sample interpolation; Shippey et al. [1998a] demonstrates another method of implementation using phase-modulation and interpolation of the baseband envelope. The timing-error approach neglects the $\cos \theta$ dependence.

---

[2]Amplitude effects caused by the shift in position are negligible as $|x| \gg |X(u)|$ for most of the image.

Figure 6.1   Wide-beam motion compensation geometry showing how data is incorrectly compensated using the normal motion compensation steps. The correction normally used, Δr, is applied as if the entire signal has been collected at broadside. The correction that should be applied, Δr_{tgt}, is a function of the angle to the target θ [Madsen 2001].

Narrow-beam motion compensation does not apply for the high-resolution, wide-beam systems coming into use. This is evident in the error of the approximation

$$\epsilon = 2X(1 - \cos\theta). \tag{6.6}$$

For one of the high resolution SAS systems currently under development [Keeter 2001], a one metre sway from nominal corrected using timing-error based compensation causes a 2λ phase error for target information at the edge of the beam. The error is even worse ($\approx 6\lambda$) for some wide-beam, sonar systems [Cook et al. 2001]. This causes blurring and destroys the phase information in the image.

6.4   Improved wide-beam sway compensation

Compensation of the entire scene via (6.4) (including angles other than broadside) requires an estimate of the angle to all points in the scene for each ping. This section discusses wide-beam motion compensation for single and multiple-receiver SAS systems.

6.4.1   Single-receiver wide-beam motion compensation

In a single-receiver system, as is typical for SAR systems, an estimate of the angle to target at each pulse is generated using short-term Fourier transformation of a limited number of pulses [Potsis et al. 2001]. 2-D space-variant spatial-filtering [Soumekh 1999] achieves the same result. With the angle estimate (generated using the combined space, spatial-frequency representation of the algorithms) wide-beam motion compensation is

possible. Madsen [2001] follows a similar approach by performing space-variant filtering on local image patches and using chirp-scaling techniques to extend the patch size.

6.4.2   Multiple-receiver wide-beam motion compensation

Wide-beam motion compensation complexity is reduced for multiple-receiver systems because many samples are collected for each sonar pulse. In multiple-receiver SAS systems, the direction to target (for all targets) can be estimated for each ping— avoiding combined space, spatial-frequency representations. In a swaying multiple-receiver system all of the receivers suffer the same sway error for a given ping. By Fourier transforming the multiple-receiver data (along the receivers) to get a beam-space representation of the data, wide-beam motion correction can be applied. The direct application of wide-beam techniques provides motion compensation for targets independent of their angle from bore-sight.

The algorithm for improved motion compensation is implemented as follows. Starting with the received echo signal, take a 2-D Fourier transform along the receiver array and the time-series data. Recognising the discrete nature of the pulse ($u \equiv u_p$ for the p^{th} ping) gives

$$e(\omega, k_\sigma, u_p) = \mathcal{F}_{t \to \omega, \ \sigma \to k_\sigma}\{e(t, \sigma, u_p)\}. \tag{6.7}$$

This Fourier transform provides information about the angle of arrival because

$$k_\sigma = k \sin \theta_\sigma, \tag{6.8}$$

where θ_σ is the angle of arrival of the signal at the receivers.

Based on the argument outlined above and using (6.8) the wide-beam correction for the displacement, X_p, at pulse p, is given by

$$e(\omega, k_\sigma, u_p) = \widetilde{e}(\omega, k_\sigma, u_p) \exp\left(j2X_p\sqrt{k^2 - k_\sigma^2}\right), \tag{6.9}$$

where $\widetilde{e}(\omega, k_\sigma, u_p)$ is the (distorted) recorded echo signal. The wide-beam compensation (6.9) can also be derived by using the bistatic modelling from Chapter 3 and making the phase centres approximation [Callow et al. 2002b]. Thus

$$E(\omega, k_\sigma, k_u) = \widetilde{E}(\omega, k_\sigma, k_u) \exp\left(j2X_p\sqrt{k^2 - \tfrac{1}{4}k_y^2}\right). \tag{6.10}$$

Recalling that $k_y \approx 2k_\sigma \approx k_u$, $2k_\sigma$ allows correction of the pulse-to-pulse[3] sway X_p to be derived from (6.10) giving the same result as (6.9).

---

[3]Again the result is approximate. k_σ provides the coarse spatial-frequency information and k_u the fine information. In a pulse-by-pulse motion compensation scheme k_y is replaced with k_σ and an inverse Fourier transform over k_u is taken.

The wide-beam compensation operation specified in (6.9) is efficient. Compared to standard narrow-beam correction, given by

$$e(\omega, \sigma, u_p) = \widetilde{e}(\omega, \sigma, u_p) \exp\left(j2kX_p\right), \tag{6.11}$$

for a narrow-beam system, $k \gg k_\sigma$ and so the wide-beam correction (6.9) approaches the narrow-beam (timing-error) correction shown in (6.11). Wide-beam compensation has only a small increase in computation due to the Fourier transform from σ to k_σ for each ping.

For systems with few receivers, corrections are approximate (but better than correction to broadside) because there are few independent angle estimates (only as many resolvable spatial-frequencies as receivers in the Fourier transform from σ to k_σ). Further improved compensation can be performed using a hybrid of the new algorithm and the single-receiver motion compensation algorithms. First stage correction is performed using multiple-receiver motion compensation and later refined using single-receiver techniques. In the limiting case of a single-receiver system, the spatial-filtering algorithms used by the SAR community are necessary (see previous section).

The multiple-receiver motion compensation algorithm is equivalent to the single-receiver motion compensation methods described earlier. The transform over receivers in (6.7) is the same as the Fourier transform over multiple pulses in the algorithm described by Potsis et al. [2001]. The advantage of the multiple-receiver algorithm is that a Fourier transform is needed only at each pulse (rather than at every along-track sample). This is possible because all of the receivers at a given pulse are in a straight line. In a single-receiver system each along-track sample has a different sway preventing this saving.

The algorithm presented here provides a way of performing wide-beam motion compensation without the computational expense of time-domain correlation. The algorithm is also efficient, consisting only of Fourier transformation and phase multiplication. A single-receiver system is a special case and the algorithm offers no improvement over traditional compensation—the single-receiver algorithms outlined in the previous section should be used. Improvement over traditional narrow-beam compensation is gained for any system with a multiple-receiver geometry. Improvement is greater the more receiver elements in the system[4].

6.4.3   Multiple pulse motion compensation

For platform motion where many pulses have a similar sway, an extension to the method described in Section 6.4.2 provides motion-compensation (MOCOMP) improvement.

---

[4]In this algorithm the improvement is limited by the limited number of receivers in a straight line; this is due to the finite, fixed extent of the receiver array limiting angular resolution.

Like the motion compensation algorithm described in the previous section, the algorithm requires consecutive along-track samples to have the same sway.

Instead of using the angular measurements from a single pulse for motion compensation (as in the previous motion compensation algorithm), several consecutive pulses are used, enhancing the angular accuracy. Compensation of the sway common to the consecutive pulses is enhanced. For differing sways in the consecutive pulses, angular resolution degrades and the algorithm fails to provide improvement[5]. If all consecutive pulses have the same sway, i.e., a constant across-track offset, alteration of x_0, the range offset parameter used in the wavenumber/chirp-scaling/range-Doppler algorithms, provides the necessary compensation. Note that the algorithm can be implemented iteratively, with different along-track straight sections used at each stage.

Motion compensation is a combined space/spatial-frequency problem—the correction of motion at a particular position requires differing amounts of correction for echoes coming from various angles (spatial-frequencies). For ideal motion-compensation, back propagation of the echoes is needed, this is implemented as a space-varying, spatial-frequency phase filter [Soumekh 1999]. Methods such as described in this section make use of additional prior information about geometry (multiple pulses in a straight line) to allow beams to be formed from a number of positions and approximate the angle-variant correction required.

The approximate motion compensation discussed shows many similarities to FFBP (see Section 4.3.1). The reason is that the motion compensation problem can be implemented as a back/forward projection (FFBP is motion compensation for a large offset sway of x_0). One benefit of performing motion compensation iteratively using the described method and reconstructing using the wavenumber algorithm—only a single stage (single ping) motion compensation is usually required, reducing complexity, and computational effort compared with FFBP.

6.4.4   Individual element motion compensation

Compensation for the movement of individual elements in a receiver array should be performed using the standard narrow-beam method. It is very unlikely that an individual element is displaced far enough that correction for broadside is inadequate. In the rare cases where broadside correction is not sufficient, a 2-D spatial filtering operation is needed (the alternative—better alignment of the receivers—is preferable). Depending on the wavenumber algorithm interpolation used, the phase-centre correction technique described in Section 5.2.1 can be used for this purpose.

---

[5]These should be compensated beforehand using the algorithm described in the previous section.

6.4.5 Across-track motion during the ping

For the purpose of this discussion only a single receiver, narrow beam sonar is considered so that the timing error approximation [Hawkins 1996] applies. It is possible to extend the analysis to multiple-receiver, wide-beam sonars in a straightforward manner. The standard motion compensation correction applied for a sonar of this type is

$$e(\omega, u) = \widetilde{e}(\omega, u) \exp\left(j2kX(u)\right), \tag{6.12}$$

where $\widetilde{e}$ is the distorted echo, k is the wavenumber, and $X(u)$ is the across-track displacement for a given along-track position u.

A problem in attempting wide swath reconstruction (when significant motion is present) is the timing difference between close and far range. When the required timing difference is time variant, the MOCOMP should also be time-variant: all ranges suffer the transmit timing-error but close-range echoes are received earlier than far-range echoes and so have a different receive timing-error. For echoes coming from the extreme far range, the receive error is the same as the transmit of the following ping[6].

One way of improving the motion compensation, other than dividing the problem into multiple range swaths, is to assume the echo signal undergoes a more complicated distortion than previously considered. The assumption that the across-track displacement, $X(u)$, varies linearly during the time the sonar pulse is in the water, allows the motion to be modelled as a phase and frequency shift [Madsen 2001]. This leads to the motion distorted signal being represented as[7],

$$\widetilde{e}(\omega, u) \approx e(\omega - a(u), u) \exp\left(-j2kb(u)\right), \tag{6.13}$$

where $a(u)$ and $b(u)$ are derived constants for a given along-track position (that depend on the displacement at the start and finish of the ping).

Using the narrow-beam equivalence of time and space in the across-track direction ($t = 2r/c \approx 2x/c$) [Hawkins 1996], a constant linear across-track crabbing will cause a shift in the temporal-frequency domain (the basic principle behind Doppler shifting). Correction of motion is a matter of frequency shifting and later phase shifting the echo data. This is equivalent to performing time shifting to correct the bulk error and using phase compensation to fix the time variant shifts.

A more complicated compensation based on chirp-scaling principle is also possible. If enough across-track movement takes place, the swath width of the pulse-compressed image, is altered. The improved compensation then rescales the time-domain (by chirping, scaling and then performing phase compensation).

---

[6]Assuming only one pulse is in the water at a time.
[7]Also making an assumption that the system is narrow band.

6.4.6 Bulk yaw compensation

Sometimes there is a need to rotate the entire synthetic aperture around the centre of the image patch [Cook et al. 2001]. Steering the entire synthetic aperture is akin to performing yaw compensation.

Recognising the yaw as a rotation leads to the observation that the wavenumber spectrum of the final image is also rotated [Owen and Makedon 1996]. Rotation of the wavenumber spectrum provides a way of compensating arbitrary yaw errors. Because of the wavenumber algorithm's properties, the rotation may be applied by remapping the wavenumber spectrum of the echo data (and allowing the wavenumber interpolation to complete the rotation).

The current compensation technique (time shifting each receiver independently) is the same as shifting the wavenumber spectrum [Owen and Makedon 1996]. Instead of performing a shift, which is a narrow-beam, narrow-band approximation, the wavenumber spectrum requires a shift and frequency scaling to remove yaw without approximation. This may be easily implemented as a frequency domain interpolation requiring only a simple, low order interpolator.

For the purposes of this discussion, a single receiver sonar with a rotated *synthetic* aperture (i.e., the data is collected on a rotated geometry) is considered. Recall that (from Chapter 3, equations (3.25),(3.26))

$$\begin{aligned}
k_x &= \sqrt{4k^2 - k_u^2}, \\
k_y &= k_u,
\end{aligned} \qquad (6.14)$$

the k, k_u pair can instead by parametrised as a radius and angle,

$$k, k_u \Rightarrow k, \Phi, \qquad (6.15)$$

where

$$\begin{aligned}
k_u &= 2k \sin \Phi, \\
k_y &= 2k \sin \Phi, \\
k_x &= 2k \cos \Phi,
\end{aligned} \qquad (6.16)$$

and $\Phi = \sin^{-1}(k_u/2k)$. If $k_u > 2k$ then Φ has imaginary components and describes evanescent waves[8]. By parameterising the reconstruction in terms of angle, it becomes obvious that the wavenumber-domain rotation caused by yaw may be compensated using modifications to the wavenumber algorithm.

When the aperture is steered by $10°$ the resulting Fourier image is also rotated

---

[8]Evanescent waves are not important in SAS imaging as they do not propagate [Morse and Feshbach 1953]. The k_u wavenumber of an imaging system is limited to $< 2k$.

(a) No rotation   (b) Rotated

Figure 6.2 Effect of rotation on the offset Fourier data collection. (a) shows data collected with normal geometry. (b) shows the data collected if the sonar has been yawed during imaging—the red dashed box indicates the original data location.

by 10°. This wavenumber-domain rotation should be considered when reconstructing and interpolating to get k_x, k_y from k, k_u. When performing the standard wavenumber reconstruction the output k_x, k_y are rotated if the collection geometry is rotated. Figure 6.2 demonstrates the effect of a yaw on the support region of the Fourier space.

Instead of using the rotated k_x, k_y, the appropriate values for k_x, k_y can be retrieved at the interpolation step of the wavenumber algorithm (see Chapter 5). A yaw-corrected version of the wavenumber algorithm can then be performed using

$$\Phi = \sin^{-1}\left(\frac{k_u}{2k}\right), \tag{6.17}$$

as

$$
\begin{aligned}
k_{y_{\mathrm{yc}}} &= 2k\sin(\Phi + \Phi_{\mathrm{yaw}}), \\
k_{x_{\mathrm{yc}}} &= 2k\cos(\Phi + \Phi_{\mathrm{yaw}}) = \sqrt{4k^2 - k_{y_{\mathrm{yc}}}^2},
\end{aligned}
\tag{6.18}
$$

Exploiting the double angle formula, $\sin\left(\alpha \pm \beta\right) = \sin(\alpha)\cos(\beta) \pm \cos(\alpha)\sin(\beta)$, in (6.18) gives,

$$
\begin{aligned}
k_{y_{\mathrm{yc}}} &= 2k\sin(\Phi)\cos(\Phi_{\mathrm{yaw}}) + 2k\cos(\Phi)\sin(\Phi_{\mathrm{yaw}}), \\
&= k_u\cos(\Phi_{\mathrm{yaw}}) + \sqrt{4k^2 - k_u^2}\sin(\Phi_{\mathrm{yaw}}),
\end{aligned}
\tag{6.19}
$$

with the corresponding $k_{x_{\mathrm{yc}}}$ as,

$$k_{x_{\mathrm{yc}}} = \sqrt{4k^2 - k_{y_{\mathrm{yc}}}^2}. \tag{6.20}$$

Another way of solving the problem that does not exploit the wavenumber algorithm is to rotate either the final output image, or the Fourier domain version of the output image. Owen and Makedon [1996] give a method suited to SAS imagery.

Standard yaw correction

When the sonar platform is subject to large lateral currents, platform crabbing occurs. Imagery is formed along the line the platform travels but the transducers point in a direction other than broadside. This crabbing appears as a shear rather than a true rotational yaw of the aperture. Pat, in his thesis [Pat 2000], describes some of the effects this has on final synthetic aperture imagery.

Now if a yaw in the recorded data, $e(t, u)$, is described by a simple shear instead of rotation (small yaw angle approximation or crabbing data) and we treat the problem as a timing-error (narrow beam approximation, $\sin \widetilde{\Phi}_{\text{shear}} \approx \widetilde{\Phi}_{\text{shear}}$) we get the following expression for the collected data,

$$e_{\text{yaw}}(t, u) = e(t - 2\widetilde{\Phi}_{\text{shear}} u/c, u), \tag{6.21}$$

where $\widetilde{\Phi}_{\text{shear}} \, u$ is the shear distance. The 2-D Fourier Transform of $e_{\text{yaw}}(t, u)$ is then given by [Owen and Makedon 1996]

$$E_{\text{yaw}}(\omega, k_u) = E(\omega, k_u + 2\widetilde{\Phi}_{\text{shear}} k). \tag{6.22}$$

We can derive the equivalent wavenumber reconstruction in straightforward fashion as

$$\begin{aligned} k_{y_{\text{yc}}} &= k_u + 2\widetilde{\Phi}_{\text{shear}} k, \\ k_{x_{\text{yc}}} &= \sqrt{4k^2 - k_{y_{\text{yc}}}^2}. \end{aligned} \tag{6.23}$$

Applying the approximations made in the derivation of (6.23) to (6.19) it is possible to show equivalence of the techniques. Assuming $\widetilde{\Phi}_{\text{shear}}$ is small in (6.19)

$$\cos(\widetilde{\Phi}_{\text{shear}}) = 1, \tag{6.24}$$

$$\sin(\widetilde{\Phi}_{\text{shear}}) = \widetilde{\Phi}_{\text{shear}}, \tag{6.25}$$

gives

$$k_{y_{\text{yc}}} \approx k_u + \sqrt{4k^2 - k_u^2}\,\widetilde{\Phi}_{\text{shear}}. \tag{6.26}$$

Now for $k_u \ll k$, a narrow-beam system, (6.26) becomes

$$k_{y_{\text{yc}}} \approx k_u + 2k\widetilde{\Phi}_{\text{shear}}, \tag{6.27}$$

with the reconstruction following and being identical to (6.23)[9].

The error in narrow-beam correction (noted previously in Section 6.3) affects traditional single-receiver yaw compensation and the wide-angle nature of the scene should ideally be taken into account. Accounting for the error is necessary when correcting large (>0.1 rad) yaw errors. The error using standard yaw compensation in wide-aspect imagery is noted in [Cook et al. 2001], although it was related in a different context.

6.4.7 Multiple-receiver yaw correction

Yaw that changes from pulse-to-pulse is also a problem for a multiple-receiver SAS system. Yawing the receiver array by only a small angle around the centre of rotation amounts to a large displacement for the receivers at the ends of the array. This displacement, like other receiver displacements, leads to blurring in the SAS image.

The data recorded by a receiver array suffering a yaw at a particular position of the sonar platform, $\Phi_\sigma(u)$, may be represented as (under the same approximations as (6.4))

$$\widetilde{e}(t, u, \sigma) \approx$$
$$\iint \frac{f(x,y)\, p\Big(t - \frac{2}{c}\sqrt{(x - \sigma/2\ \sin\Phi_\sigma(u))^2 + (u + \sigma/2\cos\Phi_\sigma(u) - y)^2}\Big)}{(x^2 + (u + \sigma/2 - y)^2)}\, \mathrm{d}x\, \mathrm{d}y, \quad (6.28)$$

which can be seen to be a rotation of the receiver array around $u = 0$ by $\Phi_\sigma(u)$. In the past, yaw has been corrected in a similar fashion to sway errors: the echo from each receiver is individually time shifted to correct for the position error, [Wilkinson 2001]. The errors in conventional pulse-to-pulse yaw correction are negligible in current SAS systems. This is due to a small[10] equivalent sway at each receiver, the case even for moderately large towfish yaws of ≈ 0.1 rad.

Following the analysis presented in the previous section, the yawed data set may be corrected via,

$$k_{\sigma_{\mathrm{yc}}}(u) = k_\sigma(u)\cos(\Phi_\sigma(u)) + \sqrt{k^2 - k_\sigma(u)^2}\sin(\Phi_\sigma(u)), \quad (6.29)$$

where u refers to the sonar pulse in question. The global (non ping-to-ping) yaw on a multiple-receiver system is more complicated. Assuming all ping-to-ping yaws have been corrected, the compensation is the same as in Section 6.4.6,

$$k_{u_{\mathrm{yc}}} = \cos(\Phi_u) + \sqrt{k^2 - k_u^2}\sin(\Phi_u), \quad (6.30)$$

---

[9]The small yaw approximation also means that $\Phi_{\mathrm{yaw}} \approx \widetilde{\Phi}_{\mathrm{shear}}$ in this discussion.

[10]This is due to the relatively short receiver arrays used in SAS systems.

and

$$k_{y_{\text{yc}}} = k_{u_{\text{yc}}},$$
$$k_{x_{\text{yc}}} = \sqrt{k^2 - k_u^2} + \sqrt{k^2 - (k_{u_{\text{yc}}} - k_{\sigma_{\text{yc}}})^2},$$

(6.31)

where k_u is the Fourier transform down the u direction.

Given that the pulse-to-pulse yaw differences are small they can often be corrected with narrow-beam correction techniques—the motion-compensation improvement does not warrant the increased computation required.

6.5   Results from simulated system

To test the proposed motion compensation algorithm, the low frequency system operated by the Coastal Systems Station (CSS, a Florida-based US-Navy research group) [Cook et al. 2001] was simulated using a standard ray-tracing based simulation. A scene consisting of discrete prominent points was used to help illustrate compensation performance. Reconstruction of the echo data was via a modified wavenumber algorithm [Wilkinson 2001]. For reference purposes a simulation was run with no path distortion added. The resulting image is shown in Figure 6.3(a).

To test motion compensation an across-track sway of Figure 6.3(b) was inserted during the simulation. When reconstructed using standard timing-error based compensation, some distortion of the point reflectors is apparent (Figure 6.3(c)). A significantly improved result is achieved when wide-beam motion compensation is used before reconstruction (Figure 6.3(d)).

6.5.1   Summary

As is clear from the results obtained from simulation of the CSS sonar multiple-receiver SAS system, the algorithm described offers significant motion compensation improvement when used with existing block reconstruction algorithms. The motion compensation improvement is enhanced when the SAS has many receivers. The algorithm requires only a trivial increase in computational cost if used in conjunction with Fourier-based reconstruction algorithms.

6.6   Summary

Motion compensation is required to allow efficient Fourier domain methods to reconstruct images taken with a non uniform geometry. The wide-beam motion model and motion compensation improvements discussed in this chapter allow improved reconstruction accuracy.
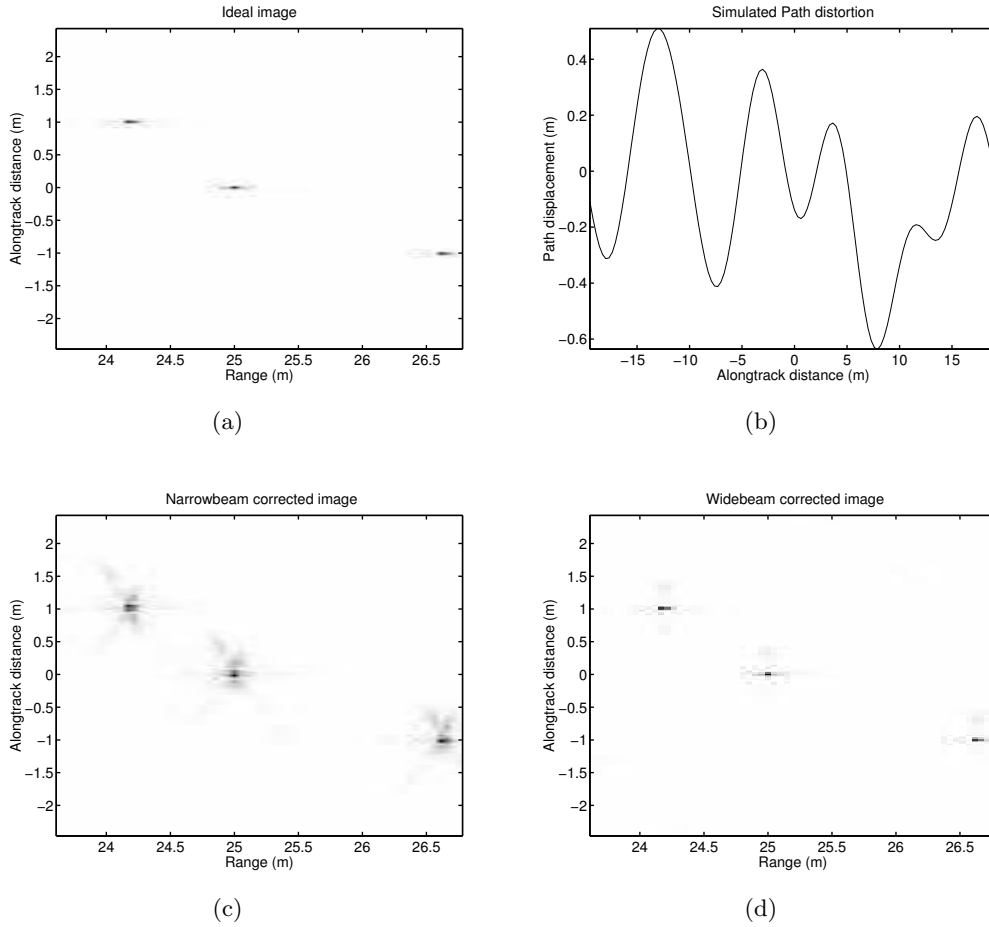
Figure 6.3   Demonstration of proposed motion compensation scheme. (a) Ideal reconstructed image with no distortion added. (b) Uncorrelated path distortion (sway) $X(u)$ as a function of along-track distance. (c) Reconstructed image with across-track path distortion shown in (b) and narrow-beam sway compensation via (6.11). (d) Reconstructed image with across-track path distortion shown in (b) and wide-beam sway compensation via (6.9).

A novel wide-beam motion compensation technique for multiple-receiver systems was suggested with the motion compensation improvement dependent on the number of receiver elements. These improvements have been demonstrated with simulated data from a wide-beam system. Single receiver improvements (following related SAR techniques) and extensions of the multiple-receiver method based on them were also presented. Using these techniques, full motion compensation is possible.

In addition, improved yaw compensation techniques for both pulse-to-pulse array yaw in multiple-receiver systems and bulk aperture yaw have been discussed. These techniques allow the wide-beam compensation of large (> 0.1 rad) yaws. Specialised extensions of the wavenumber algorithm for systems suffering yaw have been presented with traditional yaw compensation techniques derived as a special case of the wide-beam algorithms.

Autofocus fundamentals

In SAS imagery, turbulence in the medium or unknown path movements corrupt the phase of the echo signals leading to image blurring. Phase-distortions represent a major obstacle preventing the widespread use of SAS imaging. The lengthy aperture-creation times SAS requires mean that the platform often deviates a large number of wavelengths from the ideal straight path over the aperture. This deviation must be estimated to sub-wavelength accuracy to prevent image blurring.

The removal of blur-causing phase-distortions often requires data-driven phase-retrieval or autofocus techniques. These techniques are widely used in narrow-beam, narrow-bandwidth, spotlight mode SAR to improve the system performance. This chapter provides an introduction to the problem posed by phase-distortions and gives an overview of common phase estimation kernels used in autofocus algorithms.

7.1 Motivation for autofocus

Synthetic aperture creation requires a number of sonar pulses taken at well known locations. Unknown path-deviations of only fractions of a wavelength ($< \lambda/16$) cause the resulting synthetic aperture imagery to blur. Both SAS and SAR imagery suffer from this blurring, although the problem is less significant with the high pulse-repetition-frequency and short integration times of SAR systems [Hawkins 1996]. Moreover, the accuracy of a typical GPS-locked inertial navigation system (INS) is sufficient to allow close to diffraction-limited SAR imagery [Potsis et al. 2001]. In contrast, the accuracy of even the best current INS is not enough to allow diffraction-limited imagery of SAS images [Bellettini and Pinto 2002; Pinto et al. 2002; Wang et al. 2001].

Instead, autofocus algorithms must be used to estimate the phase-distortions caused by path errors. Autofocus algorithms estimate the platform trajectory and remove residual image blurring using the collected data. Another name for this type of algorithm is *micronavigation*, where the path estimates are usually smoothed using Kalman filtering and INS measurements [Pinto et al. 2002].

7.2 The autofocus problem

Synthetic aperture autofocus was originally developed in the SAR community for estimating the *Doppler-rate* error and *Doppler-centroid* of the echoes. Estimation of the Doppler-centroid was originally known as *clutterlock* and estimation of the Doppler-rate known as *autofocus*[1] [Berizzi et al. 1997; Carrera et al. 1995; Curlander and McDonough 1996; Madsen 1989; Prati and Rocca 1992]. The Doppler-rate error is due to errors in estimating the SAR platform speed (caused by orbital effects). The imagery from a system suffering a Doppler-rate error has a quadratic defocus (which map-drift and similar algorithms were designed to estimate [Carrera et al. 1995]). Doppler-centroid errors occur when the beam pointing direction moves during aperture creation. Doppler-centroid errors cause along-track blurring effects similar to those caused by low-order sway errors in stripmap systems.

As the resolution of SAR systems improved, atmosphere turbulence and high-frequency near-sinusoidal path errors started to become increasingly important. The term autofocus was extended to algorithms estimating any order of phase error, such as multi-aperture map-drift [Curlander and McDonough 1996] and PGA [Curlander and McDonough 1996; Jakowatz et al. 1996]. Currently, the term autofocus is used to describe the estimation and correction of any number of unknown parameters, including the entire unknown path.

The goal for an autofocus algorithm is to estimate a number of unknown parameters, usually the platform's path-deviation at each pulse. Autofocus is often achieved in an iterative framework, using only the recorded echo data. The related problem of micronavigation is usually aided by an on-board INS and is not generally iterative (see Section 7.4). SAS autofocus involves estimating and correcting the blurring caused by platform path error and medium fluctuation.

Phase errors must be estimated to better than $\lambda/8$ over the length of the aperture [Hayes and Gough 1992]. The constraints for motion are derived using the two-way propagation path—echos still sum coherently if position errors of less than $\lambda/16$ exist. When unknown motion/medium fluctuations cause phase-distortion greater than $\lambda/8$, echos do not sum coherently and the SAS image suffers degradation. The constraints on high frequency sinusoidal path errors are even tighter and are $< \lambda/60$ over the length of the aperture [Carrera et al. 1995, Chapter 5][Wang and Huang 1997]. Note that slowly varying phase errors across the aperture cause main-lobe broadening whereas rapidly varying phase errors result in raising of the side-lobes [Wang and Huang 1997].

---

[1]An incorrectly estimated Doppler rate results in a simple defocus in the imagery caused by a quadratic phase error. This defocus is very similar to that experienced in photographic systems. It is from this similarity Doppler-rate estimation/correction was given the name autofocus.

7.2.1   Motion errors

Cutrona [1975] showed that INS were more than an order of magnitude shy of the accuracy required to meet the strict SAS motion measurement requirements. Autofocus and micronavigation algorithms attempt to estimate unknown motion-errors and provide the accuracy unable to be obtained using an INS. Accurate motion estimation is important because unknown motion errors cause the primary limitation of short range ($< 100\,\mathrm{m}$) SAS imaging. Compensation to remove image-blurring caused by known motion errors is discussed in Chapter 6.

An alternative to using autofocus is to employ beacon positioning systems [Pihl et al. 2000; Pilbrow et al. 2002a, b; Shippey et al. 2001] where active or passive beacons are fixed to the seafloor; positioning using these beacons presents the same problem as autofocus using prominent scatterers [Shippey et al. 1998a]. Recent results indicate that general autofocus performed at least as well as beacon positioning [Shippey et al. 2001]. Positioning based on transponder echoes in beacon systems is identical to prominent point autofocus; albeit based on the assumption a beacon represents a high SNR point-scattering object.

Autofocus algorithms exploit redundancy in the data collection and scene to estimate motion—the phase distortion caused by path deviations is apparent in many parts of the scene while the phase due to the scene itself is random. The autofocus problem is like a typical system estimation problem: estimate the unknown system using a random noise input. The system in this case is the unknown distortion filter caused by the path-deviation and the random noise is the supposedly random scene. Autofocus algorithms therefore exploit the redundancy of the phase distortion across the scene to estimate motion errors.

7.2.2   Medium fluctuation (acoustic variability)

SAS imaging assumes that the imaging medium is homogeneous and that no variations occur in the speed of sound. However, the speed of sound in water varies with both time and position. This leads to the sonar pulses travelling along bent and/or twisted paths and to significant random phase delays [Urick 1975]. Moreover, these random speed of sound variations are space-varying, and image degradation results. Medium induced phase-distortions are the limiting factor in long-range SAS performance [Chang and Tinkle 2001].

The autofocus problem caused by fluctuations in the water column is subtly different from that caused by path errors. Where the motion induced phase distortions are common to all targets in the sonar beam at a given pulse, this is not so for fluctuation-caused distortions. Medium fluctuation affects separate parts of the image differently—i.e., both fast and slow water can be in the imaging beam. Moreover, the sonar beam

at a particular pulse "sees" many separate image sections, each with a different phase-distortion. Autofocus of medium fluctuation has constraints similar to those seen in the astronomical imaging problem. In particular, the spatial coherence length of the distortion is limited, limiting the length of aperture with common phase errors [Chang and Tinkle 2001; Knox and Thompson 1974]. Short coherence length, caused by medium fluctuation, thus limits the maximum size of the autofocused sea-floor patch [Chang and Tinkle 2001].

7.3   Revisiting the motion constraints

The amount of blurring SAS images suffer is dependant on both the magnitude and how rapidly varying the sway motions are. For this reason, the terminology used by Jakowatz et al. [1996] is chosen. The corrupting sway is described in terms of a polynomial function. Low-order sways represent deviations that vary slowly along the aperture and high-order sways more rapid variations. The limiting case of high-order sway is when the sway is uncorrelated i.e., is a white noise signal. Note that the relative sway order "seen" by a target is range-dependant since long range targets are in the imaging beam for a longer time period.

The commonly accepted motion constraint for blur-free synthetic aperture imagery is that the sway motions must be $< \lambda/16$ for the extent of the aperture. Whilst that constraint is viable for low-order sways, it is much tighter ($\lambda/60$) for high-order sways [Carrera et al. 1995, Chapter 5]. Moreover, the aperture length varies considerably in a typical wide-swath SAS, making the constraint difficult to grasp intuitively. A new motion constraint (giving the required pulse-to-pulse accuracy) is derived.

To calculate the motion constraints the blurring relation for a linear sway of [Callow et al. 2002a]

$$X(u) = \gamma u \tag{7.1}$$

is used. Under this model, the distorted image may be described by

$$\widetilde{f}(x,y) \approx f(x - \gamma y, y + \gamma x), \tag{7.2}$$

where $f(x,y)$ is the undistorted image and γ is a linear sway across the aperture. Extending the argument to successively smaller sub-apertures, the smallest possible sub-aperture (pulse-to-pulse differences) fits the same model.

Acceptable imagery is assumed when 95% of the image energy is contained within half the along-track resolution $\delta_y/2$. Note that this is quite different to the usual assumption that the aperture phase of the target must vary by less than $\pi/4$ along the aperture.

Applying the above constraint, γ is described by

$$\gamma < \frac{\delta_y}{2x}, \tag{7.3}$$

assuming without loss of generality that γ is a zero-mean random variable[2]. For 95% confidence, the linear sway standard-deviation σ_γ must be

$$\sigma_\gamma < \frac{\delta_y}{4x} \tag{7.4}$$

This represents a constraint on the sway per metre of aperture and is a system yaw limit on multiple-receiver sonars. To obtain acceptable imagery, ping-to-ping differences must also satisfy (7.3). Thus the standard-deviation of the ping-to-ping sway is given by

$$\sigma_{\text{sway}} < \frac{\delta_y}{4x} \Delta u, \tag{7.5}$$

where Δu is the along-track sample spacing. For a diffraction-limited image from a $D/4$ sampled sonar this can be expressed as

$$\sigma_{\text{sway}} < \frac{D^2}{32x}, \tag{7.6}$$

which gives $\sigma_{\text{sway}} \approx 70 \, \mu\text{m}$ ping-to-ping for the KiwiSAS-II sonar with an element size of $D = 0.325 \, \text{m}$, $\delta_y = 0.16 \, \text{m}$, and maximum range of $50 \, \text{m}$. These parameters give a linear sway standard-deviation of σ_γ of $\approx 0.8 \, \text{mm m}^{-1}$ which is equivalent to a yaw standard-deviation of $0.05°$ for a multiple-receiver equivalent of the KiwiSAS-II. The standard-deviation for the sway over the aperture can be calculated using the central limit theorem as

$$\sigma_{\text{sway}_L} = \sqrt{N_p} \sigma_{\text{sway}}, \tag{7.7}$$

where N_p is the number of pulses covering the aperture

$$N_p = \frac{x\lambda}{D\Delta u}. \tag{7.8}$$

For the KiwiSAS-II the sway standard deviation over the length of the aperture σ_{sway_L} $\approx 8 \, \text{mm}$ ($\lambda/6$). These constraints can be compared with the Cramér-Rao lower bounds (CRLBs) for the time-delay/phase-gradient sway estimators discussed in Section 7.6.1.

The advantage of (7.3) is that it expresses blurring constraints in terms of the design parameters of the sonar. If the system resolution is substituted into (7.3) the result is expressed in terms of the sonar's physical hardware and is independent of

---

[2]Not true in practice, but having a mean causes image shift not blurring.

frequency. Alternately (7.3) can be rewritten in terms of the SAS gain—the more gain in using SAS processing the stricter the motion requirements.

Constraints of this precision make autofocus challenging. Note however that the constraints given neglect beam-patterning and other effects; thus slackening the constraints by factors of 2–5 can often still result in acceptable imagery. The accuracy required always depends on the application.

7.4   Micronavigation/autofocus

Both micronavigation and autofocus are generic terms describing a number of algorithms that estimate path-deviations or medium fluctuation. Both families of algorithms have many similarities and only subtle distinctions.

Micronavigation is a term that arose to describe the integration of the redundant phase centres (RPC) algorithm (see Section 9.1) with INS and other aiding sensors onboard the imaging sonar [Bellettini and Pinto 2002]. The term micronavigation refers to any autofocus algorithm that operates to provide a real-time estimate of the path of the imaging platform. Micronavigation algorithms, such as RPC and shear average, typically exploit redundancy in the echo collection. In contrast, autofocus algorithms usually require certain statistical properties in the scene of interest. This distinction is easily blurred as echo redundancy and prior path-information can be used in typical autofocus algorithms and minor modifications to RPC allow for autofocus assumptions.

The use of micronavigation algorithms is likely to increase with the anticipated prevalence of Autonomous Underwater Vehicle (AUV) based SAS systems [Hagen et al. 2001; Sutton et al. 2002]. These systems require extremely high accuracy in positioning to operate effectively. Whilst an INS may not provide the accuracy needed for SAS, micronavigation algorithms designed to meet SAS measurement tolerances provide a boon to AUV operation. Using the information obtained using micronavigation for navigating an AUV is an interesting prospect.

7.5   Spotlight vs stripmap autofocus

The imagery obtained from a stripmap SAR or SAS is more difficult to autofocus than from spotlight systems. This extra difficulty is due to the blurring of stripmap images being space-*variant* [Gough et al. 2000b] whereas spotlight images have space-*invariant* blurring [Hawkins 1996; Wahl et al. 1994b]. Space-variant blurring means that standard, well researched, spotlight autofocus methods such as map-drift [Curlander and McDonough 1996], contrast optimisation [Blacknell et al. 1992], and PGA [Eichel and Jakowatz 1989; Jakowatz et al. 1996; Wahl et al. 1994a] are unable to be applied to stripmap data without modification. In the case of PGA, an extension to

the stripmap case does exist [Wahl et al. 1994b] and is referred to as phase curvature autofocus (PCA) [Pat 2000; Sutton et al. 2000]. The assumption of space invariance in the spotlight algorithms often makes the extension to stripmap systems challenging. The reason why space-invariant algorithms perform poorly in space-variant problems is straightforward. A space-invariant autofocus algorithm ensemble averages over all scatterers to estimate the path. However, in a space-invariant problem, all scatters have the same blurring. In a space-variant problem, each scatterer has a different blurring. Averaging many different path estimates results in a poor overall path estimate.

To overcome the problem caused by space-variant blurring, data from stripmap systems is often autofocused by segmenting the image into smaller along-track sections [Bonifant et al. 2000; Sutton et al. 2000; Thompson et al. 1999]. These sections have nearly space-invariant blurring and the application of normal spotlight autofocus algorithms yields some improvement. Usually a preprocessing step is required [Curlander and McDonough 1996; Pihl et al. 2000; Thompson et al. 1999] to massage the data into a form that SAR autofocusing algorithms can use. Each section of the data is then autofocused independently and the individual estimates combined to remove the distortion from the entire image [Bonifant 1999; Bonifant et al. 2000]. However, splitting the image in azimuth and autofocusing each sub-image independently is undesirable as useful information is discarded. In particular, path estimation from information at the edges of the sub-patch is inaccurate. Approaches that account for the spatial variance and retain information (such as the stripmap phase gradient autofocus (SPGA) algorithm—see Chapter 10) have better performance. To summarise, spotlight autofocus is a special case of stripmap autofocus where the space-variance of the problem is low. Spotlight algorithms often neglect the effect of space variant blurring (which must be considered for stripmap autofocus).

7.6   Autofocus techniques

This section describes the main techniques used in autofocus algorithms. Most algorithms (both micronavigation and autofocus) use some form of correlation or phase estimation. This section summarises these techniques from an autofocus perspective and in a common notation. Salient features of the algorithms are noted and related to the SAS autofocus problem.

7.6.1   Correlation of complex baseband signals

Many of the autofocus and micronavigation algorithms discussed later in the thesis have estimation techniques based on correlation of time-sequences or 2-D complex images. Most of the phase estimation kernels discussed later in the chapter perform some form of phase-only correlation for different along-track frequency bins. For example,

the Maximum Likelihood (ML) phase estimator described later in this chapter corre-
lates adjacent frequency bins—i.e., it estimates the time-delay difference at different
spatial frequencies. In addition, the correlation techniques discussed here demonstrate
a method of avoiding phase unwrapping requirements in phase estimation kernels.

This section presents a summary of the application of correlation to complex time-
sequences. The problem is treated as a correlation of a band-pass time signal, $p_0(t)$,
and its delayed echo signal, $p_1(t)$. The extension to multi-dimensional space follows
along similar lines [Bracewell 1986].

Given the two complex baseband time sequences, $p_0(t)$ and $p_1(t)$, their cross cor-
relation in time is given by

$$pp_{01}(t) = p_0(t) \star_t p_1(t),$$
$$= \int_{t'} p_0(t')p_1^*(t' - t)\, dt', \tag{7.9}$$

where $\star_t$ refers to the correlation operation. If $p_0(t)$ has a Fourier transform given by
$P_0(\omega)$ then (7.9) may be written [Bracewell 1986; Haykin 1994]

$$pp_{01}(t) = \mathcal{F}_\omega^{-1}\{P_0(\omega)P_1^*(\omega)\}. \tag{7.10}$$

In practice, the correlation obtained using (7.10) with FFT based processing gives a
different result to that of (7.9) due to circular convolution[3].

This type of correlation is usually used for estimating the time-delay between a
transmitted signal and the received echo signal. In this application $p_1(t) = p_0(t - t_0) +
n(t)$. From linearity, it follows that the spectra of the two signals are related,

$$P_1(\omega) = P_0(\omega)\exp\left(-j(\omega + \omega_0)t_0\right) + N(\omega). \tag{7.11}$$

Rewriting the correlation (7.10) gives

$$pp_{01}(t) = \mathcal{F}_\omega^{-1}\{P_0(\omega)(P_0^*(\omega)\exp\left(j(\omega + \omega_0)t_0\right) + N^*(\omega))\}. \tag{7.12}$$

Taking an inverse Fourier transform, (7.12) becomes

$$pp_{01}(t) = \mathcal{F}_\omega^{-1}\left\{|P_0(\omega)|^2 \exp\left(j(\omega + \omega_0)t_0\right)\right\} + n_p^*(t),$$
$$= pp_0(t - t_0)\exp\left(j\omega_0 t_0\right) + n_p^*(t), \tag{7.13}$$

where $pp_0(t)$ is the autocorrelation of $p_0(t)$, and $n_p^*(t)$ is the noise filtered by $p_0(t)$.
Time-delay estimators select the time where $pp_{01}(t)$ is maximum to give an estimate
of delay, $\widehat{t}_0$. The Cramér-Rao lower bound (CRLB) of the estimation may be derived

---

[3]Circular convolution problems may be mitigated by employing *guard bands* at the edge of the time
series to be correlated [Bracewell 1986]. This is more commonly called *zero padding*.

[Bellettini and Pinto 2002; Pinto et al. 2002] as[4]

$$\mathrm{CRLB} = \frac{1}{\sqrt{B_c \tau_{\mathrm{rep}}}} \frac{1}{\sqrt{1 + \frac{B_c^2}{12\omega_0^2}}} \frac{1}{\omega_0} \sqrt{\frac{1}{\mathrm{SNR}} + \frac{1}{2\,\mathrm{SNR}^2}}, \qquad (7.14)$$

where SNR is the average signal to noise ratio in the time sequences p_0 and p_1, $\sqrt{B_c \tau_{\mathrm{rep}}}$ (where B_c is the equivalent rectangular signal bandwidth in $\mathrm{rad\,s^{-1}}$ and τ_{rep} is the time between pulses) is the number of independent time samples, and ω_0 is the signal carrier frequency.

Amplitude-only envelope correlation

When correlating baseband signals and ignoring the phase of the correlation result, the correlation may be written

$$pp_{01_{\mathrm{env}}}(t) = \left| \int_{t'} p_0(t') p_1^*(t' - t)\,\mathrm{d}t' \right|, \qquad (7.15)$$

Writing $p_0(t)$ in the form of an amplitude and phase function

$$p_0(t) \equiv a(t) \exp\left(j\phi(t)\right), \qquad (7.16)$$

where the time-delayed echo signal is given by

$$p_1(t - t_0) \equiv a(t - t_0) \exp\left(j\phi(t - t_0) - j\omega_0 t_0\right), \qquad (7.17)$$

then the amplitude correlation given by (7.15) may be written

$$pp_{01_{\mathrm{env}}}(t) = \left| \exp\left(j\omega_0 t_0\right) \int_{t'} a(t') \exp\left(-j\phi(t')\right) a(t' - t + t_0) \exp\left(-j\phi(t' - t + t_0)\right)\,\mathrm{d}t' \right|. \qquad (7.18)$$

The information inherent in correlation phase is lost since only the amplitude of the correlation is used (i.e., the carrier phase shift, $\exp\left(j\omega_0 t_0\right)$, is not used). The CRLB is derived by setting the carrier frequency ω_0 zero in (7.14), giving

$$\mathrm{CRLB} = \frac{1}{\sqrt{B_c \tau_{\mathrm{rep}}}} \frac{2\sqrt{3}}{B_c} \sqrt{\frac{1}{\mathrm{SNR}} + \frac{1}{2\,\mathrm{SNR}^2}}. \qquad (7.19)$$

Envelope correlation provides a simple method for estimating the bulk time shift of a signal. However, the envelope correlation method is not as accurate as full complex correlation as it discards the information contained in the carrier phase shift.

---

[4]Making the substitution for the discrete number of independent samples, $\sqrt{N} \equiv \sqrt{B_c \tau_{\mathrm{rep}}}$, in (7.14) links the time estimation CRLBs with the autofocus phase CRLBs presented later in this chapter.

Narrow-band correlation

The time signals used in sonar imaging typically have a high carrier frequency compared with their bandwidth (i.e., have a Q $\gg$ 1.0). In this situation, correlation of the phase results in more accurate time-delay estimates than envelope correlation. As can be seen in the Fourier representation of the time-delay correlation (7.13), the correlation result is the auto-correlation of the transmitted signal, $pp_0(t)$, multiplied by a carrier phase shift $\exp{(j\omega_0 t_0)}$. The phase shift gives a direct estimate of the time-delay, t_0 (since the auto-correlation $pp_0(t)$ is a real-only signal if the baseband signals are symmetric about zero frequency [Haykin 1994; Shippey et al. 1998b]).

If the signals $p_0(t)$ and $p_1(t)$ have baseband spectra that are symmetric about zero frequency then the narrow-band or *phase-only* correlation of the signals is given by

$$pp_{01_{\text{phase}}} \approx \exp{(j\omega_0 t_0)} + n_p^*(t). \tag{7.20}$$

Thus a correlation time-delay can be measured as a carrier phase and estimated via

$$\begin{aligned}
\widehat{t}_0 &= \frac{1}{\omega_0}\mathbf{Arg}\left\{\int p_0(t)p_1^*(t)\,\mathrm{d}t\right\}, \\
&= \frac{1}{\omega_0}\mathbf{Arg}\left\{\int P_0(\omega)P_1^*(\omega)\,\mathrm{d}\omega\right\},
\end{aligned} \tag{7.21}$$

which is equivalent to the phase at zero correlation lag.

Narrow-band correlation suffers an ambiguity problem if large time-delays are to be estimated. Small time-delays are able to be measured but have a 2π ambiguity problem. This is discussed in more detail in following sections.

Narrow-band correlation, like envelope correlation, also has less accuracy than a full complex envelope correlation (or correlation of the modulated real signal). However the loss in accuracy from using narrow-band correlation is minor; particularly for high-Q systems. Narrow-band correlation has major computational savings over traditional correlation and these far outweigh the minor loss of accuracy if small time-delays are to be estimated.

Shear product

If $p_0(t)$ and $p_1(t)$ are sampled to give $p_0[n] \equiv p_0(n\Delta t)$ and $p_1[n] \equiv p_1(n\Delta t)$, the time-delay, t_0, may be estimated by

$$\widehat{t}_0 = \frac{1}{\omega_0}\mathbf{Arg}\left\{\frac{1}{N}\sum_{n=0}^{N-1}p_0[n]p_1^*[n]\right\}, \tag{7.22}$$

which is the average of the sheared product, or *shear average*[5]. This is the discrete implementation of the phase only correlation described in the previous section.

Delay estimation via (7.22) is variously known as the *shear average* or *sheared-product* and has been used in autofocus for both SAS and SAR [Fienup 1989; Johnson et al. 1995] and also in astronomical imaging[6].

Quasi-narrowband framework

Phase only correlations, while giving locally accurate results, suffer from phase unwrapping problems. This results in ambiguity in the results[7]. Sub-cycle time-delays are estimated well but an estimate can have an integer number of 2π radians error. Full correlation does not suffer this problem.

One approach is to consider that discrete complex correlation is described by [Shippey et al. 2001]

$$pp_{01}[m] = \frac{1}{N} \sum_{n=0}^{N-1} p_0[n]p_1^*[n-m], \tag{7.23}$$

where $pp_{01}[m]$ is the discrete version of $pp_{01}(t)$. The time delay estimate, $\widehat{t}_0$, can be estimated from either the peak of the cross-correlation $pp_{01}[m]$ or the phase at zero correlation lag $pp_{01}[0]$. The phase of the correlation gives an high-resolution, ambiguous indication of the shift and the amplitude peak gives a low-resolution, unambiguous shift estimate. If the bandwidth and SNR are such that the envelope correlation (7.23) gives sufficient accuracy to avoid the cycle ambiguity in phase estimation, the coarse delay may be estimated using envelope correlation and the fine delays estimated using phase correlation [Pinto et al. 2002]. This is sometimes known as quasi-narrow-band processing [Shippey et al. 1998b]. Bathymetric imaging with vertically separated receiver arrays uses a similar technique to avoid phase ambiguity [Bellettini and Pinto 2002].

The CRLB for an unambiguous narrow-band correlation is derived by setting the bandwidth B_c to zero in (7.14) (this does not apply to the $\sqrt{B_c\tau_{\text{rep}}}$ term however) and is given by

$$\text{CRLB} = \frac{1}{\sqrt{B_c\tau_{\text{rep}}}} \frac{1}{\omega_0} \sqrt{\frac{1}{\text{SNR}} + \frac{1}{2\,\text{SNR}^2}}. \tag{7.24}$$

Because the full correlation bound (7.14) and the narrow-band bound (7.24) are very similar, the quasi-narrow-band technique is often used [Callow et al. 2001c; Shippey et al. 1998a; Wang et al. 2001].

---

[5]It is also the ML phase estimator discussed in Section 7.6.2.

[6]In astronomical imaging it is equivalent to the well-known Knox-Thompson method [Knox and Thompson 1974].

[7]Caused by the basic problem of phases being measured modulo 2π.

7.6.2   Phase estimation kernels

Phase estimation kernels have many similarities with phase only correlation techniques. For example, the phase difference between pulses may be treated as a narrow-band estimate of the relative time-delay between those pulses [Wahl et al. 1991]. Phase estimation kernels provide methods for narrow-band (phase-only) time-delay estimation.

Autofocus algorithms often need to estimate the phase error across the aperture $\phi(k_y)$ or between sonar pulses—effectively determining the time shift in the image. Many different phase estimation kernels have been used for this purpose. Some of the more widely used phase estimation kernels are outlined in this section.

The performance of autofocus algorithms relies heavily on the phase estimation technique chosen. There are a number of PGA variants (such as WPGA [Ye et al. 1999], and FLOS-PGA [Tsakalides and Nikias 2001]) offering performance improvements in select conditions. The variants often only change the PGA phase estimation kernel. Phase kernels that better match the known (or estimated) statistical properties of the image can yield significant performance gains.

To aid in the readability of this section a modified, discrete notation is used, similar to that of Jakowatz and Wahl [1993]. In this notation, n is the across-track index, so that $x \equiv n\Delta x$, and l is the along-track spatial-frequency index, so that $k_y \equiv l\Delta k_y$ where Δx and Δk_y are the sample spacings in cross-track and along-track spatial-frequency respectively. p is the along-track position index so that $u \equiv p\Delta u$. In this notation the discrete range-Doppler domain version of the blurred image may be written

$$f(x, k_y) = f(x\Delta x, l\Delta k_y), \tag{7.25}$$

$$= g[n, l] \tag{7.26}$$

where the change of letter from $f(x, k_y)$ to $g[n, l]$ signals the discrete conversion.

Differentiation kernel

The original phase gradient estimation kernel for SAR autofocus was published for use with PGA [Eichel et al. 1989]. The kernel in discrete notation is given by

$$\dot{\phi}[l] = \frac{\sum\limits_{n=0}^{N-1} \text{Im}\left\{\dot{g}[n,l]g^*[n,l]\right\}}{\sum\limits_{n=0}^{N-1} |g[n,l]|^2}, \tag{7.27}$$

where $\dot{\phi}[l]$ is the average phase gradient, $\dot{g}[n, l]$ is calculated using Fourier differentiation as

$$\dot{g}[n, l] = \mathcal{F}_{p\to l}\left\{j2\pi p\ \mathcal{F}^{-1}_{l\to p}\{g[n,l]\}/P\right\}, \tag{7.28}$$

and P is the number of along-track samples. The phase is then estimated by integrating the phase gradient via

$$\varphi[l] = \int \dot{\phi}[l]\,\mathrm{d}l, \tag{7.29}$$

using numerical integration techniques.

The performance of this kernel was demonstrated to be inferior to that of the Maximum Likelihood (ML) kernel (see next two sections). The reason for poor performance is a bias at low signal to noise ratios [Jakowatz and Wahl 1993]. This bias is the result of a high signal to clutter ratio assumption in the derivation of the estimation kernel [Eichel and Jakowatz 1989; Jakowatz and Wahl 1993]. PGA often operates on imagery where the signal-to-clutter ratio is poor (-10 dB to 0 dB) and so it is better to use the (unbiased) ML estimator instead.

Eigenvector kernel

The eigenvector phase estimation kernel first appeared for PGA in [Jakowatz and Wahl 1993] as an improvement to the differentiation-based PGA kernel (see previous section). Similar techniques have appeared in ultrasonic imaging [Rachlin 1990] and adaptive beam-forming for sonar [Nielsen 1991, page 261]. Jakowatz and Wahl [1993] show that the eigenvector-based estimator is Maximum Likelihood (ML) and demonstrate the performance increase over the differentiation-based kernel. The ML phase estimation kernel used in shear average and the Knox-Thompson techniques is a special case of the eigenvector estimator.

The eigenvector phase estimator is based on taking the phase of the eigenvector corresponding to the largest eigenvalue of the (along-track) covariance matrix[8]. Under the PGA system model this is directly equivalent to the distortion phase that causes blurring.

The performance of the eigenvector kernel is improved by utilising higher-order measurements (i.e., by calculating larger covariance matrices). The CRLB for phase estimators[9] based on using the M^{th} order covariance matrix is [Jakowatz and Wahl 1993]

$$\mathrm{CRLB} = \sqrt{\frac{1 + M\beta}{MN\beta^2}}, \tag{7.30}$$

$$= \frac{1}{\sqrt{N}}\sqrt{\frac{1}{M\beta^2} + \frac{1}{\beta}}, \tag{7.31}$$

---

[8]The covariance matrix is defined for fixed n as $E\{\mathbf{g_n}\mathbf{g_n^H}\}$ where $E\{\}$ is the statistical expectation operator and $\mathbf{g_n} \equiv [g[n,0], g[n,1], \ldots, g[n,L]]^T$. Singular Value Decomposition (SVD) is a particularly good numerical technique for finding the eigenvector corresponding to the largest eigenvalue [Nielsen 1991, page 268].

[9]Time/position estimation CRLBs may be obtained by dividing those presented here by the appropriate wavenumber (ω or $2k$).

where β is the SNR for the phase estimation (often the signal-to-clutter ratio rather than the signal-to-noise ratio) and N is the number of independent time samples (equivalent to $B_c\tau_{\text{rep}}$). Calculating a M^{th} order covariance matrix requires $\binom{2}{M}$ phase-only correlations and has the benefit of a CRLB proportional to $1/\sqrt{\beta^2 M}$ for small β.

The ML phase estimator used in the Knox-Thomson and shear average methods corresponds to the eigenvector method when only a second order covariance matrix is used—i.e., only adjacent frequency bins are used. Substituting for covariance order $M = 2$ for the ML estimator (7.30) becomes

$$\text{CRLB} = \frac{1}{\sqrt{N}}\sqrt{\frac{1}{\beta} + \frac{1}{2\beta^2}}, \tag{7.32}$$

which is identical to the CRLB for phase only time-delay estimation (7.14).

In the SAR situations where PGA is typically used, the increased estimation accuracy of high-order phase estimation cannot be justified on the basis of computational cost [Jakowatz and Wahl 1993]. Even with a poorer low-order phase estimator, the low computational cost of iterating PGA means that an estimator based on only adjacent frequency bins is computationally cheaper (even though it takes more iterations). However, high-order phase estimation gives benefits when β is very small, such as for clutter-dominated scenes. This is usually the case in scenes for which PGA fails to give acceptable autofocus. Use of high-order phase estimation improves PGA's robustness. Moreover, SAS autofocus necessitates a much larger iteration cost [Hawkins 1996] and so higher-order phase estimation is preferable.

A simplification of the eigenvector estimation kernel can be derived for M=3 as (see Appendix E)

$$\begin{bmatrix} \Delta\phi_1[l] \\ \Delta\phi_2[l] \end{bmatrix} = \begin{bmatrix} \frac{2}{3} & \frac{1}{3} & -\frac{1}{3} \\ \frac{1}{3} & \frac{2}{3} & \frac{1}{3} \end{bmatrix} \begin{bmatrix} -\mathbf{Arg}\left\{\left(\frac{1}{N}\sum_{n=0}^{N-1} g[n,l]g^*[n,l+1]\right)\right\} \\ -\mathbf{Arg}\left\{\left(\frac{1}{N}\sum_{n=0}^{N-1} g[n,l+1]g^*[n,l+2]\right)\right\} \\ -\mathbf{Arg}\left\{\left(\frac{1}{N}\sum_{n=0}^{N-1} g[n,l]g^*[n,l+2]\right)\right\} \end{bmatrix}, \tag{7.33}$$

where the phase, $\varphi[l]$, is obtained from the phase differences, $\Delta\phi_1[l]$, $\Delta\phi_2[l]$, as

$$\varphi[l] = (\varphi[l-1] + \Delta\phi_1[l]) + (\varphi[l-2] + \Delta\phi_2[l]). \tag{7.34}$$

This form of the eigenvector kernel is reminiscent of the ML estimator and those estimators based on the image bispectrum.

It is possible to show (see Appendix E) that eigenvector-based phase estimators (of order ≥ 3) implement phase-closure (a particularly powerful phase-retrieval technique used in astronomical imaging [Roddier 1986]). Eigenvector-based phase estimators are

also equivalent to bispectral or trispectral (cumulent or triple and quadruple correlation) phase estimators (see Appendix F). Interestingly, eigenvector decompositions are closely related to the Karhunen-Loéve (KL) [Van Trees 1968] technique used in image processing to generate a maximum entropy image; how closely this relates eigenvector-based phase estimation to entropy-based iterative autofocus [Morrison and Munson 2002; Xi et al. 1999] remains to be seen.

Eigenvector-based techniques provide efficient, accurate, phase estimation for autofocus. SAS autofocus should use higher order eigenvector estimators to improve performance particularly as high order estimators improve phase estimation performance in the heavy clutter scenes expected in SAS imagery [Billon and Pinto 1995]. However, the WPGA estimator (see later section) should be used in preference to a high order eigenvector kernel since it is equivalent to an ∞-order eigenvector estimator with reduced computational load.

Maximum likelihood (ML) estimation kernel [Jakowatz and Wahl 1993]

The ML estimator is identical to those used in spotlight shear average [Fienup 1989], spatial correlation autofocus [Attia and Steinberg 1989], and Knox-Thompson speckle-interferometry [Knox and Thompson 1974], and is similar to that used in phase difference autofocus (PDA—see Section 8.5). This estimator operates by correlating adjacent frequency bins (or adjacent blocks of frequency bins in PDA) and can be interpreted as estimating the difference in time-delay between the adjacent frequency bins (see Section 7.6.1).

The Maximum Likelihood (ML) phase estimator calculates the phase-error across the aperture by comparing adjacent spatial-frequency bins,

$$\Delta\phi_1[l] = \mathbf{Arg}\left\{\frac{1}{N}\sum_{n=0}^{N-1}\left(g[n,l]g^*[n,l+1]\right)\right\}. \tag{7.35}$$

The error across the aperture, $\varphi[l]$, is then calculated by integrating the phase differences,

$$\varphi[l] = (\varphi[l-1] + \Delta\phi_1[l]). \tag{7.36}$$

Fractional low order statistics (FLOS) estimator

The phase estimator from FLOS-PGA [Tsakalides and Nikias 2001] is a good example of a *weighted* phase estimation approach. The estimator is designed for scenarios where the clutter amplitudes are not able to be modelled using Gaussian statistics (in particular those with heavy-tailed statistics). The estimator works by mapping signal amplitudes through a non-linear function to alter the PGA weightings and is given by

[Tsakalides and Nikias 2001]

$$\Delta\phi_1[l] = \mathbf{Arg}\left\{ \frac{1}{N} \sum_{n=0}^{N-1} \left(g^{<p>}[n,l] g^{*<p>}[n,l+1] \right) \right\}, \tag{7.37}$$

where[10]

$$g^{<p>}[n,l] = |g[n,l]|^{p-1.0}\, g[n,l]. \tag{7.38}$$

For $p = 1.0$, the altered phase estimation kernel (7.37) is the same as the ML kernel. The technique may be viewed as pre-whitening the signal so that the clutter may be treated as Gaussian noise. Another possible interpretation, when $p < 2$, is that the effect of strong-scattering targets is reduced somewhat. Improvements resulting from the technique have been shown in field collected SAR data [Tsakalides and Nikias 2001].

Weighted PGA (WPGA) estimator

WPGA uses the *calculated* target-to-clutter ratio in weighting the contributions from each individual PGA phase estimate [Ye et al. 1999]. PGA assumes that the clutter is a Gaussian process with constant variance and thus weights using the target amplitude. Note that this assumption is often violated in SAS. Using the estimated clutter variance for weighting allows a broader range of scenes to be used and improves performance on clutter-like scenes [Ye et al. 1998, 1999]. WPGA is also the Minimum error Variance (MV) phase estimator for PGA with any noise process. By estimating the clutter variance and noise power, WPGA avoids needing to model the clutter process. This improves performance where the traditional models break-down. Thus the WPGA estimator is superior to both FLOS-PGA and traditional PGA.

WPGA has the additional benefit of lowering the weighting of strong-scattering extended targets. By estimating the noise power, the noise contributed due to the extended nature of targets is considered and their relative weighting lowered. This makes the WPGA estimator robust to model variations caused by extended targets. This is particularly important in SAS imaging to improve performance when strong-scattering extended targets exist.

WPGA also benefits by avoiding the estimation of phase differences. Phase-difference (and phase gradient) estimation is an inherently noisy process, lowering the accuracy of the phase estimator [Chan and Yeo 1998]. The WPGA estimation process is performed by averaging unwrapped phases from each range-bin[11] instead of averaging phase gradients (which were originally used to avoid phase wrapping problems).

---

[10]The mapping in [Tsakalides and Nikias 2001] is the complex conjugate of the mapping presented here.

[11]Chevillon et al. [1998] describe a similar phase estimator where unwrapped phases for each frequency are averaged.

These phases are generated by unwrapping to the phase of the range-bin with the highest SNR, a *local phase unwrapping* [Ye et al. 1999]. Local phase unwrapping avoids unnecessary noise injection of differentiation and provides more accurate estimation. The CRLB for the WPGA estimator is

$$\text{CRLB} = \frac{1}{\sqrt{N\beta}}, \tag{7.39}$$

which is better than the ML or eigenvector phase estimation kernels for low SNRs (it is equivalent to an ∞-order eigenvector kernel).

WPGA currently provides the best way of performing phase estimation in a PGA framework and has improved results when compared with traditional PGA [Ye et al. 1999]. Both of the improvements WPGA offers, noise power estimation and averaging unwrapped phases, are effective and enhance PGA performance at low signal-to-clutter ratios. The noise power estimation also avoids some of the biasing effects that strong-scattering extended targets can cause. These benefits make it the estimator of choice for SAS autofocus. However, the stripmap nature of SAS autofocus prevents operation of WPGA as published—the space-variance stripmap autofocus removes some of the redundancy that local phase unwrapping relies on. More research is required to extend the use of this promising phase estimator to SAS autofocus.

Knox-Thompson

The Knox-Thompson[12] method is a popular phase estimation technique used in astronomical imaging to estimate the object phase both in 1-D and 2-D [Knox and Thompson 1974]. Where autofocus for SAR/SAS takes an ensemble of range-samples to get the common path-deviation, astronomical phase estimation uses an ensemble of random speckle images[13] and generates an estimate of the common phase of the underlying scene.

Mathematically this method is identical to the ML phase estimator. The method begins by taking the along-track Fourier transform of the image and then multiplies different frequencies to get an idea of the corrupting phase over the aperture.

Knox-Thompson estimates the phase-gradient across the aperture by,

$$\Delta\phi_a[l] = \mathbf{Arg}\left\{\frac{1}{N}\sum_{n=0}^{N-1}\left(g[n,l]g^*[n,l+a]\right)\right\}, \tag{7.40}$$

for a given integer spacing[14], a. The error across the aperture, $\varphi[l]$, is calculated by cumulatively summing the phase differences, taking into account the frequency

---

[12]Also described as shearing interferometer in some literature.
[13]Each speckle image is an image of the scene corrupted by some random phase-error.
[14]For $a = 1$ this method is the same as the ML phase estimator.

separation a. Usually in the case of Knox-Thompson, the data is integrated out from the reliable low-frequency data toward the less reliable high-frequency data

$$\varphi[l] = (\varphi[l - a] + \Delta\phi_a[l]). \tag{7.41}$$

The Knox-Thompson method was very successful in astronomical imaging of both 1-D and 2-D data. Knox-Thompson does however rely on a recursive phase estimation, estimating high frequency phase information from low frequency phase information (the cumulative summation of (7.41)). This recursion causes errors to accumulate as the integration proceeds and high frequency phase information becomes unreliable. The problem also affects the ML phase estimation discussed earlier. The bispectrum method (see next section) can alleviate this problem using a least squares frequency estimation approach.

Bispectrum method

The bispectrum method, again from the astronomical imaging community, is an extension of the Knox-Thompson presented in the previous section. The method estimates the Knox-Thompson phases over all possible separations (i.e., it varies a in (7.40)). This allows many possible estimates of the corrupting phase at a particular spatial frequency or along-track position. Mathematically, the average bispectrum is expressed as follows [Lohmann et al. 1983],

$$\Phi_{BS}[l, a] = \frac{1}{N} \sum_{n=0}^{N-1} (g^*[n, l + a]g[n, a]g[n, l]). \tag{7.42}$$

The phase across the aperture is either obtained using recursive techniques (similar to the simple integration used in KT and shear average) or a slower least squared error based solution (that requires phase unwrapping).

The recursive phase estimation algorithm often used in the bispectrum method estimates the aperture phase via [Bartelt et al. 1984]

$$\varphi[l] = \varphi[j] + \varphi[l - j] - \mathbf{Arg}\left\{ \frac{1}{N} \sum_{n=0}^{N-1} (g^*[n, l]g[n, -j]g[n, l + j]) \right\}, \tag{7.43}$$

which gives j different estimates for a particular $\varphi[l]$ allowing averaging. For the j^{th} estimate, the starting phase, $\varphi[j]$ is chosen to be [Lohmann et al. 1983]

$$\mathbf{Arg}\left\{ \frac{1}{N} \sum_{n=0}^{N-1} (g[n, -j]) \right\}. \tag{7.44}$$

There are many possibilities for different averaging schemes using the bispectrum data. Freeman et al. [1988] discusses some of the popular choices.

In a coherent system, it is best to average the (separately obtained) phases for each across-track index (i.e., the aperture phases calculated via (7.43) rather than the complex vectors). The reason is that the phase at any given range bin for a fixed separation, a, is the same as Knox-Thompson phase for that separation multiplied by a constant phase factor $g[n, a]$. This gives a constant phase offset to the aperture phase that varies with each range bin. If this is left uncorrected the averaging of (wrapped) phases fails.

Some additional modifications have appeared to allow bispectrum to work on the complex-valued imagery obtained in spotlight SAR. The most notable of these is the work done by Berizzi et al. [1996] on cumulant-based autofocus and the method used by Nikias described in the next section.

Note that the bispectrum method uses the same information as the eigenvector kernels (see Appendix F). Eigenvector kernels are preferred to the bispectrum method however the WPGA estimator outperforms both techniques.

Nikias method (HOSPA)

The phase estimation method used by Nikias et al. [2000] appears to be another extension of bispectrum to deal with complex data. The algorithm starts by calculating a slice through the complex trispectrum [Nikias and Petropulu 1993]

$$\Phi_{TS}[l, a] = \frac{1}{N} \sum_{n=0}^{N-1} (g^*[n, l+a]g^*[n, 0]g[n, l]g[n, a]). \tag{7.45}$$

In the patent [Nikias et al. 2000], the frequency separation, a, is fixed at 1 which makes the method the same as a weighted version of shear average (Section 8.2). The phase estimates are obtained in a recursive fashion using a modified version of the Bartelt-Lohmann-Wirntzer algorithm (7.43) given below

$$\varphi[l] = \varphi[j] + \varphi[l-j] - \mathbf{Arg}\left\{\frac{1}{N} \sum_{n=0}^{N-1} (g^*[n, 0]g^*[n, l]g[n, -j]g[n, l+j])\right\}, \tag{7.46}$$

where $\varphi[j]$ is chosen[15] for each separate estimate of $\varphi[l]$ to be

$$\mathbf{Arg}\left\{\frac{1}{N} \sum_{n=0}^{N-1} (g^*[n, 0]g[n, -j])\right\}. \tag{7.47}$$

The comment in the previous section regarding averaging of the bispectrum data still applies here. The bispectrum should not be directly averaged because the phase

---

[15]Originally in Nikias et al. [2000] it was incorrectly chosen to be $\mathbf{Arg}\left\{\frac{1}{N} \sum_{n=0}^{N-1} (g[n, -j])\right\}$ which gives an extra range-variant phase term of $g^*[n, 0]$. cf.with (7.43).

term in (7.47) varies with range (unknown linear shifts and unwrapping phases).

This method appears to be a modified bispectrum method and similar performance is expected. Again, as with the bispectrum method, eigenvector kernels are preferred (and WPGA over eigenvector kernels).

7.6.3   Gradient versus curvature phase kernels

The majority of the phase kernels discussed in this thesis are *phase-gradient* based—i.e., they are based on calculating phase gradients via differentiation or by some form of discrete differencing. However, some stripmap algorithms require phase curvature (phase double-difference) estimates to operate when random linear-phase trends are present. This section briefly summarises the similarities and differences between phase gradient and phase curvature estimation kernels.

Phase curvature estimators give the rate-of-change of the phase gradient. Phase curvature calculation is often a matter of running a gradient estimator on a previous phase gradient estimate. Only one kernel has been regularly used for phase curvature estimation in SAS, the ML kernel from Pat's thesis [Pat 2000]. Thus, the phase curvature, $\Delta^2\phi_1[l]$, for the ML estimator is given by

$$\Delta^2\phi_1[l] = \mathbf{Arg}\left\{\frac{1}{N}\sum_{n=0}^{N-1}\left(g[n,l-1](g^*[n,l])^2 g[n,l+1]\right)\right\}. \tag{7.48}$$

Given the phase curvature estimates the phase can be estimated via

$$\varphi[l] = (\varphi[l-1] + \Delta\phi_1[l]), \tag{7.49}$$

where

$$\Delta\phi_1 = (\Delta\phi_1[l-1] + \Delta^2\phi_1[l]). \tag{7.50}$$

The ML kernel (7.48) was derived by taking the shear average of the shear-product (ML phase gradient of the complex phase difference signal). Although only the ML kernel has been used regularly, most of the phase gradient techniques discussed in the previous section can be adapted to estimate phase curvature in similar fashion.

Phase (and thus sway) estimation from phase curvature requires a double integration. This integration results in an extra unknown integration constant relative to phase gradient techniques. Viewed from a sway estimation perspective, this results in reduced reliability of low-order phase (sway) errors[16]. The double integration needed by phase curvature estimators reduces accuracy and should be avoided.

---

[16]The low-order degradation is most easily seen by considering the combined differentiation/integration operation as a filter that increases low-order noise. Repeating the operation as phase curvature estimation requires makes low-order phase results poorer.

7.7 Overview of current SAS autofocus algorithms

This section gives an overview of the techniques currently used for SAS autofocus. Each of the major SAS autofocus divisions[17] (echo-correlation based, global optimisation, and phase-gradient based) are discussed and related to similar techniques in other fields.

7.7.1 Echo-correlation based autofocus

Echo-correlation based autofocus techniques rely on redundancy in the collected data to provide reliable motion estimates. SAS data must be oversampled in along-track for the algorithms to operate. Extensions relying on typical scene-based autofocus information can remove this requirement at the expense of accuracy and algorithmic simplicity.

The simplest of the echo-correlation techniques, shear-averaging, is a successful autofocusing technique for single receiver stripmap SAS [Johnson et al. 1995]. Shear average has been used in SAR autofocus [Fienup 1989, 2001], speckle-interferometry [Ghiglia and Mastin 1989], and real-array radar [Attia and Steinberg 1989]. Variations of the shear-average method have been employed in the field of medical ultrasound, where a very similar autofocus problem (aberration detection) exists [Behar 2002; Krishnan et al. 1997; Ng et al. 1997], and also in the SAS receiver calibration problem [Banks and Griffiths 2002; Douglas et al. 1992].

Shear average requires that the echo signals recorded from adjacent pulses contain redundancy. The redundancy is exploited to obtain navigation information. The requirement for echo redundancy means that the SAS must take along-track samples closer than $D/2$ [Johnson et al. 1995]. The other assumption shear average requires is that the sea floor may be treated as many uncorrelated point sources of similar amplitude, i.e., the pulse-compressed signal is *delta-correlated* [Fienup 1989]. In low-contrast images, the assumption of a delta-correlated signal is valid and reasonable autofocus occurs both in simulation and in practice [Callow et al. 2000, 2001b, c; Shippey et al. 1998a]. Shear average makes a useful algorithm for quickly estimating the bulk motion errors present in an image before using other algorithms to finish the autofocus [Fortune et al. 2001b]. Also, because of its low computational cost, shear average makes the basis of a good micronavigation algorithm. In the presence of a strong scattering target, the assumption of delta-correlation breaks down and the autofocus is biased. This degradation is also a problem in spotlight SAR autofocus where Doppler-centroid (clutter-lock) estimates are required [Berizzi et al. 1997]. The degradation may be reduced by using very bland sea-floor sections [Shippey et al. 1998a] or by reducing the influence of strong scattering targets [Callow et al. 2001c]. Another limitation of

---

[17]Sub-aperture correlation is discussed as a sub-division of echo-correlation.

the shear average method as originally proposed is the limit on the amount of phase-distortion measurable. Methods of overcoming the phase distortion limit for wide-band systems have been presented [Callow et al. 2001c; Shippey et al. 1998a].

The redundant phase centres algorithm (RPC), also known as the displaced phase centres (DPC) algorithm, and the displaced phase centre antenna (DPCA) algorithm, [Bellettini and Pinto 2002; Billon and Pinto 1995; Groen and Sabel 2002; Pinto et al. 2002; Raven 1981; Sheriff 1992; Tonard and Brussieux 1997] is similar to shear average in operation and is used for multiple-receiver SAS systems. This algorithm is used in most multiple-receiver SAS systems. The only major disadvantage of the algorithm is that the multiple-receiver array must travel slower than the sampling soft limit dictates (since RPC, like all echo-correlation based methods, relies on collecting redundant echo data spaced closer than $D_{\mathrm{array}}/2$). Algorithm accuracy improves as the sonar travels even more slowly although undesirable platform motion is thought to increase at low sonar speeds[18] [Johnson et al. 1995]. The RPC algorithm exploits echo-correlation from ping-to-ping and anything that unduly disrupts echo-correlation, such as signal multi-path effects and temporal phase de-correlation, has a detrimental impact on performance. A variation of RPC exists that exploits correlation of beams (i.e., beam-formed images) formed using the redundant elements [Tonard and Brussieux 1997]. This affords some accuracy improvements when the echo signals have a strong directional component.

The RPC algorithm has other variations that allow operation with non-overlapping arrays (non redundant data collection—the echoes need not be correlated) such as the *cascade* algorithm [Douglas and Lee 1993a; Silkaitis et al. 1994] and the sub-aperture image correlation techniques [Calloway et al. 1991; Groen and Sabel 2002; Nahum 1998; Shippey et al. 2001]. These techniques no longer rely on echo-correlation but upon the typical assumptions of autofocus techniques (although operation improves if echo-correlation exists). They can also suffer the same type of autofocus biasing with strongly scattering extended targets. The techniques are the stripmap equivalent of SAR's sub-aperture correlation autofocus techniques (see Chapter 8).

Echo-correlation based autofocus and micronavigation techniques provide successful methods for quickly estimating bulk sonar motion errors. Path estimation based on ping-to-ping estimation tends to leave low-order errors that cause residual blurring in the images. More complex algorithms, like the cascade algorithm which estimate motion based on more than adjacent pulses, can reduce this problem. However these techniques use the same scene-based prior information as other autofocus techniques for obtaining path information (and require more computation). Phase gradient and global optimisation are better for autofocus than complex image correlation based autofocus methods.

---

[18]The optimum sonar sampling rate for RPC operation is $D_{\mathrm{array}}/4$. Sampling more often requires improved ping-to-ping accuracy since the errors are cumulative [Bellettini and Pinto 2002].

7.7.2   Global optimisation based autofocus

Global optimisation techniques, such as the iterative estimation techniques which attempt to maximise an image-likelihood functional, have little history of use in SAS imaging. After a relatively successful introduction in spotlight SAR (where contrast optimisation has be shown to be optimal for images with compact support [Rachlin 1990]) [Blacknell et al. 1992; Gough and Lane 1998; Isernia et al. 1996a, b; Xi et al. 1999], a relatively straightforward application to stripmap data was anticipated. The space-variant nature of stripmap data has made the application of optimisation techniques difficult [Fortune et al. 2001a, b; Gough and Lane 1998].

The problems facing this type of autofocus algorithm are choosing a suitable functional to optimise (image-likelihood measure) and avoiding the problem of local solutions (solutions other than the true optimum) in efficient fashion. A suitable image-likelihood metric has been elusive, with image contrast-based [Fortune et al. 2001a, b] and local Fourier magnitude uniformity (wavefront sensing based) [Callow et al. 2002a] proving inadequate. The use of alternate autofocus methods prior to global optimisation reduces the impact of local solutions in the optimisation and thus allow a reduction in computational load.

Global optimisation provides a promising framework in which to perform SAS autofocus. Ideally, any useful information about the SAS system can improve autofocus performance. Currently, the computational burden of optimisation type techniques make researching them difficult. Advances in computer technology coupled with the flexible nature of global optimisation will make it a key future autofocus technique.

These techniques are not discussed further in this thesis and the interested reader is referred to [Callow et al. 2002a; Fortune et al. 2001b; Gough and Lane 1998; Isernia et al. 1996a; Morrison 2002] for more information.

7.7.3   Phase gradient/curvature based autofocus

Phase gradient autofocus (PGA) is an extremely successful autofocus technique used in spotlight SAR [Eichel et al. 1989; Eichel and Jakowatz 1989; Jakowatz et al. 1996; Wahl et al. 1994a]. However it is unable to be used with the wide-beam, wide-band, stripmap imagery prevalent in the SAS field without extensive modification.

One possible approach toward allowing stripmap PGA operation is to subdivide the image into small patches were the blurring is space-invariant and use traditional PGA on each of the image patches. The main task external to PGA is calculating the overall path estimate from each of the image patches. Mosaic PGA (mPGA) is one such algorithm that performs subdivision and has been successfully demonstrated on stripmap SAS imagery [Bonifant 1999; Bonifant et al. 2000]. As similar approach is described by Thompson et al. [1999].

Another extension of PGA for stripmap autofocus is phase curvature autofocus (PCA). The technique was first proposed by Wahl et al. [1994b] to extend PGA to narrow-band, narrow-beam stripmap systems. As demonstrated in Hawkins [1996], PCA as originally published is suitable for systems with no range curvature; this same work describes modifications required to ameliorate the range curvature problem. In practice, PCA has been found difficult to use [Pat 2000; Sutton et al. 2000] and is not in wide-spread use. The reason for this is likely to be the double integration of phase curvature to get a measure of the phase error. A similar problem was found in an implementation of rank one phase estimation (ROPE) [Snarski 1996] an autofocusing technique with a similar kernel to PGA (and hence PCA). Some research has been conducted into methods of alleviating this difficulty by fitting a limited order polynomial basis in the processing of PCA [Pat 2000].

The phase matching autofocus (PMA) algorithm [Gough et al. 2000a] gives a PGA/PCA hybrid algorithm aimed at alleviating the difficulties caused by the double integration of PCA. The algorithm exploits some additional information about the Doppler spectrum of the signal to remove one of the integrations of PCA; this results in a phase gradient type algorithm.

All of the phase-gradient based algorithms assume that the average target Fourier phase in the scene is zero. This is equivalent to having a delta-correlated scene. Note that this does *not* necessarily require that the scene consist only of point targets[19]. Other autofocus techniques assume constant target position with varying imaging direction. This implies the scene looks identical from different imaging angles and is equivalent to having zero average target Fourier phase. This makes phase gradient based algorithms equivalent[20] to their image correlation counterparts.

The main benefit of phase gradient based algorithms is that extremely fast operation is possible. The reason for this is that (at least in typical SAR scenes) only 5% of the scene contains > 90% of the energy [Carrera et al. 1995, Chapter 6]. Using only that 5% of the scene for autofocus, computational load is reduced significantly with only minor performance degradation. Another significant benefit of phase gradient based algorithms is their inherent flexibility. It is difficult to use the weighting estimation techniques such as WPGA and FLOS-PGA in correlation-based methods.

7.8   Summary

Autofocus and micronavigation techniques to estimate platform path and medium fluctuation are an essential aspect of SAS imaging. The motion and fluctuation constraints that must be met for successful SAS imaging have been summarised. Most importantly

---

[19]Stoyle [1998] describes the electromagnetic scattering phenomena leading to point-scattering and the implications of the point scatterer assumption for SAR and ISAR autofocus.

[20]Equivalent in their overall autofocus technique if not in implementation.

for short range SAS imaging, unmeasured platform motion must be less than $\lambda/8$ (for low-order sway motions).

Correlation of complex signals has been explored and expressions for the expected accuracy of time-delay estimates have been provided. Additional expressions for the accuracy of phase-only and amplitude only correlation are derived. The link between phase-only correlation and the ML phase estimation kernel was investigated and the techniques were found to be equivalent (and have the same accuracy). The benefit of phase-only correlation is simple weighting of different sections of the time series.

Phase estimation techniques from a number of fields were explored in a common notation and their respective accuracies summarised. In addition, the relationships between, and the operation of, the techniques was discussed. The phase estimation kernels used in Knox-Thompson and bispectrum-based astronomical imaging, as well as those used in shear average autofocus and receiver self-calibration, are equivalent to eigenvector based phase estimation. Weighted phase estimation kernels provide large benefits for autofocus use by rejecting contributions from extended targets. For spotlight autofocus the WPGA phase estimation kernel was found to be best although requires modification for stripmap operation. Further research should enable its use. Stripmap phase estimation should currently use eigenvector phase estimation with M=2–6 (M=2 is the commonly used ML estimator).

The minor differences between micronavigation and autofocus have been summarised. In this thesis the term micronavigation is used to describe algorithms primarily relying on echo-redundancy whilst the term autofocus for those relying on image redundancy. The distinctions between spotlight and stripmap autofocus have been analysed with reference to the difference in blurring models. Explanations for the stripmap autofocus strategies employed to date are given based on this analysis. An overview of the major autofocus groups in common use in the SAS community has been presented summarising the algorithms that embody those groups.

Spotlight autofocus

The majority of stripmap autofocus algorithms are based on spotlight SAR algorithms. For this reason it is worthwhile summarising some of the most popular spotlight algorithms. The wavenumber transform for spotlight imaging is derived and close similarities between phase difference autofocus (PDA) and the phase gradient autofocus (PGA) family of algorithms are demonstrated.

8.1 Spotlight autofocus blurring model

Blurring models are essential for understanding the similarities and differences between spotlight and stripmap autofocus. Path-errors have different effects on the reconstructed synthetic aperture images for the two imaging modes. Blurring models help to describe those differences and provide insight into the autofocus problem.

The distorted pulse compressed image, $\widetilde{s}(t, u)$, is modelled by applying unknown path perturbations $X(u)$ to the ideal collected data $s(t, u)$. Thus

$$\widetilde{s}(t, u) = s(t, u) \exp\left(j2k_0 X(u)\right), \tag{8.1}$$

where the path errors are modelled as phase errors (a narrow-band approximation—see below). The co-ordinate mapping[1] that relates along-track spatial frequency k_y to aperture position u (the wavenumber transform, see Appendix D),

$$u = y - \frac{k_y(x - x_0)}{k_x}, \tag{8.2}$$

is used where the mapping is valid for small path perturbations for the image point x, y. It is assumed the spotlight imaging system is a narrow-band, narrow-beam system and the position errors are modelled by phase errors in the data collection process [Hawkins 1996; Jakowatz et al. 1996]. This assumes that the envelope shift in the recorded data is negligible—a valid assumption in narrow-band SAR. Employing tomographic system approximations—treating the system as narrow-beam, ($k_y \ll k_x$, $k_x \approx 2k$), narrow-

---

[1]To obtain the notation used elsewhere in this thesis x_n should be substituted for $(x - x_0)$.

band, $(k \approx k_0)$, and having a small swath-extent compared to the standoff range, $(y,\ x \ll x_0)$—the co-ordinate mapping between k_y and u (8.2) becomes[2]

$$u \approx \frac{k_y x_0}{2k_0}. \tag{8.3}$$

The spotlight wavenumber transform (8.3) shows that u is related to k_y via a fixed scaling which is space-invariant, i.e., does not change with scene position. One consequence of the space-invariant coordinate transformation is that the range-Doppler domain of the final image is closely related to the pulse-compressed echo data [Hawkins 1996, page 159]. Thus for a spotlight synthetic aperture system

$$s(t,u)\big|_{u=\frac{k_y x_0}{2k_0}} \equiv f(x, k_y), \tag{8.4}$$

where $x \approx ct/2$. This implies a tomographic SAR system can be considered as directly collecting Fourier domain data. This also means that aperture position is equivalent to spatial frequency. Often this duality of collection causes confusion in the discussions of autofocus algorithms.

Using the equivalence of domains described by (8.4), the blurring caused by a path error in $s(t,u)$ can be modelled as

$$\widetilde{s}(t,u) = s(t,u) \exp\left(j2k_0 X(u) \right),$$
$$\widetilde{f}(x, k_y) = f(x, k_y) \exp\left(j2k_0 X\left(\frac{k_y x_0}{2k_0} \right) \right), \tag{8.5}$$

where $\widetilde{s}(t,u)$ is the distorted pulse compressed signal and $X(u)$ and $X(k_y x_0/(2k_0))$ are equivalent via the spotlight wavenumber transform (8.3). If the phase errors are rapidly varying, there are large position errors, or the tomographic approximations fail, the two image domains are not related and the image blurring becomes space variant.

Using the spotlight SAR mapping (8.3), the image blurring is seen to be the convolution of the original image with a fixed blurring function. The distorted image is given by

$$\widetilde{f}(x,y) = f(x,y) \odot_y h(y), \tag{8.6}$$

where $\odot_y$ is a convolution in the y dimension only. The blurring function for tomographic systems, $h(y)$, is directly related to $X(u)$ and can be written

$$h(y) = \mathcal{F}^{-1}_{k_y \to y}\left\{ \exp\left(j2k_0 X\left(\frac{k_y x_0}{2k_0} \right) \right) \right\}\Bigg|_{u \equiv \frac{k_y x_0}{2k_0}}, \tag{8.7}$$

---

[2]Even though the mapping itself is only valid for the image point x, y, the small size of spotlight patches compared with the offset range allows the mapping approximation shown here to extend over the entire scene.

which is space-invariant. The convolution (8.6) is the result of point-spread-*invariant* blurring—all points in the image are corrupted by exactly the same function. The assumption of point-spread-invariant blurring is valid only for the plane-wave approximation of spotlight SAR.

Given that the above approximations are usual for spotlight mode systems and the convolution (8.6) holds, we wish to estimate $f(x, y)$ and perhaps $X(u)$ from the distorted raw data $\widetilde{f}(x, y)$ using autofocus techniques.

8.2   Shear average for spotlight systems

The shear average method was first published in the open literature by Fienup in his 1989 paper [Fienup 1989] as an adaptation of the shearing interferometer used in optics. A similar method was shown by Wahl et al. [1991] in a comparison of correlation and autofocus techniques. The method uses the along-track Fourier transform of the spotlight image,

$$f(x, k_y) = \mathcal{F}_y\{f(x, y)\}. \tag{8.8}$$

In plane-wave spotlight mode, this is equivalent to the measured pulse compressed data. Shear average estimates the phase differences across the aperture via

$$\widehat{\Delta\phi}(k_y) = \mathbf{Arg}\left\{\int f(x, k_y) f^*(x, k_y + \Delta k_y)\, \mathrm{d}x\right\}, \tag{8.9}$$

where the corrupting phase is calculated by summation of the differences via

$$\widehat{\varphi}(k_y) = \widehat{\varphi}(k_y - \Delta k_y) + \widehat{\Delta\phi}(k_y), \tag{8.10}$$

where Δk_y is usually set to use adjacent along-track frequency bins.

A number of similar algorithms exist where alternate phase estimators are used in the same fashion, replacing the ML estimator of shear average. Examples of this type of algorithm are: HOSPA [Nikias et al. 2000], ROPE Snarski [1996], bispectral estimation, and cumulant methods [Berizzi et al. 1996]. The phase estimators for these algorithms are summarised in Section 7.6.2. Often these algorithms provide an increased accuracy at the expense of increased computation.

Shear average for spotlight systems is equivalent to PGA with the windowing and circular shifting steps neglected. Leaving out those steps causes a decrease in the signal-to-clutter ratio, lowering algorithm performance [Wahl et al. 1991]. The benefit is efficient operation; by missing the circular shifting and windowing steps, image formation is not required and a number of Fourier transforms may be avoided. The algorithm alone is not likely to provide diffraction-limited imagery in typical systems [Fienup 2001] but provides a useful first-cut autofocus solution.

8.3   Map-drift autofocus (MD)

Map-drift (MD) is a parametric autofocus technique for estimating quadratic phase errors. The algorithm operates by correlating sub-aperture images—the intensity (or amplitude) image from the first half of the spatial Doppler spectrum is correlated with the equivalent image from the second half[3]. When a quadratic phase-error exists, the images will be shifted relative to each other. By measuring the correlation peak-shift the amount of quadratic phase error can be determined.

The distorted range-Doppler image is given by

$$\widetilde{f}(x, k_y) = f(x, k_y) \exp\left(j\alpha k_y^2\right), \tag{8.11}$$

where $f(x, k_y)$ is the undistorted range-Doppler domain image. Note that the collection-domain representation

$$\widetilde{s}(t, u) = s(t, u) \exp\left(4j\alpha u^2 k_0^2 / x_0^2\right), \tag{8.12}$$

is equivalent to (8.11) in a spotlight system (see Section 8.1). The blurred sub-aperture images are given by

$$\widetilde{f}_0(x, y) = \mathcal{F}_{k_y}^{-1}\left\{f(x, k_y - k_{y_{\max}}/2) \exp\left(j\alpha k_y^2 + \alpha k_{y_{\max}}/4\right)\right\}, \tag{8.13}$$

$$\widetilde{f}_1(x, y) = \mathcal{F}_{k_y}^{-1}\left\{f(x, k_y + k_{y_{\max}}/2) \exp\left(j\alpha k_y^2 + \alpha k_{y_{\max}}/4\right)\right\}. \tag{8.14}$$

The quadratic phase-error α may be estimated by measuring the relative shift in the images. This is performed by searching for the peak in the cross-correlation of the images

$$f_{01}(x, y) = \left|\widetilde{f}_0(x, y)\right| \star_y \left|\widetilde{f}_1(x, y)\right|, \tag{8.15}$$

$$= \mathcal{F}_{k_y}^{-1}\left\{\mathcal{F}_y\left\{\left|\widetilde{f}_0(x, y)\right|\right\}\mathbf{Conj}\left\{\mathcal{F}_y\left\{\left|\widetilde{f}_1(x, y)\right|\right\}\right\}\right\}, \tag{8.16}$$

where $\mathbf{Conj}\left\{\right\}$ refers to the conjugation operation.

Once α is estimated, the quadratic error is removed and the algorithm iterated. For spotlight imagery, the correlation must be of the intensity or magnitude images—a full complex correlation is not possible because each sub-aperture image occupies a different part of the Doppler spectrum (and is thus uncorrelated).

The need for iteration in the MD algorithm is because the sub-aperture images, (f_0, f_1), are blurred as well as shifted. The accuracy of the shift estimation is directly related to scene contrast—the higher the scene contrast the better the estimation.

---

[3]MD is thus very similar to the image correlation stripmap autofocus methods. In a stripmap setting the sub-aperture images have different spatial extent but as long as some overlap is present MD can operate.

Where there is sub-aperture image blurring, the initial phase-error estimates are poor because the image contrast is lower. With improving image accuracy, the shift estimation is enhanced. The MD algorithm typically converges in 2–3 iterations using 5% of the range bins containing the most scene energy [Carrera et al. 1995].

8.4   Multi-aperture MD

Multi-aperture MD (MAMD) is an extension of MD allowing estimation of third and higher orders of phase-errors. Multi-aperture MD adds the ability to estimate higher-order phase errors at the expense of increased computation and loss of estimation accuracy.

The technique begins, as MD, by subdividing the aperture; instead of dividing into two sub-apertures, multi-aperture MD divides the aperture according to the order of phase error to be estimated. Intensity images from each of the sub-apertures are formed and cross-correlated. The peak shifts in the correlation images are then used to estimate the corrupting phase error. Carrera et al. [1995] outlines a method for improving phase-estimation performance by also calculating correlations of non-adjacent sub-apertures.

Interestingly, estimation of phase from image shifts is the same problem as estimating the phase from phase derivatives. The peak shift in the correlation image is proportional to the overall phase gradient (linear phase shift). Thus multi-aperture MD estimates the average phase gradient over each sub-aperture[4]. Calloway et al. [1991] provides a more detailed description of the similarities between MD and phase-gradient based algorithms and concludes that PGA is better than multi-aperture MD in the presence of high-order phase errors.

Whilst multi-aperture MD can improve algorithm performance by allowing the estimation of higher order phase errors, the technique is still parametric. At the outset of the autofocus, the maximum order of phase error need to be estimated and fixes the number of sub-apertures used. Moreover, the robustness of the phase estimates is inversely proportional to the number of sub-apertures. As the number of sub-apertures increases, the signal-to-noise decreases in each image lowering the correlation peak and affecting the accuracy of the peak determination. Carrera et al. [1995, page 254] and Jakowatz et al. [1996, page 250] recommend a maximum of approximately 5 sub-apertures.

8.5   Phase difference autofocus (PDA)

The phase difference autofocus (PDA) algorithm is a parametric algorithm in the same vein as MD [Carrera et al. 1995]. The algorithm can be extended to higher order phase-

---

[4]The information that MAMD uses is the same as that phase-gradient based algorithms use.

error estimation in the same way as multi-aperture MD[5]. The major benefit of PDA is that it provides similar performance to MD without requiring iteration.

PDA finds the peak in a modified cross-correlation image, $f_{01}(x, y)$, which is calculated as (including the quadratic phase-error (8.11))

$$f_{01}(x, y) = \mathcal{F}_{k_y}^{-1}\left\{ \widetilde{f}_1(x, k_y) \widetilde{f}_0^*(x, k_y) \right\}, \tag{8.17}$$

$$= \mathcal{F}_{k_y}^{-1}\{ f(x, k_y + k_{y_{\max}}/2) f^*(x, k_y - k_{y_{\max}}/2) \exp(j2\alpha k_{y_{\max}} k_y) \}, \tag{8.18}$$

where $\widetilde{f}_0(x, k_y)$ and $\widetilde{f}_1(x, k_y)$, the blurred sub-aperture images, are given by (8.13) and (8.14) respectively[6]. When a quadratic phase-error exists, the linear phase shift in the Doppler domain results in an along-track shift in the peak of $f_{01}(x, y)$. The peak-shift is related to the quadratic error coefficient α by

$$\Delta y = 2\alpha k_{y_{\max}}. \tag{8.19}$$

PDA exploits *complex* correlation of *blur-free* images. The conjugation of $f_0(x, k_y)$ acts to cancel out any phase errors common to both $\widetilde{f}_0(x, k_y)$ and $\widetilde{f}_1(x, k_y)$. negating the need to iterate the algorithm. PDA's complex correlation of *demodulated* images can also be used with stripmap image correlation techniques and should improve performance over intensity image correlation.

PDA provides robust autofocus in situations where polynomial, space-invariant phase errors exist. The PDA algorithm is able to replace MD in all situations and provides similar levels of performance without the need for iteration.

8.6   Phase gradient autofocus (PGA)

Phase gradient autofocus (PGA) is a technique that first appeared in the open literature in 1989 with the paper by Eichel et al. [1989]. It is a technique for autofocusing spotlight SAR imagery that may be described using the plane wave formulation[7] (see Section 8.1). The phase gradient algorithm (PGA) is a particularly successful method for obtaining a blur-free image estimate from corrupted raw data—it has become the standard for spotlight mode autofocus. The algorithm is summarised as follows:

---

[5]Interestingly, the multi-aperture PDA phase kernel is identical to the ML phase estimator (see Section 7.6.2) when the number of sub-apertures is equal to the number of along-track spatial-frequency bins. (In that situation, the correlation of PDA is unnecessary; the phase of the (1 pixel) result gives the correlation peak shift in the same way as narrow-band correlation estimates time-delay.)

[6]It is important to note that the sub-aperture range-Doppler domain images $\widetilde{f}_1(x, k_y)$ and $\widetilde{f}_2(x, k_y)$ are multiplied directly. These images, by definition, are centred on zero Doppler. This ensures the complex correlation of sub-band images used in PDA is not of sub-aperture images but of *demodulated* sub-aperture images. Direct along-track correlation of sub-aperture images fails unless the sub-aperture regions overlap.

[7]In this mode the SAR autofocus problem is very similar to the problems encountered in speckle astronomy [Eichel and Jakowatz 1989; Fienup 1989].

1. In imitation of the method of shift and add [Bates and Cady 1980], the brightest
   points of the image estimate are circularly shifted to the centre of the image.
   This removes any unknown linear shifts that may otherwise affect the phase
   estimation[8]. For every value of cross-track x, the most prominent point for all
   y in $\widetilde{f}(x,y)$ is selected (the m^{th} target) and the image shifted in y to the centre
   row of a new image,

   $$g_m(x,y) = \widetilde{f}(x, y - y_m(x)). \tag{8.20}$$

2. Along-track windowing is performed to improve the phase-estimation later in the
   algorithm. This removes clutter and weak targets which acts as noise in the phase
   estimation. The windowed image is given by

   $$g(x,y) = w(y)\ g_m(x,y). \tag{8.21}$$

   where the window function $w(y)$ limits the extent of each prominent point to
   a region only a few pixels wide. The windowing step is important in overall
   operation of PGA and is discussed in more detail in Section 8.6.2. At this stage,
   the shift and add method estimates the point-spread function by summing the
   data in the cross-track direction.

3. The along-track Fourier Transform of the windowed image is taken. The range-
   Doppler domain of this shifted, windowed image is described by

   $$g(x, k_y) = \mathcal{F}_y\{g(x,y)\}. \tag{8.22}$$

   Assuming the cross-track range bins are statistically independent[9] and the path-
   deviation can be modelled as a constant phase shift for each range bin, the system
   may be modelled as [Jakowatz and Wahl 1993],

   $$\begin{aligned}
   g(x, k_y) &= a(x, k_y) + \mathcal{N}(x, k_y), \\
   g(x, k_y + \Delta k_y) &= a(x, k_y + \Delta k_y)\exp\left(j\Delta\phi_1(k_y)\right) + \mathcal{N}(x, k_y + \Delta k_y), \\
   &\vdots \\
   g(x, k_y + \Delta_L k_y) &= a(x, k_y + \Delta_L k_y)\exp\left(j\Delta\phi_L(k_y)\right) + \mathcal{N}(x, k_y + \Delta_L k_y),
   \end{aligned} \tag{8.23}$$

   for a given along-track spatial frequency k_y, where $a(x, k_y)$ are amplitudes and
   phases representing the targets, $\Delta\phi_L(k_y)$ is the phase difference between $g(x, k_y)$
   and $g(x, k_y + \Delta_L k_y)$, and $\mathcal{N}(x, k_y)$ are the unknown noise signals. The phases

---

[8]Random linear phases (removed by centre shifting) would cause the averaging step of PGA to fail.
These multiplicative phase errors prevent the coherent addition of phase gradients. The problem is
explained using the same argument as for why phases do not coherently average when unknown phase
offsets exist.

[9]This assumption is important and is borrowed from the speckle imaging assumption that individual
speckles are independent.

contributed by the targets (the phases of $a(x, k_y)$) are assumed constant with along-track spatial frequency, k_y. Any phase differences are caused by the (phase) signature target in question (useful for target classification purposes [Jakowatz et al. 1996]) and additive noise. For PGA to operate, this changing phase information must be zero mean—any bias or common phase information will be interpreted as a phase error requiring removal. Note the constant phase offset contained in the $a(x, k_y)$ terms prevent simple averaging of the aperture phases [Chan and Yeo 1998]. This multiplicative phase constant varies randomly with range and prevents coherent phase averaging.

4. Phase gradients with respect to k_y are calculated and averaged to avoid the problem noted above. The phase-gradient is calculated via

$$\widehat{\Delta\phi}(x, k_y) = \mathbf{Arg}\left\{g(x, k_y)g^*(x, k_y + \Delta k_y)\right\}, \tag{8.24}$$

and the Maximum Likelihood (ML) phase gradient is calculated using[10]

$$\widehat{\Delta\phi}(k_y) = \mathbf{Arg}\left\{\int_x g(x, k_y)g^*(x, k_y + \Delta k_y)\,\mathrm{d}x\right\}. \tag{8.25}$$

5. The phase gradient estimate is integrated to calculate the distorting phase error

$$\widehat{\phi}(k_y) = \int_{k_y'} \widehat{\Delta\phi}(k_y')\,\mathrm{d}k_y'; \qquad \widehat{\phi}(0) \equiv 0. \tag{8.26}$$

This step is usually calculated as a discrete summation

$$\widehat{\phi}[Q] = \sum_{q'=1}^{Q-1} \widehat{\Delta\phi}[q']; \qquad \widehat{\phi}[0] \equiv 0, \tag{8.27}$$

where q is the discrete equivalent of k_y.

6. To prevent image shifting and skewing, the linear phase and constant offset phase (that PGA is unable to estimate) are removed. The phase estimate is also interpolated (upsampled) to provide the same number samples as the image. (The original has fewer samples due to the initial Fourier transform often operating on a truncated dataset).

7. The equivalent path error is calculated from the phase error using

$$\widehat{X}(k_y) = \frac{\widehat{\phi}(k_y)}{2k_0}, \tag{8.28}$$

---

[10]A weighted least squares average is possible and is discussed in Section 8.6.3.

which is then removed from the distorted range-Doppler image via

$$\widetilde{f}_{\text{new}}(x, k_y) = \widetilde{f}(x, k_y) \exp\left(-j2k_0 \widehat{X}(k_y)\right). \tag{8.29}$$

PGA then takes the new image $\widetilde{f}_{\text{new}}(x, y)$ and iterates starting back at the step 1. The window size is reduced with subsequent iterations depending on the type of scene used. A summary of window selection strategies is discussed by Wahl et al. [1994a]. Typically PGA converges within 4–5 iterations and provides estimation of high-order phase errors unable to be estimated using the parametric MD and PDA algorithms.

8.6.1 Centre shifting

PGA's centre shifting step is an important part of the overall algorithm and performance decreases if it is neglected [Wahl et al. 1994a]. Centre shifting consists of taking the strongest scatterer in a given range-bin and barrel-shifting[11] it (in along-track) to the centre of the image. The resulting image then has all of the strong targets aligned at $y = 0$.

Centre shifting in PGA is an imitation of the shift-and-add algorithm [Bates and Cady 1980] and if considered naively, is performed for the same purpose (that of directly estimating the point-spread-function via an along-track summation). Instead, the most important consequences of the centre shifting are in the estimation of the window width and the removal of random linear phase trends. Selection of the window width using the shifted image is a straightforward non-coherent summation along the range-bins and picking the -10 dB point (discussed in the following section). QPGA [Chan and Yeo 1998] shows modifications to remove centre shifting from the window width estimation of PGA—although it does employ centre shifting later in the algorithm for computational savings. Removal of the random linear phase trend (constant phase gradient term) from each target signature improves phase estimation. Averaging of the (random) constant phase gradients decreases performance (for the same reason that phase averaging fails when constant random phases are averaged). Jakowatz et al. [1996] suggests that this linear phase removal is desirable. Better phase estimation is possible if the constant phase gradient error is estimated and removed (as in done in PMA). Phase estimation using the phase-only eigenvector kernel also benefits from removal of linear phase trends as this prevents unwrapping difficulties later in the algorithm. If neither the centre shifting or windowing operations are performed, PGA is identical to the spotlight shear average algorithm.

To summarise, the centre shifting operation of PGA provides a way of simplifying the window width selection and also giving a convenient visual indication of autofocus

---

[11]Performed for computational efficiency and to help alleviate circular convolution effects.

progress. Circular shifting aids accuracy by improving the averaging of phase gradients.

8.6.2 Windowing

Scene clutter acts as noise in PGA and the accuracy of the PGA algorithm is dependent on the signal-to-clutter ratio of the image scene [Jakowatz and Wahl 1993]

$$\beta = \frac{\sigma_a^2}{\sigma_n^2}, \tag{8.30}$$

where σ_a^2 and σ_n^2 are the signal and clutter variances respectively. The windowing operation of PGA attempts to improve accuracy by maximising β for each prominent scatterer. This is performed by placing a window around the blurred prominent scatterer such that clutter in the same range bin is rejected. Windowing masks the image with an amplitude function placed around the prominent scatterer. The window usually chosen for PGA uniformly weights the image inside the window and rejects the image outside—i.e., a rectangular window.

Windowing has many effects, the most important of which is limiting the order[12] of phase error estimated [Warner et al. 2000]. This makes the width of the window a critical parameter in PGA. In particular, if the width of the window does not fully encompass the blurring of the scene, PGA will not converge to the correct solution (see PGA focused images in [Barbarossa and Scaglione 1998]). Conversely, if the width of the window is too large, unnecessary noise is injected into the estimation also degrading performance [Jakowatz et al. 1996, page 254]. Warner et al. [2000] presents an in-depth discussion on the effects of window width on PGA performance. An interesting side-effect of stripmap autofocus (see Chapter 10) is that the sway-order implications of the window size are lifted. This however comes at the expense of requiring a complicated wavenumber transform.

A common method for selecting the window width is to use a non-coherent average in across-track of the centre shifted image. This average attempts to estimate the extent of the blurring from the prominent scatterers in the scene. The width of the non-coherent average is approximately the width of the point-spread-function. The window width is then calculated based on the -10 dB point of the average [Jakowatz et al. 1996]. Iteration reduces the amount of blurring, and the window width has a corresponding reduction. This metric for window width determination is not suitable for clutter-like images[13] (most SAS images) and instead the window width set at the outset and made progressively smaller (usually by a factor of between 0.5–0.9 each iteration).

---

[12]Note that this is actually caused by a limiting the along-track spatial frequencies of the phase error.

[13]In clutter images the blurring extent is not evident as the blurring affects only the contrast of the scene.

In SAS autofocus scenarios, window width should be chosen based on prior information about the amount of phase error present in the image. When PGA follows a micronavigation algorithm, information about the amount of residual phase distortion is available from the algorithm (straightforward extensions can retrieve this information) [Bellettini and Pinto 2002].

Window width selection is a critical PGA implementation issue that has a large effect on end performance. This performance impact is more prevalent when using less accurate phase estimation methods. For SAS usage, the initial window size should be selected based on prior information about the magnitude of phase errors and progressively reduced with successive iterations—a *progressive* window [Warner et al. 2000].

8.6.3   Phase estimation

Phase estimation is perhaps the most important part of the PGA algorithm. More effective phase estimators can ameliorate some of the effects of poor window selection. The importance of the phase estimation is probably underscored by the large number of PGA variations and similar autofocus techniques varying only the phase estimator. Due to its importance in many autofocus algorithms phase estimation is described separately in Section 7.6.2.

8.7   PGA variants

PGA has a number of variations used to improve accuracy in specific situations or offering efficiency improvements. A number of PGA's variations are detailed in this section.

8.7.1   2-D PGA

Warner et al. [2000] notes that there is nothing preventing the PGA procedure from being used in situations where non-separable 2-D phase errors occur. All of PGA's operations, except the integration of 2-D phase differences to get the 2-D phase error, are able to be applied. The difficulties in the operation of 2-D PGA are concentrated in a 2-D phase unwrapping required to replace PGA's integration. This problem is identical to 2-D phase unwrapping for interferometry where residues cause unwrapping difficulties.

8.7.2   FLOS-PGA

Fractional lower order statistics (FLOS)-PGA is a variation of PGA that substitutes a different phase estimation kernel (the FLOS kernel described in Section 7.6.2, [Tsakalides

and Nikias 2001]). Performance is improved in environments where the clutter process is non Gaussian for only a minor increase in computation.

8.7.3 Weighted PGA (WPGA)

WPGA is another variation of PGA that alters the phase estimation kernel [Ye et al. 1999]. In this case the kernel described in Section 62 is used. The phase kernel performs optimally for any clutter process by estimating the clutter-noise variance. Additionally, the phase kernel is able to lower the contribution from strong-scattering targets that would otherwise bias the autofocus result. This method represents the optimum PGA method for SAR applications.

WPGA relies on a number of spotlight SAR assumptions and the phase estimator is not currently usable for SAS autofocus. More research is required to extend WPGA to SAS.

8.7.4 Quality PGA (QPGA)

QPGA is an attempt to perform PGA without iteration [Chan and Yeo 1998]. To do so it selects a large number of *quality* targets rather than the strongest target per range-bin. The quality estimate is based on a technique similar to the noise-power estimation in WPGA. As in WPGA, the use of a quality of target metric can prevent biasing caused by strong-scattering extended targets.

The algorithm operates as a single iteration of PGA, except multiple targets per range-bin are selected, and circular shifting is avoided. The use of a higher-order eigenvector phase-estimator combined with using only targets that meet a certain quality threshold results in autofocus improvement. Both improvements, as in the similar WPGA, help autofocus performance when the signal-to-clutter ratio is low. Altered windowing strategies are employed because of the non-iterative nature of QPGA; usually the window width is chosen based on a prior estimate of the amount of blurring.

QPGA offers a PGA-based autofocus technique suitable for real-time implementation. Iterative QPGA offers benefits over traditional PGA that approach those of WPGA. QPGA relies on a number of spotlight SAR assumptions in the calculation of the target quality and requires modification for stripmap SAS use.

8.7.5 Mosaic PGA (mPGA)

Mosaic PGA (mPGA) [Bonifant 1999; Bonifant et al. 2000] is a PGA variation for stripmap autofocus. The algorithm takes stripmap imagery and slices it in along-track to get overlapping image sections that are autofocused individually. The individually autofocused sub-images are then combined into a fully autofocused image. mPGA thus avoids the problem caused by space-variance in stripmap imagery by operating

on space-invariant sub-swathes. Interestingly, mPGA is based on narrow-band and narrow-beam (traditional spotlight assumptions) but incorporates the along-track target position. Modifications are required to allow wide-beam, wide-band SAS operation.

8.8 Summary

Spotlight SAR autofocus algorithms often provide inspiration for similar stripmap algorithms and the relationship between the two modes has been explored. The sway induced blurring in spotlight systems was derived from the stripmap wavenumber transform and under common spotlight approximations shown to be space-invariant—i.e., all scatterers suffer the same distortion. This is identical to the commonly used spotlight autofocus blurring model.

MD, MAMD, and PDA, have been described in a common notation and the similarity between the sub-aperture correlation based PDA and the phase gradient based PGA explored. PDA operates by estimating the phase gradient over an entire sub-aperture (using correlation) whereas PGA estimates the phase gradient for each along-track spatial frequency. PDA should always be used in preference to MD as it offers similar performance with less computation.

A detailed summary of the PGA algorithm has been presented with a discussion of each of the major components of the algorithm. PGA variants have also been summarised and their differences highlighted. WPGA should be used in preference to the other PGA variant algorithms.

Micronavigation using reverberation based autofocus

The algorithms described in this chapter estimate the platform path using the recorded echo data (usually without image reconstruction taking place). The algorithms discussed require some data redundancy in azimuth, i.e., the SAS must travel somewhat slower than sampling considerations alone would allow[1]. With sufficient data redundancy (and combined with a sonar's on-board inertial navigation system) the navigation accuracies on a pulse-to-pulse basis are in the sub-millimetre order [Bellettini and Pinto 2002].

This chapter primarily covers three bulk-motion estimation algorithms: the redundant phase centre (RPC) algorithm (the mainstay of multiple-receiver bulk autofocus), image correlation (an extension of RPC), and shear average (equivalent to RPC for monostatic SAS systems). A discussion on the relationships between the algorithms and extensions to the shear average algorithm are presented. Results of the shear average extensions operating on both simulated and field-collected data are also presented. The algorithms described in this chapter are best applied to SAS data before other autofocus techniques as a way of quickly removing the bulk of the image distortion.

9.1   Redundant phase centre (RPC) algorithm (DPC, DPCA)

The redundant phase centre (RPC) algorithm [Sheriff 1992] is used for estimating the across track sonar sway from the recorded echo data. It is also known as the displaced phase centre (DPC) algorithm [Groen and Sabel 2002; Raven 1981] and displaced phase centre antenna (DPCA) [Bellettini and Pinto 2002].

RPC operates by exploiting redundancy in the data obtained by ensuring vernier-array SAS data is oversampled in azimuth. With the accuracy of RPC being directly related to the amount of redundancy in the data (oversampling factor), the efficacy of the algorithm is dependent on the platform speed. RPC provides the mainstay of vernier-array autofocus for the SAS community with almost all multiple-receiver SAS

---

[1]This is the case when considering $D/2$ sampling. A SAS operating using the $D/4$ sampling constraint discussed in Section 2.6.5 automatically provides enough redundancy for autofocus purposes.

implementations using an implementation of the algorithm in some form. An in-depth analysis of RPC and a comparison of the theoretical accuracy versus that required for SAS imagery is presented by Bellettini and Pinto [2002].

9.1.1   RPC operation

The operation of RPC is as follows:

1. Using the phase centre approximation (see Section 3.2.2) the positions of the receivers that collect redundant data between pings are calculated. This step depends on the along-track velocity of the sonar. If suitably aligned hydrophones do not exist, the data is interpolated in along track to provide aligned pseudo-hydrophones. Typical hydrophone alignment and phase centre redundancy can be seen in Figure 9.1.

2. The time-delay differences between the redundant hydrophones (these hydrophones are sometimes called the redundant sub-arrays) are calculated by cross-correlating the hydrophone signals. Unambiguous time-delay estimates are derived using the methods shown in Section 7.6.1.

3. The sonar's sway and yaw are determined from the time-delay estimates: sway is estimated using the pulse-to-pulse time-delay common to all redundant pairs, and yaw is estimated using the gradient of the time-delays.

For typical SAS systems, the yaw estimation accuracy is poor leading to degradation in the accuracy of sway estimation [Bellettini and Pinto 2002]. This lack of yaw estimation accuracy is due to the extremely small time-delays yaw causes. To make sway estimation as accurate as possible, yaw is often better estimated using an on-board INS where possible [Wang et al. 2001]. Platform surge inaccuracies also cause sway and yaw estimation difficulties. However, these effects can be avoided by estimating the surge using a similar cross-correlation technique in the along-track direction [Shippey et al. 2001]. Simple interpolation schemes are then used to allow RPC by forming pseudo-redundant-pairs. A typical multiple-receiver system estimates surge and sway using RPC and uses an INS for estimating yaw.

Array calibration is extremely important when accurate RPC is desired. Any distortions in the receiver array cause significant estimation error in the algorithm. Array distortion is straightforward to estimate. Many calibration algorithms exist with the most promising appearing similar to spotlight SAR autofocus algorithms. For example, the algorithm presented in [Douglas et al. 1992] has many similarities to shear average—it operates by cross-correlating the individual element signals with a reference element, instead of cross-correlating adjacent elements[2]. Array calibration is the

---

[2]Banks and Griffiths [2002] presents a variation of this algorithm.
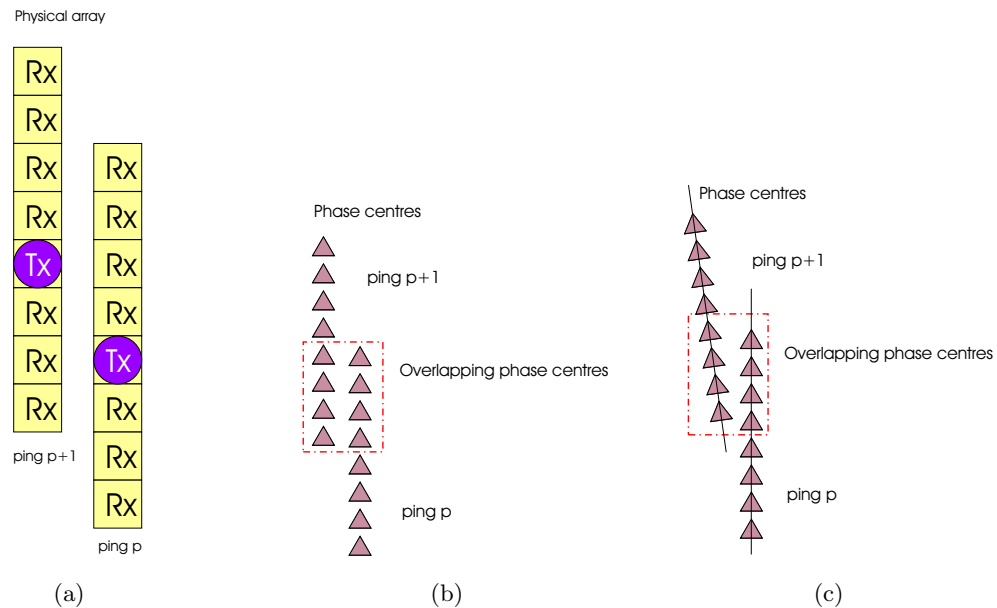
Figure 9.1 Redundant phase centres algorithm geometry. (a) The physical arrangement of an over-sampled multiple-receiver SAS. (b) The inter-pulse redundant phase centres of a multiple-receiver SAS. (c) The redundant phase centres allow platform surge, sway, and yaw to be estimated [Bellettini and Pinto 2002; Shippey et al. 2001].

same problem as single receiver autofocus. With calibration errors, all targets suffer the same blurring, i.e., the blurring is space invariant. When the blurring is space invariant, spotlight SAR algorithms, such as PGA and spotlight shear average (Chapter 8), provide accurate sway estimation—the same information required for array distortion calibration. By using straightforward autofocus algorithms for calibration, slow time variation in the array characteristics and receiver locations can be estimated. This is particularly important for arrays subjected to large physical stresses as those caused by pressure or wave-motion which can alter the shape of the array. Once the receiver array is calibrated and the data compensated, image blurring is reduced and RPC operates close to predicted accuracy [Banks and Griffiths 2002].

9.1.2 Directivity

RPC determines sway and yaw estimates based on the assumption of broadside scattering since strongly scattering targets at angles other than broadside cause estimation bias. The bias is due to a position shift of $X(u)$ appearing as a timing shift in the echo data of $X(u)\cos\theta$, where θ is the angle to target (see Chapter 6). A related problem occurs even in uniform clutter—because energy comes from all angles not just broadside. The sway underestimation error ϵ caused by assuming a narrow-beam geometry

Figure 9.2   Beam-width problem for RPC. Energy comes from all angles in the beam resulting in a $\int_{-\theta_{3\text{dB}}/2}^{\theta_{3\text{dB}}/2} a(\theta) \cos \theta \, d\theta - X(u)$ underestimation of sway. The bias is able to be removed if the beam-width is estimated. In general, the integration is further biased by strong scattering targets—beam-to-beam correlation is required to remove this type of bias [Tonard and Brussieux 1997].

is given by

$$\epsilon = \int_{-\theta_{3\text{dB}}/2}^{\theta_{3\text{dB}}/2} a(\theta) X(u) \cos \theta \, d\theta - X(u), \qquad (9.1)$$

where $a(\theta)$ is the energy received at a particular angle. Figure 9.2 demonstrates one of the bias-causing situations. Ideally, RPC would be able to estimate the angle to each scatterer independently and be able to remove the estimation bias.

When several redundant phase centres exist, the angle to each part of the scene can be independently determined. The method of correlating beams described by Tonard and Brussieux [1997] implements a technique based on this principle. The beam correlation method forms a set of beams from the pulse-to-pulse redundant sub-arrays and correlates these beams instead of the raw echo signals. Estimation accuracy improves with increasing number of elements in the sub-arrays (a consequence of reducing error in angle estimation). Beam correlation RPC has similarities[3] to the image correlation techniques discussed in the next section but still requires echo-correlation. Range correlations of the beams provide the sway estimation and beam-to-beam correlations provide yaw information.

9.2   Image correlation

Other extensions of the RPC algorithm are embodied in image correlation techniques [Banks and Griffiths 2002; Gough and Miller 2002; Groen and Sabel 2002; Shippey

---

[3]This is because a pulse-to-pulse polar image is reconstructed from the redundant section of the array.

et al. 2001]. These methods make low resolution images on a pulse-by-pulse basis (i.e., a multiple-receiver system is still required[4]) and correlate those images instead of the raw echo data. Importantly, phase information is usually discarded in these correlations for the same reasons as in map-drift (see Section 8.3). Similar correlation methods are used for aberration correction in the ultrasonic imaging field [Karaman et al. 1995; Trahey and Nock 1992]. The additional benefit of using *image* correlation is that no redundancy in data collection is required, i.e., no overlapping phase centres are required—although if overlapping phase centres exist the algorithm performance improves.

Redundancy in data collection aids the operation of image correlation algorithms and avoids some of the problems discussed later in the section. Redundancy *must* exist in the image itself if there is no echo redundancy. This image redundancy is evident as image structure; if the image consists solely of sea-floor clutter, image correlation algorithms fail [Billon and Pinto 1995]. Failure on clutter-only images is also evident in the sub-aperture correlation autofocus techniques such as map-drift and phase difference autofocus (see Section 8.3, [Huxtable and Geyer 1993]). Redundancy is *essential* if autofocus is to be used, whether in the echo data itself (a reliable source of redundancy if the receiver arrays overlap), or in the statistical nature of the scene.

A reliance on image redundancy gives the conceptually straightforward image correlation and sub-aperture correlation techniques significant drawbacks. In particular, when there is no array overlap, image correlation techniques suffer similar drawbacks to other autofocus techniques. The main drawbacks are[5]:

Speckle-noise — Speckle-noise is independent in each image so the speckle patterns do not correlate, raising the correlation noise level. The lack of speckle correlation makes motion parameter determination more difficult and lowers overall accuracy. The effects can be reduced by using multi-look images.

Scene contrast — Image correlation techniques rely on intensity variation in the scene (other than image variation caused by speckle noise). A lack of variation (low scene contrast) lowers the correlation energy and increases the influence of noise on parameter estimation. This was also noted by Billon and Pinto [1995] who suggest an INS may always be needed if navigation is to operate in bland scenes.

Non-Lambertian scattering — Image correlation requires that the underlying scene looks the same from different imaging angles. Complicated scatterers can have wildly varying target strength with varying angle (non-Lambertian scattering).

---

[4]The image correlation techniques may be generalised for use with single receiver SAS systems in a straightforward fashion. This is performed in a similar manner to the operation of sub-aperture processing in SAR autofocus—images from a series of adjacent pulses are formed and correlated.

[5]See [Moreira et al. 1997] for a discussion of general sub-aperture correlation autofocus drawbacks.

In addition, the target peak may shift with angle. Either of these effects cause false motion estimates.

Shadowing — A strong scattering target partially in shadow of another target appears for some angles and not others. Image shadowing thus causes difficulties for the same reasons as non-Lambertian scattering.

Image correlation techniques share other similarities with sub-aperture correlation/aperture division autofocus techniques. Like sub-aperture correlation, image correlation requires magnitude or intensity images to be correlated to ensure high correlation coefficients. Direct correlation of the complex images is prevented by an along-track wavenumber shift (similar to that seen in SAR interferometry [Gatelli et al. 1994]). However, the faster phase-based spotlight sub-aperture correlation algorithms (such as phase difference autofocus (PDA) [Curlander and McDonough 1996]) do allow complex correlation. Thus it is likely that image correlation could employ similar techniques to enhance accuracy. Image correlation effectively provides a stripmap equivalent of spotlight sub-aperture correlation algorithms such as map-drift[6]. This equivalence allows image correlation algorithms to be compared with traditional spotlight SAR algorithms and useful parallels to be drawn.

9.2.1   Cascade algorithm

The cascade autofocus algorithm [Douglas 1993; Douglas and Lee 1993a; Silkaitis et al. 1994] is another image correlation/RPC based autofocus algorithm used for SAS imaging. Originally, the algorithm only estimated sway [Douglas and Lee 1993a], whilst Silkaitis et al. [1994] demonstrates surge, sway and yaw estimation. The later paper also notes that large yaw errors can severely affect the rest of the algorithm (RPC also suffers this problem).

As published, the cascade algorithm is suited to narrow-band systems—the algorithm employs phase-only correlations for estimating sonar parameters. This restriction may be lifted when using the correlation techniques discussed in Section 7.6.1. The cascade algorithm is identical to the other image correlation algorithms when general correlation techniques are employed.

9.3   Shear average

Shear average appeared in the SAS literature in 1992 as a technique for estimating sway in monostatic SAS systems [Johnson 1992] and in a later paper in 1995 [Johnson et al.

---

[6]The techniques used in map-drift are not restricted to spotlight autofocus—the correlation can use the overlapping parts of the image in the same way as image correlation. Map-drift is different from image correlation only in the method used to map the image shifts to a distorting phase—it uses spotlight assumptions.

1995]. The technique was developed independently of the work of Fienup in the SAR autofocus field [Fienup 1989] but in essence is a stripmap version of the shear averaging algorithm that he presents. (In a spotlight system the algorithms are equivalent due to the duality of the image range-Doppler and data collection domains.) The algorithm is also the single-receiver equivalent of RPC.

Shear average generates path distortion estimates by correlating time-signals from adjacent pulses. Originally, the correlation used was a phase-only correlation (sufficient for most systems because the pulse-to-pulse sway differences are usually smaller than $\lambda/2$). Phase unwrapping (if the phase is not wildly varying) or the use of non-coherent shear average extends the method to cases when the distortion is severe (see Section 9.3.4, [Callow et al. 2001c]). If the scene is statistically homogeneous and the system is oversampled in along-track, sway motions can readily be estimated using the correlation estimates [Johnson et al. 1995].

Similar techniques have been used in ultrasonic imaging when estimating near-field distortions. Other techniques rely on correlating element signal to a reference element signal instead of correlating adjacent element signals [Douglas et al. 1992]. Reference-based correlation has been shown to reduce low-order estimation errors [Behar 2002; Krishnan et al. 1997] but has poor high-order estimation[7]. However, reference-correlation combined with element-to-element based estimates provides accuracy for both high-order and low-order information [Behar 2002]. This type of technique is equivalent to the phase-closure phase kernels discussed in Section 7.6.2. The techniques tried in the ultrasonics field require modification to operate with wide-beam SAS systems although do provide insight into alternate autofocus techniques.

Shear average assumes certain echo statistics to obtain autofocus measures. The assumption required is that the sea-floor is delta-correlated [Fienup 1989] (i.e., that the scene is statistically homogeneous) and that the scene contrast is low (so the echoes themselves are spatially delta-correlated). On clutter images this assumption appears valid and reasonable autofocus occurs both in simulation and in practice [Pihl et al. 2000; Shippey et al. 1998a]. In the presence of a strong-scattering target, the delta correlation assumption breaks down[8] and the autofocus quality degrades. This degradation is similar to a known problem in spotlight SAR autofocus with the quality of some parameter estimates decreasing with increasing image contrast [Berizzi et al. 1997; Madsen 1989]. One way of mitigating this problem is to alter the weighting of the phase estimates to reduce the contribution of the strong targets [Callow et al. 2001b, c]; this is discussed in further detail in Section 9.3.1.

Shear average is the single-receiver equivalent of RPC. If the echo data is sampled

---

[7]These techniques are only viable if there is sufficient cross-correlation energy. This is not the case in a $D/2$ sampled image of a statistically homogeneous scene.

[8]Even though the image containing strong-scattering targets is still delta-correlated, the pulse-compressed image may not be. Shear average requires that the *pulse-compressed* image show delta-correlation properties.

Figure 9.3 Subdividing a single receive element into portions. (a) The physical hardware of the single receiver SAS. (b) In this framework a single receiver can be shown to be the summation of individual elements in a multiple-receiver array. The phase centre idea can be used to show redundancy in the data collection (i.e., from the overlapping receiver section). Shear average correlates the response of the entire array—both the overlapping and non-overlapping sections—to estimate the across-track shift between pulses.

more often than $D/2$ then there is enough redundancy[9] for the algorithm to operate. This echo redundancy is a function of the amount of receiver overlap.

A multiple-receiver sonar with an omni-directional transmitter is equivalent to a monostatic sonar if all the receiver elements are summed[10]. A monostatic sonar may thus be treated as a multiple-receiver sonar that only has access to the sum of the element signals. This equivalence is illustrated in Figure 9.3. Using the same phase centres approach used in RPC (see Section 9.1), it is possible to show that when sampling more often than $D/2$ some of the equivalent phase centres overlap. In a single-receiver system the true echo signal is the sum of the equivalent overlapping and non-overlapping phase centres. The overlapping phase centres provide redundancy for the algorithm's operation while the non-overlapping phase centres raise the noise-floor. As for RPC, the redundant part of the signal allows for navigation. Shear average only has access to the entire array's echo signal and is not able to remove the contribution of the non-overlapping receivers. Shear average is thus unable to estimate sonar yaw, or use advanced extensions of RPC such as beam or image correlation that employ the signals from individual elements. This results in some major drawbacks:

Sway underestimation — shear average is a narrow-beam formulation, increasing

---

[9]Redundancy exists even in data normally considered under-sampled—however there is not much. This is solely due to the unusual sampling constraints on SAS systems (see Section 2.6.5).

[10]The effect of the transmit beam-pattern is neglected but as the transmit beam-pattern of a monostatic sonar is identical to the receive beam-pattern the effect is straightforward to model.

beam-widths cause sway to be underestimated. This problem is the same as the equivalent bias in RPC. In homogeneous clutter, the amount of underestimation is quantifiable and can be corrected.

Beam directivity — If the majority of the energy in the signal comes from an angle, θ, then the sway is underestimated (caused by the $\cos\theta$ factor biasing the sway estimate (cf. with Chapter 6 and Section 9.1.2)). A bias in the direction of energy also causes the sway-underestimation-correction described above to fail (the directivity bias cannot be detected by forming beams as is possible with RPC).

Non-overlapping collection/strong scatterer biasing — shear average is unable to determine which parts of the signal come from the overlapping receiver section and so must correlate the entire signal. In homogeneous clutter, the non-overlapping receiver section does not correlate and acts to increase the noise floor. When significant scene variation exists—such as caused by a strong scattering target—the non-overlapping receiver sections have *significant* correlation. The along-track displacement in the non-overlapping sections of the array combined with hyperbolic range-migration causes an echo time-shift from each array section. The effect appears as a sway motion. This time-shift is only weakly related to platform sway and is largely a function of along-track displacement between pulses. Motion estimates are biased by strong scattering targets and scene intensity variation. The biasing may be reduced by reducing the contribution strong scattering targets have on the result using weighted correlation (see Section 9.3.1).

System yaw — shear average cannot determine the difference between a system yaw and a sway. A constant yaw leads to the estimation of a linear sway. When the yaw varies in more complicated fashion, incorrect sways are estimated (the sway estimated is the integral of the yaw). This effect causes the major limitation in autofocusing field-collected data using shear average. Few systems collect data without any system yaw and even though it causes little image blurring; it prevents accurate sway estimation.

Aside from the drawbacks mentioned above, the shear average technique, like RPC, is able to focus both propagation errors and tow-fish position estimation based errors. The algorithm works efficiently (especially the phase only version) and is particularly suited for use as a bulk motion estimator of use before other, often processing-intensive, autofocus is employed [Fortune et al. 2001b]. Additionally, because only adjacent pulses are stored and correlated, the algorithm is suited to real-time operation.

9.3.1 Amplitude weighting

Shear average estimates are biased when the target scene contains strong-scattering reflectors. The cause of the biasing is two-fold: beam-directivity causes sway underestimation, and the non-overlapping receiver section causes quadratic bias. The first source of biasing results since shear average assumes the majority of the signal energy comes from broadside. A cosine θ relation between sway and signal timing shift causes bias when targets are not at broadside. The second bias is from non-overlapping receivers contributing to the estimate if strong-scattering targets cause scene correlation. Scene-correlation is mistakenly detected in shear average algorithms as a path-deviation; its effect is to straighten the range curvature of a strong-scattering target leading to hyperbolic biasing in the region of that target. The biasing has a detrimental effect on autofocus performance, causing residual blurring.

Choosing a region of sea-floor without targets of interest (no strong-scattering targets) is one method of reducing the bias problem in autofocus situations [Shippey et al. 1998a]. A more sophisticated method is to weight differing sections of sea-floor to reduce the contribution from strong scattering targets. One possibility is to ignore the strong-targets altogether; this is normally done using an additional weighting in the summation described by (7.23). When this weighting is included, the shear average estimator is given by

$$
\widehat{t}_0 = \frac{1}{\omega_0}\mathbf{Arg}\left\{ \frac{\sum\limits_{n=0}^{N-1} \beta[n]\ p_0[n]p_1^*[n]}{\sum\limits_{n=0}^{N-1} \beta[n]} \right\},
\tag{9.2}
$$

where $\beta[n]$ is a weighting function to reduce the effect of noise and strong-scattering targets.

Originally the weighting, $\beta[n]$, was chosen to be 1, to give the maximum likelihood (ML) estimator [Johnson et al. 1995]. This has the effect of putting increased weighting on strong-scattering targets—those targets that suffer the least from additive noise but cause biasing. Giving a large weighting to the targets which cause bias is undesirable.

Using the prior information that strong-scattering targets skew the estimation and weak targets suffer from additive noise, a weighting of

$$
\beta[n] = \frac{1}{\alpha + |p_0[n]p_1^*[n]|},
\tag{9.3}
$$

is chosen where the α parameter is selected so that the peak weighting is well above the noise floor. The idea behind a weighting of this form is that the contributions from strong-scattering targets are reduced and that noise affected targets are suppressed. This can be looked at as weighting the estimator so that there is a high probability

(a)   (b)

(c)   (d)

Figure 9.4   Various amplitude weightings chosen for shear average phase estimation. Graph shows weighting placed on a target of cross-correlation intensity $|p[n]|^2$—this is equivalent to $\beta[n]\,|p[n]|^2$. (a) Maximum Likelihood (ML) estimator [Johnson et al. 1995]. (b) Suppresses noise (low intensity data) via (9.3). (c) Remove the effect of large targets via (9.4). (d) Weight all targets evenly via (9.5)

the clutter used for the estimation is delta-correlated. The weighting chosen is but one of many possible and methods using the local image statistics for a bias-reduction weighting may prove more worthwhile. Examples of other weightings tried are:

$$\beta[n] = \frac{1}{\alpha + |p_0[n]p_1^*[n]|^2},$$ (9.4)

[Callow et al. 2001c] which attempts to remove the effect of strong-scattering targets altogether, and

$$\beta[n] = \frac{1}{|p_0[n]p_1^*[n]|},$$ (9.5)

which weights the contribution of all targets equally. The weighting (9.4) was found to give the best results on field-collected data.

Ideally, the weightings would be used in the reconstructed image domain with data mapped through an inverse reconstruction technique[11] before using autofocus. The benefit in the approach would be to further localise the blurring-causing strong-

---

[11]It is possible to derive versions of the common image reconstruction algorithms that have the ability to generate raw data from a given reconstructed image.

scatterers but at the cost of increased processing and the loss of the ability to use simple phase estimation. Shear average is often used for bulk motion compensation (prior to other autofocus) and the added accuracy is not worth the computational burden. However, using a weighting in shear average when operating on pulse-compressed data provides a significant performance improvement with negligible computational cost.

9.3.2   Shear average Results

This section provides results from the operation of the modified shear averagealgorithm on simulated and field collected imagery.

Simulated imagery

Simulated sonar data was generated with background clutter modelled as many point targets. Strong targets and a regular extended target were then added to the clutter background giving a simulation image similar to some of the real data obtained with the KiwiSAS-II sonar. KiwiSAS-II sonar parameters were used in the simulation (see Appendix A). Additive white Gaussian noise was added to the simulated echo data at 30 dB lower than the average clutter signal. Path distortion was included as a timing-error, with the distortion used shown in Figure 9.5(b). The path distortion was generated using filtered white Gaussian noise with the amplitude and bandwidth chosen to match the path distortions observed in field-collected data.

Figures 9.5(a) and 9.5(c) show the ideal pulse-compressed and reconstructed images and the path distortion used for testing is shown in Figure 9.5(b). Figure 9.5(d) shows the reconstructed imagery with uncompensated path-deviations. Severe image distortion is apparent in the images as smearing in the along-track direction.

Smearing is also apparent in the range direction, this is due to the low Q of the system in question ($Q \approx 1.5$). The range smearing is caused by motion errors altering the along-track Doppler spectrum. The Doppler information is used to determine the range shift necessary to focus a target in the reconstruction. When distorted incorrect range shifts are used causing range smearing. Further description of the cause of the range smearing is provided in [Hawkins 1996] and in [Sutton et al. 2002].

After autofocusing using the original shear average, the resulting image, Figure 9.5(e), is also distorted. Although reduced with respect to Figure 9.5(d), the distortion obliterates scene detail. The poor autofocus is due to the biasing effect of large targets. Once the new amplitude weighting (9.4) is included in the shear average, autofocus performs better. Figure 9.5(f) shows the improvement in autofocus performance after the biasing by large targets is reduced. Residual blurring is evident and some of the targets locations are displaced. Poor autofocus performance is expected near the edges of the scene.

(a)   (b)

(c)   (d)

(e)   (f)
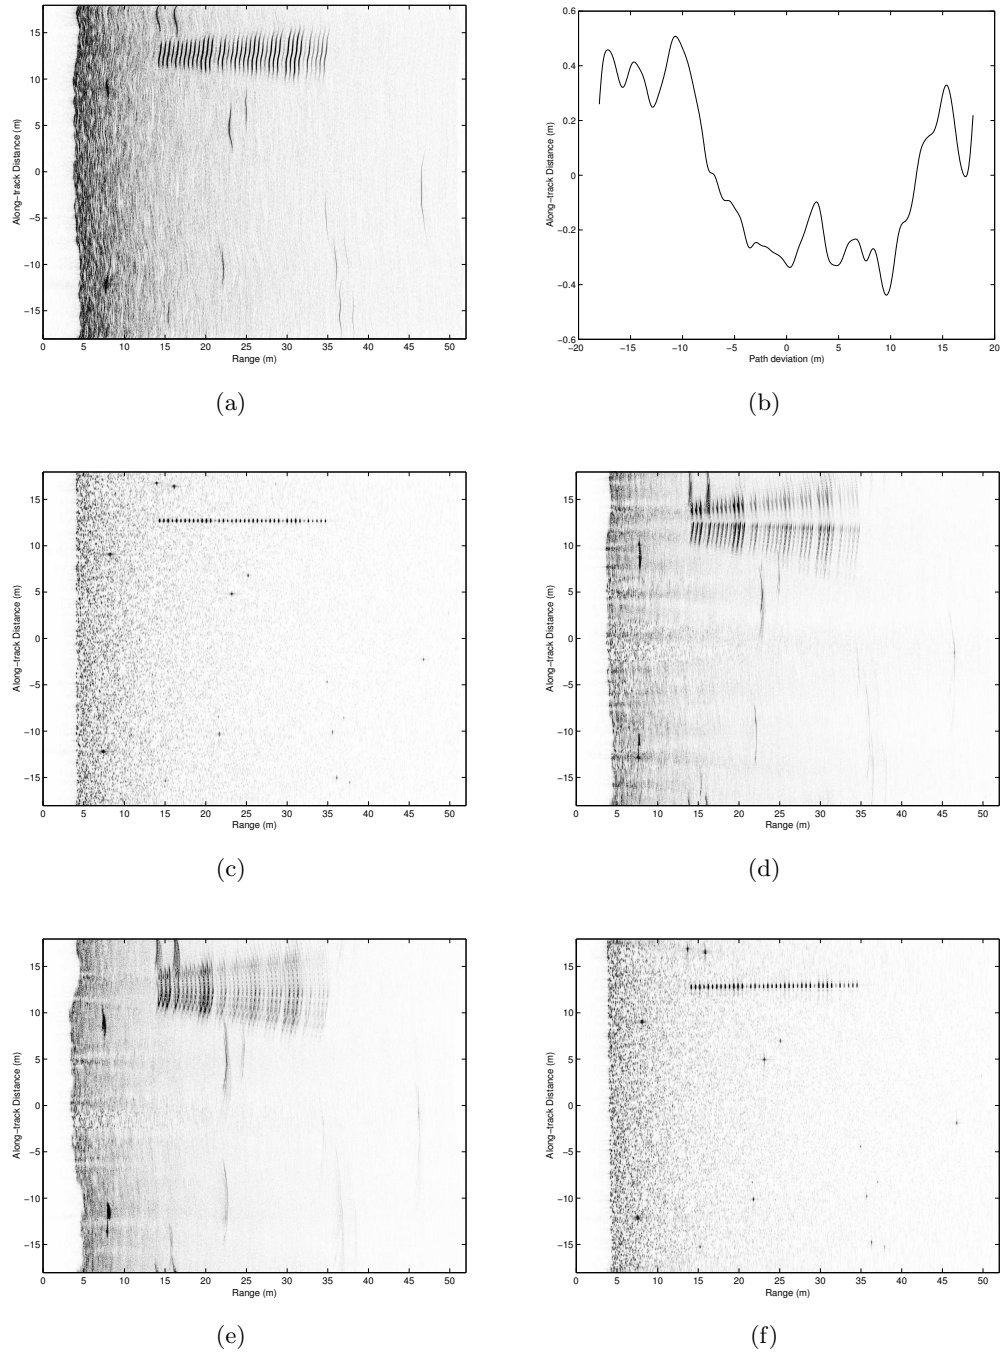
Figure 9.5   Operation of shear average on simulated data. (a) Raw pulse-compressed data with no distortion added. (b) Sway distortion added to the raw sonar data (a) as a timing error. (c) Ideal reconstructed image. (d) Reconstructed image with across-track path distortion shown in (b) (no autofocus). (e) Reconstructed image after autofocus using Johnson95 shear average [Johnson et al. 1995]. (f) Reconstructed image after autofocus using shear average with amplitude weighting of (9.4).

Field-collected data

Field-collected data was obtained in Sydney Harbour during July 2001 using the KiwiSAS-II sonar. The data presented is wide-band, low-Q data (bandwidth = 20 kHz and f_0 = 30 kHz) with a receive aperture of D = 0.33 m giving a null-to-null beam-width of about 20°. Along-track sampling is at approximately D/3 corresponding to an along-track velocity of $1.4\,\mathrm{ms}^{-1}$.

Figure 9.6 shows the imaged scene without any autofocus applied. Two cross-talk bands are apparent in the image, the first at x= 0.5 m from acoustic cross-talk and another at $x \approx 0$ m from electronic cross-talk. The image is a low-grazing angle image with the sonar flying approximately 1.5 m above a compacted sand sea-floor (apparent from the first-bottom bounce). Note the strong-scattering, extended sonar calibration rail used in the experiment evident at $y = 10$ m. The rail consists of a number of small hydrophones spaced at 1 m intervals [Ferguson and Cleary 1999] these should appear as small reflectors spaced 1 m apart. The other targets in the scene are unknown. Also of note in the image is a pulse-compression artefact (from circular convolution in range) apparent at $51\,\mathrm{m}^{12}$.

Figure 9.7 demonstrates the improvement in the final image once shear average autofocus has been applied. The scene contains strong-scattering targets and the weighting from (9.4) has been used. The autofocused image is much sharper over the region y = 5–15 m and has a resolution 3–4 times better than before autofocus. Less dramatic sharpening is seen over the lower half of the image $y <$ -5 m and slight degradation occurs at approximately $y = 0$ m. Timing-error compensation (without slant-range correction) has been used and this is apparent in the movement of the cross-talk bands. The uneven performance (and better performance over the strong-scattering targets) is attributed to system yaw. The shear average method interprets yaw (incorrectly) as sway. The removal of this incorrect sway causes residual biasing.

Results when using the Johnson95 shear average weighting are similar to those in Figure 9.5(e) and are significantly worse than those illustrated here. This is due to the biasing caused by the strong-scattering reflectors apparent in the image.

The data presented in this section was collected with the assistance of the Australian Defence Science Technology Organisation (DSTO).

9.3.3   Multi-band estimation

Multi-band shear averaging is a way of reducing the problem of phase unwrapping and ambiguity in shear average autofocus. With the current incarnation of the KiwiSAS

---

[12]This artefact is visible because side-lobes of the chirp-signal correlation (autocorrelation side-lobes) have been circularly wrapped during Fourier domain correlation. The artefact is also elevated in strength because of TVG applied to the image. Its appearance could be avoided by using guard bands during Fourier correlation or the using time-domain correlation.
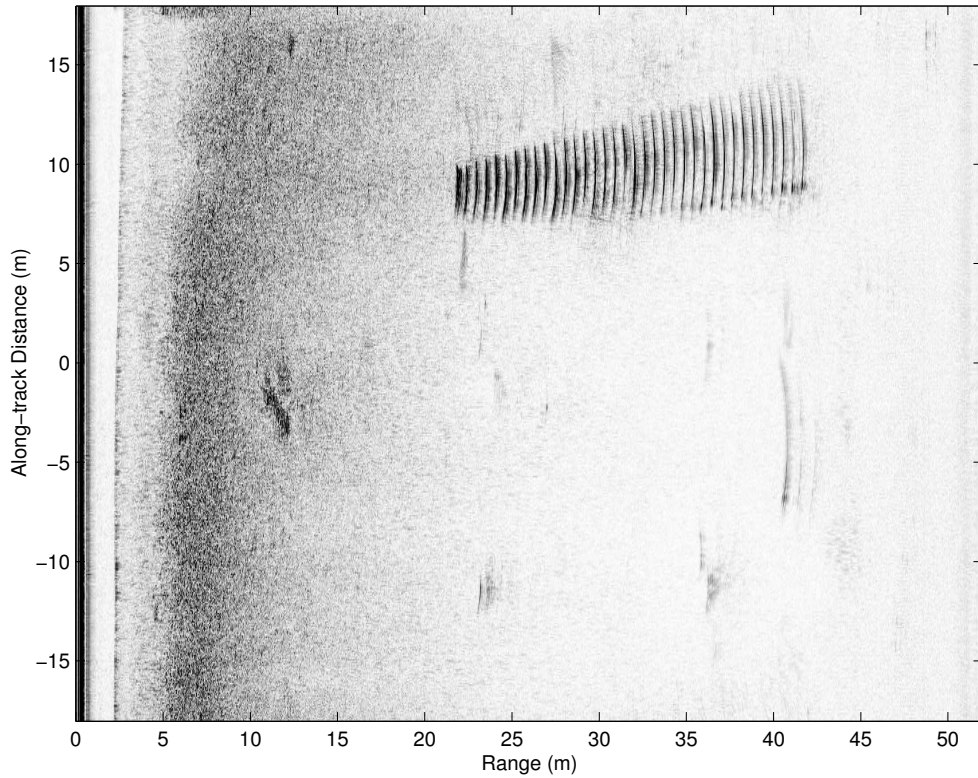
Figure 9.6 Reconstructed Sydney Harbour data before autofocus.

simultaneous high-frequency and low-frequency scene imagery is available. The advantage this gives is the ability to initially phase unwrap using low-frequency data. This reduces the problem with 2π ambiguities as the carrier phase is much lower. Following the first step, high-frequency data is used for shear average path-deviation estimates.

There are a number of advantages of using the high-frequency-band data for autofocus. Once ambiguities have been removed using the low-frequency data, shear average autofocus can attain more accurate results due to the shorter carrier wavelength. The Cramér-Rao lower bound on the estimator variance is given in (7.24),

$$\text{CRLB} = \frac{1}{\omega_0} \frac{1}{\sqrt{B_c \tau_c}} \sqrt{\frac{1}{\text{SNR}} + \frac{1}{2\text{SNR}^2}}, \tag{9.6}$$

where SNR is the clutter to noise ratio[13]. So in ideal conditions, higher frequencies give better autofocus. The improvement for the KiwiSAS-III is a factor of three. In the high-frequency-band the raw echo data also better fits the delta-correlated model. This is because the sonar beam is narrower so the biasing effect of the targets is reduced.

---

[13]For the KiwiSAS-II, blur-free imagery requires a clutter-to-noise ratio of >40 dB. Ratios of this order for effective *clutter*-to-noise levels are very challenging.

Figure 9.7   Sydney Harbour data after shear average autofocus using weighting of (9.4)

(a)   (b)

Figure 9.8   Operation of multi-band shear average on simulated data of Figure 9.5. (a) Estimated sonar sway using 30 kHz and 100 kHz shear average separately. Note the large error in the 100 kHz result caused by incorrect phase unwrapping in the ambiguous phase-only estimate. (b) Estimated wobble with 30 kHz shear average and the combined 30k/100 kHz estimate. The 100 kHz imagery allows much more precise sway estimates when the bulk sway is estimated using lower frequency data.

CNR	SA-J95	SA-W	SA-NW	Non-coherent
No noise	1.3	0.50	0.17	0.019
30 dB	1.3	0.20	0.16	0.019
15 dB	0.58	0.39	0.17	0.019
0 dB	0.70	2.7	1.6	0.031

Table 9.1   The RMS error (in metres) in bulk sway estimation of autofocus methods at various clutter to noise ratios (CNR). The simulated scene consisted entirely of uniform clutter and additive white Gaussian noise. Note the usual phase unwrapping step in the phase estimators has been removed. SA-J95, Johnson95 shear average algorithm with original weighting. SA-W, shear average algorithm with the weighting of (9.4). SA-NW, shear average algorithm with no weighting applied. Non-coherent, algorithm proposed in Section 9.3.4.

9.3.4   Non-coherent estimation

In the absence of INS data, it is possible to have gross un-corrected path-deviations in the sonar image. Severe blurring occurs when gross uncorrected path-deviation exists; in some cases completely obscuring scene detail. As a result, a fast algorithm for the removal of gross distortion is required. In addition, many other autofocus algorithms, such as statistical autofocus [Fortune et al. 2001a] and phase matching autofocus (PMA) [Gough et al. 2000a], require less computation when starting close to the final solution. Non-coherent autofocus provides a sound starting estimate for the gross path-deviations at little computational cost.

Non-coherent autofocus operates by correlating adjacent pulse-compressed sonar pulses. Non-coherent autofocus then picks the sub-pixel shift in the correlation-peak *without* using the phase information. Large path-deviations can now be considered without phase-unwrapping (see Section 7.6.1). This makes noncoherent autofocus ideal for removing bulk-motion errors from data as a first processing step.

One added benefit of using noncoherent autofocusing is that the correlation length is extended. This allows path-deviation estimation to work when cross-correlating sonar pulses spaced at more than $D/2$ apart (where D is the extent of the receive aperture). This provides means of autofocusing under-sampled synthetic apertures. Surprisingly, it is precisely the biasing effect of strong targets, covered in Section 9.3.1, that allows this advantage. These targets are visible for many pulses and they provide information on gross path-deviations. Multiple noncoherent estimates for differing separations can increase performance (using the same principles as phase closure to improve performance—see Section 7.6.2).

Table 9.1 shows the performance of the noncoherent autofocus against that of other shear average based algorithms. For this simulation additive white gaussian noise was added to a uniform clutter scene. The noise was added at clutter-to-noise ratios of 0 dB, 15 dB, and 30 dB. The values recorded in the table are the root mean squared (RMS) errors of the estimates versus the injected path-deviation (in metres).

The results indicate that non-coherent autofocus outperforms phase-based shear average algorithms in estimating gross path-deviation. The poor performance of the phase methods is due to the lack of phase unwrapping and the resultant estimator ambiguity. The phase techniques provide much better accuracy when used in conjunction with non-coherent techniques or if phase unwrapping is applied. Phase techniques with appropriate unwrapping are better than non-coherent techniques by the system Q-factor—e.g., phase autofocus of the KiwiSAS-III 100 kHz imagery is 6-10 times more accurate than non-coherent autofocus. For comparison purposes the phase method CRLB at a CNR of 30 dB is 0.0013 and at 0 dB is 0.0078; amplitude CRLBs are 0.0072 and 0.041 for the same CNRs. It can thus be seen that in the presence of gross path deviation the non-coherent estimation runs close to its CRLB whereas the ambiguity in position caused by a lack of successful phase unwrapping severely degrades traditional, phase-based, autofocus performance.

9.4   Summary

Inertial navigation systems are not accurate enough to allow SAS imaging alone, and micronavigation techniques are a method of supplementing INS accuracy using the received data. Typically these techniques rely on exploiting redundancy in the echo data (often using echo-correlation).

The commonly used RPC echo-correlation algorithm has been summarised and linked to related image correlation and shear average techniques. Beam correlation and image correlation provide extensions of RPC with improved ability to operate on directive data. Image correlation also allows operation without requiring receiver overlap and so can operate with $D/2$ sampled data. The autofocus implications of non-redundant collection on image correlation have been explored and were found to be equivalent to those affecting traditional autofocus algorithms. Image correlation algorithms provide a common ground between autofocus relying on echo-correlation and those relying on scene redundancy.

Single receiver shear average was discussed and modifications allowing operation on field-collected data in the presence of large path motions presented. Other shear average extensions take advantage of the simultaneous dual frequency capability of the KiwiSAS-III and it large bandwidth to remove phase wrapping ambiguities. Both of these techniques enhance the accuracy of shear average. Experimental results from using the modified shear average algorithm on both simulated and field-collected data have been included. From these results it is apparent that shear average alone is not enough to provide diffraction-limited imagery in the presence of motion errors. Similar results were reported for spotlight SAR operation [Fienup 2001; Wahl et al. 1991]. This and its computational efficiency mean the shear average algorithm is best suited to a "first-cut" or bulk autofocus solution prior to other means of autofocus.

Stripmap autofocus

This chapter outlines procedures for implementing phase gradient based stripmap autofocus. Initially, the stripmap blurring model and the failure of spotlight algorithms are discussed. The wavenumber transform used in the stripmap blurring model is considered and the approximations inherent in its operation discussed. Summaries of traditional stripmap autofocus techniques are presented and comments made regarding their use in SAS autofocus. A new algorithm, *stripmap phase gradient autofocus* (SPGA), is presented and traditional algorithms are shown to be special cases of SPGA. The salient operations of SPGA are discussed in detail from a SAS autofocus perspective. Analysis is provided for improving the use of the wavenumber transform and investigating autofocus difficulties such as aperture undersampling and range offsets. Finally, results of using SPGA and various traditional autofocus techniques on simulated and field-collected data are presented.

10.1 Stripmap blurring model

The model derived in this section describes the blurring of a localised section of a stripmap image. Modelling image blurring in stripmap systems is more difficult than in similar spotlight systems. The convolutional nature of spotlight imagery no longer holds and expressions for blurring become complicated. Range-variance and the along-track position of targets must both be considered when modelling stripmap SAS blurring.

The along-track sonar position, u, and the 2-D spatial-frequency variables for image point at position $(x_n,\ y_n)$ are related by the stripmap wavenumber transform (see Appendix D)

$$u = y_n - \frac{k_y x_n}{k_x}, \tag{10.1}$$

$$\omega = \frac{c}{2}\ \sqrt{k_x^2 + k_y^2}, \tag{10.2}$$

These equations relate a given point in the reconstructed image to the pulse compressed

raw data for that point via

$$S_n(\omega, u) = f_n(k_x, k_y)\big|_{k_x = \sqrt{4k^2 - k_y^2}, k_y = \frac{k_x(y_n - u)}{x_n}}, \tag{10.3}$$

which has an inverse mapping of

$$f_n(k_x, k_y) = S_n(\omega, u)\big|_{\omega = \frac{c}{2}\sqrt{k_x^2 + k_y^2}, u = y_n - \frac{k_y x_n}{k_x}}, \tag{10.4}$$

where $f_n(x, y)$ is the image at point x_n, y_n and $S_n(\omega, u)$ is the Fourier transform of the pulse compressed data corresponding to that point.

The relationships above are equivalent to the coordinate transforms used by the wavenumber algorithm (see Chapter 5) and hold within the approximations used for that algorithm so long as the target patch is a single point. The analysis can be extended to a small patch around the point in question by assuming that the position offsets are large compared to the patch size and that the receiver is in the far field of the patch. Note that these assumptions resemble those used in spotlight autofocus but apply equally well to the small region around a single point in a stripmap image as long as the patch is small. The technique is very similar to digital spotlighting reconstruction methods.

The mapping between the small image patch and pulse compressed data is used to derive the effect of sway on individual pieces of the reconstructed image. The wide-beam, stripmap model of sway $X(u)$ for the image point at an angle θ to broadside (see Chapter 6) is given by

$$\widetilde{S}(\omega, u) \approx S(\omega, u)\exp\left(j2kX(u)\cos\theta\right). \tag{10.5}$$

The geometrical properties of the stripmap wavenumber transform (called the wavenumber transform for the remainder of the thesis) are invoked to describe the blurring of individual sections of the reconstructed image via (10.5). Mapping the blurred pulse-compressed data for a small image section $\widetilde{S}_n(\omega, u)$ through the wavenumber transform and substituting

$$k_x = 2k\cos\theta, \tag{10.6}$$

allows the blurring to be expressed as

$$\widetilde{f}_n(k_x, k_y) \approx f_n(k_x, k_y)\exp\left(jk_xX\left(y_n - \frac{k_y x_n}{k_x}\right)\right). \tag{10.7}$$

Note that the wavenumber transform, and thus the blurring model, is approximate when sway exists. This is because the stationary phase point of the signal is altered by the sway. Section 10.1.1 discusses this phenomenon in more detail.

The implication of the blurring model (10.7) is that for an image patch at a given

location, (x_n, y_n) it is possible to calculate target blurring given a path distortion $X(u)$. Conversely, given the blurring, it is then possible to estimate a path that caused that blurring.

10.1.1 Wavenumber transform and the small sway approximation

The true blurring model for phase distortions is more complicated than (10.7) implies. The derivation of the blurring model assumes that the phase gradient caused by the distortion is small—so that the stationary phase model[1] can be used to derive the wavenumber transform. Failure of the above blurring model is particularly evident in stripmap systems. This can be seen, for example, when the sway consists of a linear crabbing (i.e., a linearly increasing sway): the measured along-track spectrum is shifted by a constant amount (see Chapter 6) and the image skews.

The blurring model limitations are uncovered by considering the interaction of sway and the measured along-track spatial frequency. The measured along-track spatial frequency, k_y', is a function of both the true along-track spatial frequency, k_y, and the sway, $X(u)$. Shifts of the measured along-track spatial Doppler spectrum are caused by sway-induced phase gradients. When significant sway exists, these spectral shifts cause a failure in the approximations used in deriving (10.7). The accuracy of the blur modelling is degraded in such conditions.

The error caused by frequency shifting can be corrected with better modelling. The sway-induced difference between $S_n(\omega, u)$ and $\widetilde{S}_n(\omega, u)$ is a phase modulation term

$$\exp\left(jk_x X(u)\right).$$

In this case, the instantaneous frequency shift Δk_y in the blurred pulse-compressed image, $\widetilde{S}_n(\omega, u)$, is given by

$$\Delta k_u = \Delta k_y = \frac{\mathrm{d}}{\mathrm{d}u}\mathbf{Arg}\left\{\exp\left(jk_x X(u)\right)\right\} \tag{10.8}$$

$$= \frac{\mathrm{d}}{\mathrm{d}u}(k_x X(u)), \tag{10.9}$$

where $\mathrm{d}/(\mathrm{d}u)(k_x X(u))$ is the sway-induced phase-derivative measured as a function of u. The measured instantaneous along-track spatial frequencies can now be expressed as the sum of the original instantaneous spatial frequencies and those caused by sway

$$k_y' = k_y + \Delta k_y. \tag{10.10}$$

This expression combined with the traditional wavenumber transform accurately predicts the shifting of targets under a linear sway. Thus it can be shown that a linear

---

[1]Note that the approximation also affects geometry based derivations.

sway of $X(u) = \gamma u$ results in an along-track position shift of

$$\Delta y = \frac{\Delta k_y x_n}{k_x} \tag{10.11}$$

$$= \gamma x_n. \tag{10.12}$$

This is the same expression for position-shift derived using different methods [Callow et al. 2002a]. Higher order phase variations cause parts of the target to appear at different along-track positions—the familiar image blurring to be removed via autofocus processing.

For autofocus the inverse model is more interesting:

$$k_y = k'_y - \Delta k_y. \tag{10.13}$$

For small Δk_y this can be used to estimate the true spectrum from the distorted measurements. This is described further in Section 10.3.6.

The failure of the stationary phase approximation with large phase gradients leads to some of the major difficulties of stripmap autofocus. In particular, targets imaged with a system undergoing a linear sway appear at incorrect along-track positions. Better modelling of the blurring process allows more accurate phase-error removal by using an improved wavenumber transform (i.e., phase-gradient estimates are translated into sway distortions more accurately).

10.1.2   Wide-band vs narrow-band blurring

Much of the SAS and SAR autofocus literature considers the narrow-band blurring appropriate for high-Q systems. This treatment makes the assumption that the image blurring is contained in 1-D, i.e., blurring only occurs in the along-track direction.

Blurring in wide-band (low-Q) systems is 2-D, causing some of the traditional autofocus assumptions to break down. Despite this, the 2-D blurring encountered in wide-band systems is not a limiting factor in autofocus design. Bandwidth subdivision can be applied to simplify the problem to multiple narrow-band 1-D autofocus problems [Chevillon et al. 1998].

A better method of avoiding problems caused by 2-D blurring is simply to consider the effect during algorithm development. The unified stripmap PGA (SPGA) framework discussed later in the chapter outlines methods for accounting for the 2-D blurring. Importantly, the *last* iterations of autofocus algorithms only need to remove small distortions[2]. When the distorting sway is small compared to the across-track resolution, wide-band blurring can be treated as narrow-band using the quasi narrow-band formulation. This condition is always true for near-focused imagery since blur-free

---

[2]Assuming algorithm convergence.

imagery requires $X(u) < \lambda/16$, whereas $\lambda < \delta_x$ and thus the sway is much less than the resolution $(X(u) \ll \delta_x)$. Therefore, the important final iterations of autofocus algorithms only need account for narrow-band, 1-D blurring.

To summarise, the impact of neglecting the wide-band nature of the autofocus problem is not severe. However, autofocus performance is degraded in the *initial* iterations. This slows algorithm convergence. The slower convergence does not affect the final autofocus result except to increase computation.

10.1.3   Wide-beam vs narrow-beam blurring

Wide-beam systems use longer apertures than narrow-beam systems, making the SAS navigation constraints more stringent. Additionally, wide-beam systems are more likely to suffer from the narrow-beam timing-error approximations often made in autofocus algorithms. These differences are enough to alter autofocus performance and should be considered in autofocus algorithm design.

The narrow-beam approximation is used in most previous autofocus algorithms, such as those outlined by Hawkins [1996], Chevillon et al. [1998], and Pat [2000]. It is a good approximation when considering past SAS systems with beam-widths narrower than $10°$. The approximation is valid for both the KiwiSAS-II and KiwiSAS-III sonars, and also the rail-based SAS described by Chevillon et al. [1998]. However, the approximation is inappropriate for newer SAS systems, such as those proposed by Keeter [2001] and Hagen et al. [2001]. Wide-beam autofocus techniques are needed for wide-beam sonars.

Designing autofocus algorithms for wide-beam systems requires replacement of the timing-error approximation. The timing-error approximation is covered in detail in Section 6.3. Summarising the ideas of that section, the differences in blurring model can be described by making the narrow-beam approximations $\cos\theta \approx 1$, and $k_x \approx 2k$ in (10.5). Thus the narrow-beam blurring model is given by

$$\widetilde{S}(\omega, u) \approx S(\omega, u) \exp\left(j2kX(u)\right). \qquad (10.14)$$

Note that medium fluctuations should not be modelled as position errors via (10.14) or via (10.5).

Wide-beam blurring has two effects on autofocus: algorithms must account for the altered phase response when calculating the path motion, and secondly motion compensation must account for the extra beam-width. The former is straightforward if the wide-beam blurring model (10.5) is used instead of (10.14). Motion compensation appropriate for wide-beam systems is discussed in Chapter 6.

Attempting narrow-beam autofocus on wide-beam data results in a performance degradation. Algorithms using narrow-beam modelling require more autofocus itera-

tions to converge. Wide-beam modelling in contrast is more accurate and requires little extra computation; it should always be used when autofocusing wide-beam data.

10.1.4 The failure of the spotlight blurring model for stripmap systems (PGA failure)

The spotlight blurring model (as used by the algorithms in the previous chapter and discussed in Section 8.1) is unable to describe the blurring observed in stripmap systems. The model failure prevents the unaltered use of spotlight autofocus algorithms. This section demonstrates the model failure and the inability of spotlight PGA to resolve stripmap path errors.

The spotlight wavenumber transform is obtained from the stripmap wavenumber transform (10.1) using the approximations $y_n \approx 0$, $k_x \approx 2k_0$ (see Section 8.1), this is summarised as

$$u = y_n - \frac{k_y x_n}{k_x}, \tag{10.15}$$

$$\approx -\frac{k_y x_0}{2k_0}. \tag{10.16}$$

The most important of these approximations is that $y_n \approx 0$ [Bonifant 1999]. The neglected along-track target offsets y_n change the scale and shift between aperture position u and spatial frequency k_y into a simple scaling. In stripmap systems, the neglected term is responsible for the different blurring at different along-track positions since each target sees a shifted sway (the sway relative to its along-track position y_n). The other major spotlight approximation (range-invariant blurring) neglects the varied scaling between along-track spatial frequency and aperture position—i.e., it fails to account for targets at longer range having larger aperture coverage (slower azimuth-chirp rates [Thompson et al. 1999]). These approximations are important if large swath widths are to be autofocused. Traditional PCA is designed to operate under both of the approximations mentioned. Further narrow-band and narrow-beam approximations are also made but have a relatively minor effect (the effects of these are discussed in Sections 10.1.2 and 10.1.3).

The effect of stripmap blurring is demonstrated in Figures 10.1(a) and 10.1(b). Figure 10.1(a) shows a reconstructed stripmap image suffering motion blurring. Note that the blurring of each target is different—i.e., the blurring is space-variant. In particular, the blurring at longer range is more severe. This is caused by the varying scale between aperture position and spatial frequency—targets at further range "see" a longer aperture so the blurring order is higher for those targets than targets at short range (a longer aperture has more cycles of sway than a shorter aperture). To demonstrate the failure of spotlight autofocus algorithms, Figure 10.1(b) overlays the PGA derived sway against spatial frequency for each of the targets in the image. Note

that each target shows a different sway estimate. The range redundancy required for PGA operation disappears since the mapping between spatial frequency and the injected sway is no longer a simple scaling. PGA thus fails to provide reasonable sway estimates for stripmap systems.

10.1.5 Phase curvature autofocus (PCA) blurring model

The stripmap PCA blurring model [Hawkins 1996, pp 165–168] is equivalent to using the spotlight wavenumber transform and accounting for the y_n term and the range variance in blurring. Substituting narrow-band, narrow-beam approximations into (10.1) leads to the blurring model

$$u = y_n - \frac{k_y x_n}{k_x}, \tag{10.17}$$

$$\approx y_n - \frac{k_y x_n}{2k_0}. \tag{10.18}$$

The PCA model, like the spotlight model, assumes that path distortions can be detected as phase terms in the spatial frequency domain. Instead of using the wavenumber transform directly, PCA chirps the image in azimuth (using the chirp rate $x_n/(2k_0)$); this spreads spatial frequency information to the extent of the original aperture coverage. The azimuth-chirping operation is described by

$$f_c(x,y) = f(x,y) \odot_y \exp\left(-j2k_0\left(\sqrt{x^2 + y^2} - x\right)\right), \tag{10.19}$$

and can be seen to be a narrow-band approximation to inverse reconstruction (see Section 10.2.2). The PCA blurring model then assumes that sway motions are apparent as time shifts to the azimuth-chirped image (10.19), i.e., [Hawkins 1996, page 166]

$$\widetilde{f}_c(x,y)_{\text{narrow-band}} = f_c(x,y) \exp\left(j2k_0 X(y)\right). \tag{10.20}$$

The wide-band version of (10.20) is given by [Hawkins 1996, page 171]

$$\widetilde{f}_c(x,y)_{\text{wide-band}} = f_c(x,y) \odot_x \delta(x - X(y)). \tag{10.21}$$

This implies sway errors can be detected as a time shift in the azimuth-chirped image.

Failure of the PCA blurring model

The wide-band PCA blurring model of (10.21) has flaws when used for autofocus. Most importantly, the narrow-beam approximations are severe, causing the model to fail.

The assumption that the *azimuth-chirped* image is time shifted by sway is invalid for most SAS systems. Instead, the *pulse-compressed* data is time shifted by the sway (under the timing-error approximation). Moreover, since the azimuth-chirped image is

(a)

(b)

Figure 10.1 Spotlight blurring model breakdown for stripmap autofocus. (a) Blurred stripmap image. Note how the blurring is different for targets at different along-track and range positions. Spotlight autofocus is designed to treat space-invariant blurring rather than the space-variant blurring seen here. (b) Overlay of point response sways. The phase error versus spatial frequency is different for each target. The mapping between aperture position u and spatial frequency is no-longer a simple scaling and range redundancy disappears. Without redundancy in range (as occurs with stripmap data) spotlight autofocus procedures fail.

generated from the pulse-compressed image via the interpolation of the range-Doppler algorithm[3] the blurring models are quite different. The PCA blurring model neglects the effect of the interpolation. In addition, sway causes frequency shifting in the pulse-compressed data affecting the interpolation[4]. This shifting is also neglected in the PCA blurring model. The wide-beam, wide-band model discussed in Section 10.1 avoids the approximations of the PCA blurring model and provides a better description of the blurring process.

The problems with the PCA blurring model are highlighted by simulated imagery from a system corresponding to the KiwiSAS-II sonar. A series of point scattering objects was simulated and corrupted with the sway depicted in Figure 10.2(a). Figure 10.2(b) is the azimuth-chirped reconstructed image without sway injected. Figure 10.2(c) has had sway injected via the wide-band PCA model of (10.21) and Figure 10.2(d) has had sway injected as a timing-error in the pulse compressed data. Note the significant difference between Figures 10.2(c) and 10.2(d). These would be identical if the model were accurate for this system. Thus the KiwiSAS-II system has a wide enough beam-width that the PCA blurring model fails[5]. The wide-band, wide-beam wavenumber transform used throughout the remainder of the chapter provides for the effect of sway on the reconstructed image better than the PCA model.

10.2   Traditional stripmap autofocus algorithms

This section summarises some of the traditional SAS autofocus algorithms in use. The algorithms are presented in a consistent notation with the aim of demonstrating similarities to the SPGA algorithm to be described in Section 10.3. The algorithms have been derived and developed from traditional spotlight SAR techniques of PPP and PGA. The SAS versions are often extensions of PPP/PGA to account for wide-band, wide-beam, stripmap imagery.

10.2.1   Prominent point positioning (inverse filter)

Prominent point positioning (PPP) is an autofocus technique used in both SAR and SAS imagery [Muff et al. 1995; Shippey et al. 2001; Soumekh 1999] and is sometimes called *inverse filtering* [Jakowatz et al. 1996]. A variation of PPP that calculates the distortion in the wavenumber domain (for use in wide-band SAS systems) was presented by Chevillon et al. [1998].

---

[3] The last step of the range-Doppler algorithm is to remove the azimuth-chirp used in (10.21)(see Section 4.5, [Hawkins 1996, page 166],[Gough and Hawkins 1997]).

[4]This is one description for the cause of 2-D smearing of blurred wide-band imagery—see [Sutton et al. 2002] and Section 10.1.1.

[5]Note that the (narrow beam) timing-error approximation still holds for the KiwiSAS-II system.
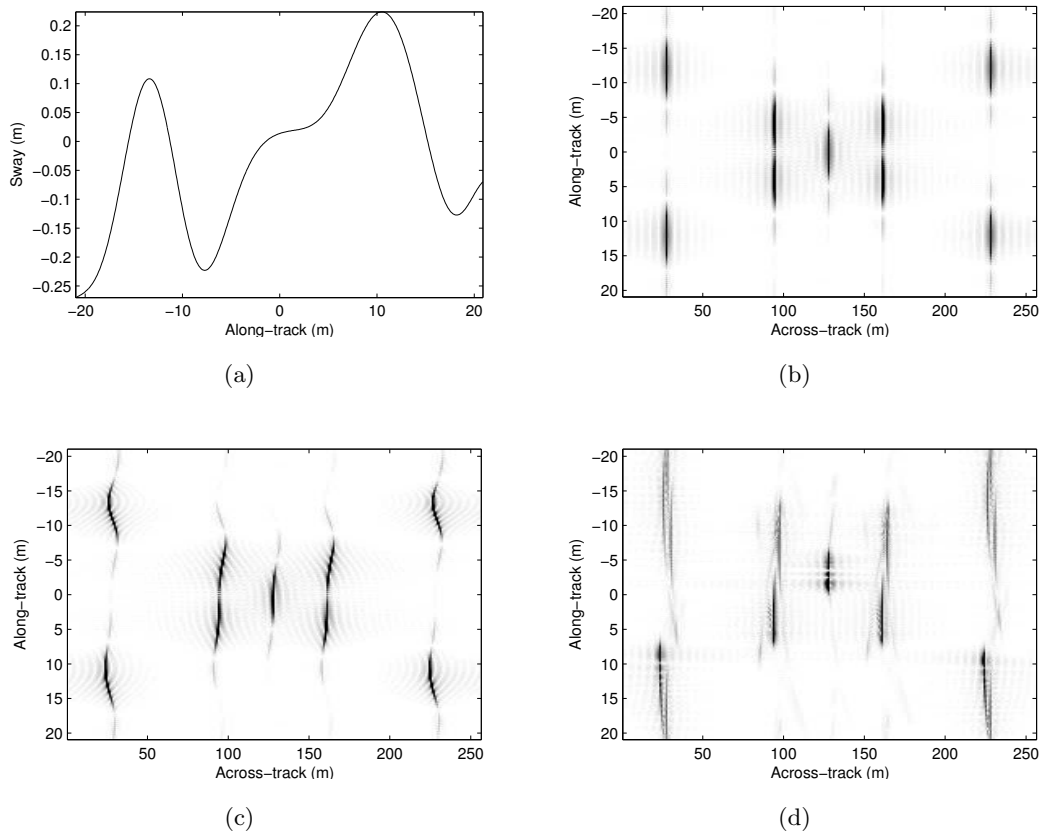
Figure 10.2 Failure of the PCA blurring model. (a) Sway distortion used in the simulation. (b) Azimuth-chirped ideal image (image without sway injected). (c) PCA blurring model—sway injected into the azimuth-chirped ideal image. (d) Narrow-beam blurring model—sway injected as timing-error into the pulse compressed data before reconstruction. Note the differences between (c) and (d). In particular, note that the targets have a larger extent in (d). This is due to the sway induced phase causing a change in the measured spatial frequencies and affects the azimuth-chirp (and reconstruction).

PPP assumes that the brightest target in the image is a point scatterer[6]. The technique tracks the aperture phase of the target and estimates the platform path such that there is no phase deviation from that of an point scatterer. In spotlight mode SAR, where images have space-invariant blurring, PPP provides successful path estimation and blur removal for the entire image. With the selection of multiple points[7], PPP is able to estimate sway, surge, and heave motions [Carrera et al. 1995].

Single scatterer PPP is loosely equivalent to spotlight PGA when using only the strongest range bin and an altered phase estimator. PPP does not average target signatures removing the need to calculate phase gradients [Carrera et al. 1995]. This allows the use of phase averaging and thus increases the accuracy of the phase estimates by avoiding the inherently noisy phase gradient calculation. However, accuracy is sacrificed by not averaging estimates from a number of targets and by neglecting

---

[6]Stoyle [1998] describes the implications of this assumption in SAR imaging.

[7]In spotlight SAR systems.

windowing. The result is that PPP has lower accuracy than PGA for most scenes— particularly those predominantly containing clutter. PPP is a special case of PGA using only a single scatterer and has an associated loss of accuracy.

Targets imaged using stripmap systems only provide information over a small section of the aperture. Thus PPP for stripmap systems requires at least one prominent point for each independent synthetic aperture created—i.e., at least one prominent point should be visible for the entire imaging duration. Moreover, combination of motion estimates from a number of prominent points is difficult so PPP is rarely used for a full aperture motion estimate (an example of PPP use in stripmap SAS is presented by Shippey et al. [2001]). The use of the SPGA framework described in Section 10.3 simplifies the problem. Furthermore, the SPGA framework can help solve other problems related to estimate weighting and reliability.

PPP represents a simple, computationally-inexpensive technique for SAR/SAS autofocus that is closely related to both image correlation (see Section 9.2 and [Wilkinson 2001]) and the phase-gradient based autofocus techniques discussed later in the chapter. PPP suffers the same fundamental limitations due to scene-content as the other autofocus methods but lacks the benefit of averaging sway estimates from multiple scatterers.

10.2.2   Phase curvature autofocus (PCA)

Phase curvature autofocus (PCA) is an extension of phase gradient autofocus (PGA) [Eichel and Jakowatz 1989; Jakowatz et al. 1996; Wahl et al. 1994a] for stripmap geometries. The technique was first proposed by Wahl et al. [1994b] to extend PGA to narrow-band stripmap SAR systems. As demonstrated in Hawkins [1996], PCA as originally published is suitable for narrow-band systems with no range curvature; this same work describes modifications required to extend the method to systems with significant range curvature. Hawkins [1996, pages 165–173] reported promising results for this PCA extension on simulation data where the blurring was contained in 1-D. A discussion of PCA's shortcomings with wide-band systems was also presented but simulation results were not.

PCA is intended for use in narrow-band systems where the sway error can be treated as a phase only function. This is demonstrated by the assumption that the image blurring is contained in 1-D [Wahl et al. 1994b]. PCA uses phase curvature range redundancy of the corrupted, azimuth-chirped image to estimate sway. It averages random phase curvature components from the scene to get the common phase curvature error caused by sway. The common phase curvature error is then double integrated to estimate the sway.

The PCA algorithm is described using the following steps:

1. PCA starts by finding the largest point-like-target at each across-track position x. The m^{th} target position is described by (x_m, y_m). The targets are then windowed to exclude weaker targets at the same range

$$g(x,y) = \sum_m w(y - y_m)\widetilde{f}(x,y), \qquad (10.22)$$

   where $w(y)$ is the window function chosen. Unlike PGA, the windowed targets are not circular shifted.

2. The windowed image, $g(x,y)$, is convolved with azimuth chirps appropriate for each across-track position; that is,

$$c_0(x,y) = g(x,y) \odot_y \exp\left(-j2k_0\left(\sqrt{x^2 + y^2} - x\right)\right), \qquad (10.23)$$

$$\approx \mathcal{F}_{k_y}^{-1}\left\{g(x,k_y) \cdot \exp\left(-jx\sqrt{4k_0^2 - k_y^2}\right)\right\}. \qquad (10.24)$$

   At this point, the windowed and azimuth-chirped data, $c_0(x,y)$, can be regarded as the result of the range-Doppler algorithm without the final phase multiplication[8]. Phase information previously contained in blurred points in the image $\widetilde{f}(x,y)$ has been spread out (to the width of the beam) in the along-track direction in the azimuth-chirped image $c_0(x,y)$.

3. The next step in the PCA algorithm is to multiply the windowed and azimuth-chirped data with the conjugate of the spatial chirp used for the convolution.

$$c_1(x,y) = c_0(x,y) \exp\left(j2k_0\left(\sqrt{x^2 + y^2} - x\right)\right). \qquad (10.25)$$

   Equation (10.25) removes the hyperbolic phase variation of each point target by phase matching for a target at along-track position $y_m = 0$. The process leaves a linear phase trend across each target that is dependent on its along-track position y_m [Pat 2000].

4. The phase curvature (second derivative in along-track direction) of the phase is calculated using [Pat 2000]

$$\widehat{\Delta^2\eta}(x,y) = c_1(x, y - \Delta y)c_1^*(x,y)^2 c_1(x, y + \Delta y), \qquad (10.26)$$

   where Δy is usually chosen to be one sample. The phase curvature calculation is necessary to remove both the unknown offset and linear phase trends to allow across-track averaging.

---

[8]This is again an approximation requiring a slowly varying error phase function for the reasons outlined in Section 10.3. Sutton et al. [2002] discusses the same effect in PGA based autofocus. Note the narrow-band approximation $k \approx k_0$ in the model.

5. The phase curvature for each across-track is averaged in a weighted ML fashion [Pat 2000]

$$\widehat{\Delta^2 \phi}(y) = \mathbf{Arg} \left\{ \int_x \widehat{\Delta^2 \eta}(x, y) \, \mathrm{d}x \right\}. \tag{10.27}$$

6. The average phase curvature is then integrated twice to give the phase error estimate,

$$\widehat{\phi}(y) = \int \widehat{\Delta \phi}(y') \, \mathrm{d}y'; \qquad \widehat{\phi}(0) \equiv 0, \tag{10.28}$$

with

$$\widehat{\Delta \phi}(y) = \int \widehat{\Delta^2 \phi}(y') \, \mathrm{d}y'; \qquad \widehat{\Delta \phi}(0) \equiv 0, \tag{10.29}$$

where the integration is usually calculated as a double summation (since the phase curvature was calculated using differences)

$$\widehat{\phi}[Q] = \sum_{q'=1}^{Q-1} \widehat{\Delta \phi}[q']; \qquad \widehat{\phi}[0] \equiv 0, \tag{10.30}$$

with

$$\widehat{\Delta \phi}[Q] = \sum_{q'=1}^{Q-1} \widehat{\Delta^2 \phi}[q']; \qquad \widehat{\Delta \phi}[0] \equiv 0. \tag{10.31}$$

7. The estimated error, $\widehat{\phi}(y)$, is converted into a sway estimate

$$\widehat{X}(u) \equiv \left. \frac{1}{2k_0} \widehat{\phi}(y) \right|_{y=u}. \tag{10.32}$$

and removed from the original pulse-compressed image via

$$\widehat{s}(t, u) = \widetilde{s}\left(t + \frac{2}{c} \widehat{X}(u), u \right), \tag{10.33}$$

or in the frequency domain using

$$\widehat{S}(\omega, u) = S(\omega, u) \exp\left(j2k\widehat{X}(u) \right). \tag{10.34}$$

8. The motion compensated raw data is reconstructed and the autofocus algorithm iterates from step 1. At each iteration, the width of the applied window $w(y)$ is decreased as for PGA [Wahl et al. 1994a]. The algorithm continues to iterate until the estimated position error $\widehat{X}(u)$ converges.

The principles behind PCA are identical to those of PGA. PCA operates by comparing points in the scene with point scatterers. The along-track "unfolding" of the

blurring using chirp convolution is a stripmap equivalent of the Fourier transform employed by PGA. The chirp convolution and later phase demodulation implements a coordinate transform from y to u (see the PCA blurring model above). This is similar to the chirp-Z method of performing the Fourier transform [Poularikas 1996]. Thus both PGA and PCA effectively estimate path errors in the along-track spatial frequency domain.

The operation of PCA can be improved by recognising that the along-track spreading of the target response is an approximation of the imaging process (i.e., the inverse of the reconstruction). The along-track spreading converts the blurred image into an approximation of the blurred pulse compressed data. Hawkins [1996, page 166] makes the assertion that the azimuth-chirped PCA image is the same as the range-migration corrected data (the range-migration corrected data is the same as obtained by performing the coordinate transform $\mathcal{T}\{\}$ of the range-Doppler algorithm on pulse compressed data but not performing the phase-multiply by $q(x, k_y)$). The assertion is based on the "unfolding" outlined above. However, the assertion fails if gross sway distortion occurs. Sway motion affects the phasing of the signal in the along-track direction; this in turn causes a shift in along track spatial frequency k_y. Reconstruction (which requires a well known k_y value) causes envelope shifting if sway is present [Sutton et al. 2002] and leads to the modelling errors discussed in Section 10.1.5 and [Hawkins 1996, pages 171–172]. Full inversion mitigates this problem and allows direct comparison of blurred and ideal targets in the pulse-compressed data domain.

The reference target comparison implicit in both PCA and PGA is best performed in the domain where the distortion occurs—the pulse-compressed data. The approximate wavenumber transform in PGA and the approximate inversion of PCA act to compare prominent scatterers with ideal scatterers in that domain. PCA performance can be improved by performing a full inversion of an "ideal" prominent scatterer image and comparing that with the raw pulse-compressed image. The performance improvement is from some of the approximations of traditional PCA being lifted (at the cost of increased computation). (PMA can be implemented in a similar method.)

PCA's other modifications are designed to ameliorate problems caused by the unknown along-track target position in stripmap imaging. The unknown along-track positions cause an unknown linear phase trend to modulate the phase signature of each target. This modulation prevents simple phase gradient averaging. PCA negates the effects of this linear phase trend by estimating and averaging phase curvatures. However, this results in poorer phase estimates due to the double differentiation of phase to get phase curvature leading to a reduction in overall performance. The major differences between PCA and PGA are to account for the stripmap nature of the problem.

PCA has seen little use so far in the SAS imaging field and claims have been made that it is difficult to use [Pat 2000; Sutton et al. 2000]. The double integration of

phase curvature in PCA leads to lowered phase estimate accuracy—possibly causing the observations of difficulty. Image regions with sparse target cover can cause major phase errors and are another possible cause for poor results. Better target selection for PCA, similar to that used in QPGA, or selecting more targets improves performance and reduces the problems with phase estimation accuracy and poor target coverage. The SPGA algorithm described later in this chapter can be used with a phase curvature based estimator and requires fewer approximations in its derivation.

2-D PCA

Hayes et al. [2002] describes a wide-band, wide-beam generalisation of PCA. The algorithm requires fewer approximations and has a similar computational burden to PCA. These modifications lift the requirement for the image blurring to be contained in 1-D. The algorithm appears similar to an unpublished, proprietary algorithm mentioned in the SAS tutorial at OCEANS 2001 [Dynamic Technology Incorporated 2001]. Again this algorithm can be described within the SPGA framework of Section 10.3.

10.2.3   Phase matching autofocus (PMA)

Phase matching autofocus is a modification of PGA/PCA to allow the use of additional prior knowledge [Gough et al. 2000a].

PMA is based upon a comparison between prominent image scatterers and point scatterers at the *corrected* positions of the prominent scatterers in the image. To this end, PMA is equivalent to PCA but with an additional step and phase gradient estimation.

PMA is described as follows:

1. Initially the prominent scatterers for each range in the image are found and windowed

$$g(x, y) = \sum_m w(y - y_m) \widetilde{f}(x, y). \tag{10.35}$$

2. The windowed image is chirped in the same fashion as PCA,

$$g_t(x, y) = g(x, y) \odot_y t(x, y), \tag{10.36}$$

where $t(x, y)$ is the chirping function given by

$$t(x, y) = \exp\left(-j2k_0\left(\sqrt{x^2 + y^2} - x\right)\right). \tag{10.37}$$

3. A prominent point image, $p(x, y)$, is generated. This image is generated by placing

points in the *corrected* positions of prominent points in the image,

$$p(x,y) = f(x,y)\delta(x, y - \widehat{y}_m), \tag{10.38}$$

where $\delta(x,y)$ is a 2-D delta function, the image $f(x,y)$ is used to scale the ideal points to be the same amplitude as the actual image points, and $\widehat{y}_m$ is an estimate of the corrected along-track position of the m^{th} target.

To find the correct position of a prominent point in the image, the linear phase across the target needs to be estimated[9]. This is because a linear phase-slope over the (local) aperture will shift the position of the reconstructed image. A simple linear sway error will shift but not distort[10] the target image. This effect is also well known in other fields of imaging ([van Dam and Lane 2002a]). A way of compensating for this effect, is to estimate the linear shift for each prominent target (like tip/tilt sensors in optics). Linear phase shifts in the spatial domain cause a frequency shift in the spatial Doppler domain. It is possible to estimate this linear phase by looking for the centroid of the along-track spatial frequency distribution. This technique is well known in the SAR literature as *Doppler centroid estimation* [Berizzi et al. 1997; Curlander and McDonough 1996; Madsen 1989]. Section 10.3.3 has a more detailed discussion of the techniques. In the original method, the peak amplitude of the azimuth-chirped image was used [Gough et al. 2000a]. Measuring this parameter is equivalent to measuring frequency shift due to the sinc-like beam-pattern response (albeit more susceptible to noise effects).

4. The prominent point image, $p(x,y)$, is now chirped with the same function used on the windowed image, $g(x,y)$,

$$p_t(x,y) = p(x,y) \odot_y t(x,y). \tag{10.39}$$

   The phase history (the phase behaviour in the along-track direction) between the chirped images $g_t(x,y)$ and $p_t(x,y)$ can now be matched.

5. Phase gradients with respect to y are calculated, matched and averaged. The complex gradient with respect to y for each cross-track position x is calculated via

$$\widehat{\Delta\eta}_g(x,y) = g_t(x,y)g_t^*(x, y + \Delta y), \tag{10.40}$$

$$\widehat{\Delta\eta}_p(x,y) = p_t(x,y)p_t^*(x, y + \Delta y). \tag{10.41}$$

---

[9]It is assumed that all the effects of yaw have been previously removed.

[10]Distortion occurs if echoes alias, creating unwanted side-lobes to appear [Berizzi et al. 1997; Pat 2000].

6. The phase gradient is calculated using

$$\widehat{\Delta\phi}(y) = \mathbf{Arg}\left\{\int_x \widehat{\Delta\eta_p}(x,y)\widehat{\Delta\eta_g}^*(x,y)\,\mathrm{d}x\right\}. \tag{10.42}$$

7. The phase gradient estimate is integrated to calculate the distorting phase error

$$\widehat{\phi}(y) = \int_{y'} \widehat{\Delta\phi}(y')\,\mathrm{d}y'; \qquad \widehat{\phi}(0) \equiv 0. \tag{10.43}$$

   although (as for PGA/PCA) it is usually implemented as a summation

$$\widehat{\phi}[Q] = \sum_{q'=1}^{Q-1} \widehat{\Delta\phi}[q']; \qquad \widehat{\phi}[0] \equiv 0. \tag{10.44}$$

8. The path estimate, $\widehat{X}(u)$, is used to modify the original pulse-compressed image $\widetilde{s}(t,u)$ via

$$\widehat{X}(u) \equiv \frac{1}{2k_0}\widehat{\phi}(y=u), \tag{10.45}$$

   and

$$\widehat{s}(t,u) = \widetilde{s}\left(t + \frac{2}{c}\widehat{X}(u), u\right). \tag{10.46}$$

9. The motion compensated raw data is reconstructed and the autofocus algorithm iterates from step 1. At each iteration the width of the applied window $w(y)$ is decreased as for PCA [Gough et al. 2000a]. The algorithm continues to iterate until the estimated position error, $\widehat{X}(u)$, converges.

PMA is equivalent to PCA with a phase gradient estimation kernel. Estimation of the linear slope associated with a target allows phase gradients instead of phase curvatures to be averaged. This increases accuracy at the additional computational expense of calculating the local linear phase of a target[11]. The method has an analogue in the local phase unwrapping of the WPGA phase estimator which estimates the local phase offsets to allow phase averaging rather than phase gradient averaging. PMA performs better than PCA (with accurate linear phase estimates) because it discards less information.

10.3   Stripmap phase gradient autofocus (SPGA)

The algorithm described here unifies the discussion of stripmap autofocus. Instead of deriving an algorithm based on narrow-band/narrow-beam approximations a model

---

[11]This operation assumes that minimal effects of yaw are present in the collected data—possibly an unsustainable assumption in SAS imagery.

with as few approximations as possible has been used. This allows the algorithm to be used for most single-receiver SAS autofocus tasks. Traditional autofocus algorithms may be derived from this algorithm by applying appropriate approximations. Thus, all of the algorithms previously discussed in the chapter can be described in the framework of the algorithm presented here.

The stripmap phase gradient autofocus (SPGA) algorithm consists of a number of subproblems that are described in separate sections later in the chapter. The SPGA algorithm operates as follows:

1. Starting with the blurred image $\widetilde{f}(x, y)$ a number of prominent targets in the image are selected and the coordinates (x_m, y_m) recorded. Typically a fixed number of bright scatterers are selected. Selection is discussed further in Section 10.3.1.

2. 2-D window functions are generated via

$$w_m(x, y) = \text{rect}\left(\frac{x - x_m}{W_x}\right) \text{rect}\left(\frac{y - y_m}{W_y}\right), \tag{10.47}$$

where W_x and W_y are the across-track and along-track widths respectively. These widths are chosen to be as small as possible whilst still encompassing the image blurring—see Section 10.3.2.

3. Region of interest images are formed by masking the blurred image with the window

$$\widetilde{f}_m(x, y) = \widetilde{f}(x, y) w_m(x, y). \tag{10.48}$$

4. A 2-D Fourier transform is applied to each region of interest

$$\widetilde{f}_m(k_x, k_y) = \mathcal{F}_{x,y}\left\{\widetilde{f}_m(x, y)\right\}. \tag{10.49}$$

5. The true target position, $(\widehat{x}_m, \widehat{y}_m)$ of the individual targets is estimated using the techniques discussed in Section 10.3.3. The current choice to estimate the spatial Doppler shift is via centroiding.

$$\Delta k_y = \frac{\iint k_y \left|\widetilde{f}_m(k_x, k_y)\right|^2 \, dk_y \, dk_x}{\iint \left|\widetilde{f}_m(k_x, k_y)\right|^2 \, dk_y \, dk_x}, \tag{10.50}$$

The spatial Doppler shift estimate, $\widehat{\Delta k_y}$, is mapped into an estimate of $(\widehat{y}_m - y_m)$ using the wavenumber transform ($\widehat{x}_m$ is assumed to equal x_m)[12].

---

[12]Estimation of $\widehat{x}_m$ may allow the averaging of phases instead of phase gradients, improving estimation kernel performance.

6. The region of interest images are phase modulated to correct for the shift in target position

$$\widetilde{f}'_m(k_x, k_y) = \widetilde{f}_m(k_x, k_y) \exp\left(-jk_x(\widehat{x}_m - x_m) - jk_y(\widehat{y}_m - y_m)\right). \qquad (10.51)$$

7. A coordinate transform is applied using the wavenumber transform

$$\chi_m(k_x, u) = \mathcal{SC}\left\{\widetilde{f}'_m(k_x, k_y)\right\}, \qquad (10.52)$$

where the modified wavenumber transform coordinate mapping $\mathcal{SC}\{\}$ is given by

$$k_x = k_x, \qquad (10.53)$$

$$k_y = k_x\left(\frac{\widehat{y}_m - u}{\widehat{x}_m}\right). \qquad (10.54)$$

8. Phase gradients of $\chi_m(k_x, u)$ in the along-track direction are calculated and averaged for all m. See Section 10.3.5 for a more thorough discussion on the phase-estimation problem. Currently the ML phase gradient estimator is employed

$$\widehat{\Delta\phi}(u) = \mathbf{Arg}\left\{\int_{k_x} \widehat{\Delta\eta}(k_x, u)\, \mathrm{d}k_x\right\}, \qquad (10.55)$$

where

$$\widehat{\Delta\eta}(k_x, u) = \sum_m \chi_m(k_x, u)\chi_m^*(k_x, u + \Delta u), \qquad (10.56)$$

is used to calculate the individual phase gradients.

9. A final phase estimate, $\widehat{\phi}(u)$, is generated from the phase gradients via cumulative summation or various other methods depending on the phase estimation kernel selected. When the ML phase gradient estimator is used, this is summarised by

$$\widehat{\phi}[u] = \sum_{q=1}^{u-1} \widehat{\Delta\phi}[q]; \qquad \widehat{\phi}[0] \equiv 0. \qquad (10.57)$$

10. The final phase-estimate $\widehat{\phi}(u)$ is converted into a sway estimate via

$$\widehat{X}(u) = \frac{\widehat{\phi}(u)}{2k_0}. \qquad (10.58)$$

11. The new sway estimate is used to motion compensate the image (see Chapter 6) and the algorithm iterates from step 1 until the sway estimate is less than some threshold.
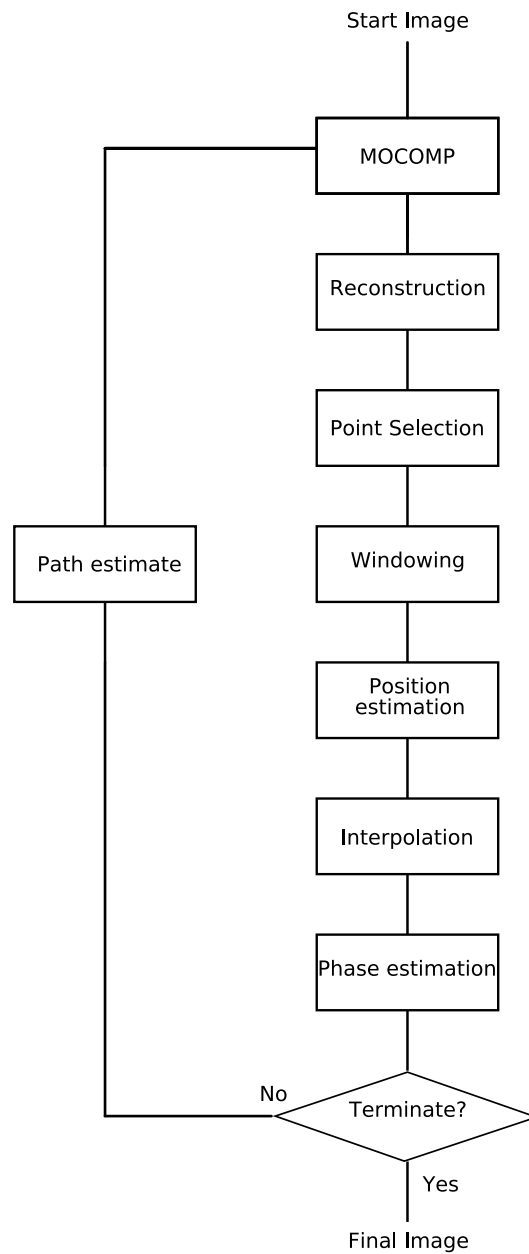
Figure 10.3   SPGA algorithm (stripmap autofocus framework). The key improvements SPGA offers are: phase gradient based estimation instead of phase curvature based estimation, and wide-beam, wide-band modelling allowing autofocus of 2-D blurring. This is possible due to the estimation of the true target position and improved blur modelling. The result is more accurate phase estimation performance.

10.3.1   Target region selection

Target region selection (selecting the point about which the region is centred) is the most straightforward aspect of SPGA: the coordinates of the target regions do not have to be estimated to sub-pixel accuracy and if poor candidate targets are selected their influence is reduced by the better phase-estimation kernels (such as WPGA's kernel). However, enough points are required to achieve the desired autofocus accuracy.

SPGA currently selects a predetermined number of target regions in order of descending energy. Regions are selected from any range bin and only need be separated by the *current* resolution at a particular iteration (estimated from the level of residual blurring expected[13]). Full aperture coverage is essential, so enough points must be selected to ensure coverage of each section of the aperture. Alternately, the selection strategy suggested for mPGA could be used, see Bonifant [1999]. Typically 100–1000 points would be selected for a 50 m by 50 m SAS image. This number should provide aperture coverage and allow accurate autofocus performance. Note that the selection of alias lobes from undersampled imagery should be avoided to prevent autofocus bias, see Section 10.5 for more detail.

Autofocus accuracy improves as more points are selected, thus having the aperture covered with multiple targets is desirable. However, a trade-off between accuracy and processing cost results, although the trade-off is not as straightforward as it appears. QPGA for example, selects 4-8 times more targets than traditional PGA but is accurate enough to autofocus without iteration [Chan and Yeo 1998]. Thus selecting more targets can *reduce* computational burden—particularly with the large iteration cost involved in stripmap reconstruction. This observation is often overlooked—autofocus accuracy improves even when using points with low signal-to-clutter ratios.

Another strategy for improving accuracy is to select widely separated points [Zavattero 1999]. This is performed to ensure that the points selected are independent i.e., are not different parts of the same object. Phase-estimation improves since the averaging of random target phase is not biased by a few strong-scattering extended targets. To some extent, the 2-D windowing used in SPGA alleviates this problem by ensuring sufficient along-track and across-track separation.

The last mechanisms for SPGA target selection improvement discussed here are those used by the QPGA algorithm:

- Select many targets discarding those with signal-to-clutter ratios lower than a given threshold.

- Using image wrapping[14] for region extraction.

---

[13]Weighted phase estimation allows an estimate of the residual blurring from the previous iteration and can be used for estimating the resolution. See the following section.

[14]QPGA's image wrapping circularly repeats the image so that regions at the edge of the image also select data from the opposite edge of the image.

The first QPGA method trades processing cost against accuracy and ensures that the points with the highest signal-to-clutter ratio are used. Weighted phase kernels (such as the WPGA kernel) perform this calculation inherently and so do not require this method—however, the QPGA method still reduces computation by reducing the number of candidate targets and associated processing. QPGA's image wrapping for region extraction mimics PGA's circular shifting operation in an efficient fashion. This aids operation when information from targets at the edges of the scene is smeared to the extent that it circularly wraps to the other edge of the scene (a property of Fourier-based reconstruction techniques). By allowing the extracted region to contain energy from both edges of the scene, the blurring is still fully encompassed, improving autofocus performance. QPGA's second improvement is used in SPGA.

10.3.2   Windowing and window width selection

SPGA windows individual targets to ensure that the phase signature of only one scatterer at a time is used. If a number of targets exist within the window, the path estimation gives poor results. As the algorithm iterates, the along-track window size decreases, reducing the likelihood of multiple targets within a single window. Selecting many targets at different locations improves sway estimation.

Window width selection is an important task for any PGA-based algorithm. The window acts to improve the signal-to-clutter ratio of the phase estimation and has a large impact on algorithm performance. In PGA, the window width also limits the order of the estimated sway. Therefore, the window width requires careful selection to avoid discarding useful autofocus information (see Section 8.6.2). This is apparent if the window does not fully encompass the target blurring; the algorithm is unable to estimate high order aperture phase components leading to residual blurring (see [Warner et al. 2000]). SPGA, like PGA, starts with a wide window that decreases over time. This window must encompass all blurring, allowing estimation of errors contained in both low-order and high-order blurring components. Usually, the initial window width is selected to exceed the expected blurring.

SPGA has slightly different windowing requirements to those of traditional PGA/PCA style algorithms due to the along-track position estimation. SPGA requires that the window width is selected to allow accurate Doppler centroid estimation. This means that a small number of along-track side-lobes of the target peak must be encompassed. Windows are typically 2-3 times bigger for SPGA than for PGA/PCA with equivalent blurring[15]. Curvature-based phase estimation removes the requirement for using larger along-track window width than PGA/PCA. Interestingly, SPGA's linear phase estimation and correction might not be required in the final iterations so the re-

---

[15]This matters more at convergence as blurred imagery contains overlapping side-lobes and mainlobe.

quirement for larger window widths at convergence is not certain. Further investigation is necessary. Another slight difference between SPGA windowing and PGA windowing is that with SPGA the window width should be range-variant since a fixed sway causes increasing blurring extent with range.

Target windowing is further complicated by along-track undersampling during collection. Even with $D/2$ sampling, alias grating lobes contain a significant amount of target energy (although the energy is smeared out). When a grating lobe is selected as an autofocus target, SPGA estimates a sway that removes the smearing[16]. This can cause incorrect overall sway estimates and thus grating lobes should be rejected from autofocus estimation windows (the problem is discussed further in Section 10.5). The predictable positioning of grating lobe targets, combined with grating lobe reduction techniques should allow grating lobe rejection. Further research is required to explore these possibilities.

The main difference between SPGA windowing and traditional PGA windowing is that SPGA uses 2-D windowing. This is due to the wide-band, wide-beam nature of SPGA's formulation. Despite that difference, window width determination is still the same: select a window width such that the (2-D) blurring is encompassed but self-clutter from other targets is reduced [Warner et al. 2000]. 2-D windowing thus requires the window width in both along-track and across-track to be determined[17]. When the residual sways are small, or the stripmap system is narrow-band, blurring is contained in 1-D. 1-D blurring allows the across-track window to be a single bin wide and SPGA windowing becomes equivalent to PGA windowing.

Traditional PGA window width selection techniques are not suitable for SAS autofocus. These typically measure the -10 dB point of the centre shifted image. These techniques work well when a number of prominent scatterers occur in the image [Wahl et al. 1994a] but fail in scenes consisting mostly of clutter. With clutter images the window width is decreased by a fixed fraction each iteration. This is needed since blurring only causes contrast loss in clutter scenes—no peak blurring is evident. SAS images have few prominent scatterers [Billon and Pinto 1995] and need to use successive window width reduction.

Typically, window widths are initially selected to encompass the maximum expected blurring and are decreased by 20%–50% each iteration. Window width selection based on the estimation of residual blurring from the previous iteration may be possible with the WPGA estimator[18] and is an area where additional research is required.

---

[16]SPGA combines weighted estimates which alleviates the problem to some extent.

[17]The problem is simpler than implied since along-track and across-track blurring are coupled; an estimate of only the along-track blurring can determine the expected across-track window width.

[18]It should be possible with any estimator in conjunction with the CRLB calculations of Chapter 7. WPGA is preferred as complicated reflector structure is considered noise.

Figure 10.4 The effect of a linear phase shift on the Doppler spectrum of a target. Note that the spectrum depicted is oversampled with respect to the SAS sampling constraints ($\approx D/8$ sampling). (a) No linear sway. (b) Linear sway of γx. More complex sway motions change both the position and shape of the Doppler spectrum.

10.3.3 Along-track position estimation

The biggest difference between SPGA and other PGA/PCA based algorithms is the estimation of the along-track position of the targets before phase-estimation. This allows phase gradient averaging instead of phase curvature averaging and improves phase-estimation. Target position estimation is necessary because phase gradient estimation is unable to estimate the linear phase trend caused by target shifting [Jakowatz et al. 1996; Rachlin 1990; Wahl et al. 1994b]. Unknown linear phase trends prevent averaging of phase gradients and must be removed. In contrast, spotlight imagery has an unknown linear phase trend common to all targets so phase gradient averaging is possible.

The along-track target position is estimated using prior knowledge of the beam-pattern and spatial-frequency coverage of stripmap systems. As shown earlier in this chapter, linear phase shifts cause frequency shifting. Thus the linear phase trend across the target may be estimated by measuring the shifting of the target's along-track spatial Doppler spectrum.

Figure 10.4 shows the effect a linear sway has on the along-track spatial Doppler spectrum. Note that the shifting is predictable and causes spatial Doppler aliasing. Where severe linear trends exist the aliasing causes ambiguity.

Three methods for determining the Doppler spectrum shift have been employed with SPGA to date. These are:

Doppler centroiding — Doppler centroiding operates by estimating the centroid of

the power spectrum averaged over range (or k_x). This may be described by

$$\Delta k_y = \frac{\iint k_y \left| \widetilde{f}_m(k_x, k_y) \right|^2 \, \mathrm{d}k_y \, \mathrm{d}k_x}{\iint \left| \widetilde{f}_m(k_x, k_y) \right|^2 \, \mathrm{d}k_y \, \mathrm{d}k_x}, \tag{10.59}$$

where $\widetilde{f}_m(k_x, k_y)$ is the wavenumber-domain data of the target patch. The centroid of the power spectrum is preferred to the centroid of the amplitude spectrum since linear trends cause energy-shifting and the amplitude can be adversely affected by higher order sway. Note that the centroid estimate also causes an underestimation of the shift due to the circular repetition of the Doppler spectrum caused by Doppler aliasing. (Spatial frequencies are aliased in a circular fashion, as can be seen in Figure 10.4.) This aliasing can be resolved using a circular centroid estimator of the form

$$\Delta k_y = \frac{k_{y\max}}{\pi} \mathbf{Arg} \left\{ \iint \left| \widetilde{f}_m(k_x, k_y) \right|^2 \, \mathrm{d}k_x \exp \left(j \frac{2\pi k_y}{k_{y\max}} - k_{y\min} \right) \mathrm{d}k_y \right\}. \tag{10.60}$$

Modified Doppler centroiding – Using the observation that the centre of the target scene is not necessarily the centre of energy[19] and that Doppler centroid estimation aims to estimate the linear trend via energy-shifting, a modified centroid estimation can also be used. This uses the centroid of the image scene and correcting the normal Doppler centroid result. Centroiding of the target scene improves the estimation when multiple targets are selected in a region of interest—the limiting condition on autofocus performance. Similar observations might also aid the correlation estimator discussed next.

Doppler correlation — Doppler correlation is another method of estimating the Doppler shift and does not suffer underestimation due to Doppler aliasing. Doppler correlation operates by correlating the (averaged) spectrum of the target patch with the expected spectrum. This operation is described by

$$\Delta k_y = \max_{k_y} \left\{ \int \widetilde{f}_m(k_x, k_y) \, \mathrm{d}k_x \star_{k_y} A(k_y) \right\}, \tag{10.61}$$

where

$$A(k_y) = \left(D_{\mathrm{t}} \, \mathrm{sinc} \left(\frac{k_y D_{\mathrm{t}}}{2\pi} \right) D_{\mathrm{r}} \, \mathrm{sinc} \left(\frac{k_y D_{\mathrm{r}}}{2\pi} \right) \right), \tag{10.62}$$

is the expected amplitude spectrum.

Of the methods, modified Doppler centroid estimation is preferred because of its improved accuracy and reduced computational requirements.

---

[19]Often the case where multiple targets are selected inside a given window.

It is worth noting that the noise performance (accuracy) of position estimation is different to that of the phase estimation kernels described earlier. Clutter suffers the same linear shift as the selected target so has the same spatial Doppler spectrum. The result is that the performance of the estimator is determined by the signal-to-noise ratio (SNR) not the signal-to-clutter ratio (β). The difference can be as much as 20-40 dB.

A strong warning regarding tow-fish yaw must be raised at this point. Tow-fish yaw also shifts the along-track Doppler spectrum[20] although does not cause a shift of target position in the image (see Chapter 6). The phenomenon is noted in some SAR Doppler centroid estimation research [Prati and Rocca 1992]. Thus the spectrum shifts associated with tow-fish yaw must be demodulated and corrected so that only sway-induced spectrum shifts are present. While the KiwiSAS-II and KiwiSAS-III platforms have been designed to keep yaw to a minimum [Hayes and Gough 1992; Johnson et al. 1995], yaw effects are still apparent in KiwiSAS-II imagery (see Section 9.3.2). This is generally apparent as position dependent grating lobe structure and Doppler centroid offsets. Yaw should be estimated (using the methods described in this chapter) and removed via the techniques outlined in Chapter 6. This is important for constant yaw in particular.

In the presence of large (or rapid) unmeasured tow-fish yaw the SPGA algorithm should calculate and average phase curvatures instead of phase gradients (cf. with Section 10.2.2). Alternately, the local linear phase trend across each target can be estimated from the phase curvature based motion estimate (once the global linear trend is estimated using the techniques described above).

Estimating the true position of the targets allows the use of phase gradient estimation. This is possible when platform yaw is negligible or accurately measured[21]. When unknown yaw is present, workarounds allow the use of phase gradient estimation. In addition, phase curvature based estimation can be used with a decrease in performance. Finally, phase gradient based estimation is sensitive to the accuracy of the along-track target position estimates, any technique for improving those estimates will benefit the SPGA algorithm.

10.3.4   Wavenumber transform coordinate change

The modified wavenumber transform coordinate transform $\mathcal{SC}\{\}$ acts to transform the phase error estimates from the individual targets into a space-invariant domain prior to averaging. This is a typical step in space-variant image processing [Sawchuk 1972]. The success of the wavenumber transform in accomplishing this results from its similarity to the reconstruction process. Both PGA and PCA use approximations of

---

[20]This is the case in single receiver SAS. The effect of yaw on multiple-receiver imagery is much more complex although the same principles apply.

[21]Yaw estimates provided by an INS tend to have a high degree of accuracy [Bellettini and Pinto 2002].

the wavenumber transform tailored to the environment they are designed for: PGA, a Fourier transform; and PCA an along-track chirp spreading[22].

The wavenumber transform maps image blurring to the phase error in the pulse compressed data that caused the blurring[23]. Once in the pulse compressed data domain, the phase errors have redundancy in along-track and may be averaged. Any error in estimating the along-track position of the target adversely affects the average. However, the difference between the image target position and the estimated true target position is small for later iterations and has little effect on the average. The use of the wavenumber transform allows phase error information from multiple targets to be combined in sensible fashion.

10.3.5   Phase estimation

Phase estimation is the heart of the SPGA method. Like PGA, small improvements to the phase estimation kernels can yield large improvements in the end result. The phase estimation kernel used depends on the accuracy of the along-track position estimation. When the along-track position estimates are accurate, a phase gradient method should be used. However, when the along-track position of the targets is unknown or inaccurate, a phase-curvature kernel must be used. Even with unknown along-track positions, a phase gradient kernel can be used once an estimate of the phase error is obtained using phase curvatures (see Section 10.3.3).

The kernel currently used for SPGA is the ML phase gradient estimator. Other possible phase estimators are summarised in Section 7.6.2. WPGA provides a better phase estimator than the ML estimator suggested. However, the weightings used in WPGA's derivation are calculated based on spotlight SAR assumptions and the iterative scheme employed cannot currently be used in a stripmap framework.

10.3.6   Improved blur modelling—sidestepping the limitations of the wavenumber transform

When significant or rapidly varying sway motions exist, the wavenumber transform does not accurately map between wavenumber domain and pulse-compressed data. This is discussed in Section 10.1.1. The largest effect is that unknown linear sways cause target shifting. The algorithm convergence rate is adversely affected since the estimated blurring function does not map accurately into a sway motion. Given the estimated blurring function, enhanced modelling of the wavenumber transform can restore that accuracy.

---

[22]PCA uses the chirp-scaling principle to perform an equivalent coordinate transformation to the wavenumber transform.

[23]A variation of the wavenumber transform is used (not performing the mapping from k_x to ω to allow wide-beam phase averaging).

It is important to remember that the initial wavenumber transform maps to an incorrect along-track position u_{sway}. The steps to calculating the error are as follows:

1. Given the position gradient $\dot{X}(u_{\text{sway}})$ (calculated via the phase estimation above) measured as a function of the distorted along-track position u_{sway}, the associated frequency shift Δk_y may be calculated via

$$\Delta k_y(u_{\text{sway}}, k_x) = k_x \dot{X}(u_{\text{sway}}), \tag{10.63}$$

where

$$\dot{X}(u_{\text{sway}}) = \frac{\mathrm{d}X(u_{\text{sway}})}{\mathrm{d}u_{\text{sway}}}, \tag{10.64}$$

$$= \Delta_{u_{\text{sway}}}\big(X(u_{\text{sway}}) - X(u_{\text{sway}} - \Delta_{u_{\text{sway}}})\big). \tag{10.65}$$

is the ping-to-ping sway difference. Note that this has been derived by calculating the instantaneous frequency shift via the instantaneous phase derivative.

The true instantaneous spatial Doppler frequency, $k_y(u_{\text{sway}})$ (as opposed to the incorrect estimate using the wavenumber transform) may now be derived via

$$k_y(u_{\text{sway}}) \approx k_{y_{\text{dist}}}(u_{\text{sway}}) - \Delta k_y(u_{\text{sway}}, k_x), \tag{10.66}$$

$$\approx \frac{k_x}{x_m}(y_m - u_{\text{sway}}) - k_x \dot{X}(u_{\text{sway}}), \tag{10.67}$$

which represents an extension of the wavenumber transform for larger sways. The approximations made in the derivation represent a linearisation of the problem about the point in question—more accurate representations are not necessary in an iterative autofocus framework. The apparent target shift along the aperture from the true along-track position u can be now calculated via the wavenumber transform

$$\widehat{u}(u_{\text{sway}}) = y_m - \frac{k_y(u_{\text{sway}})x_m}{k_x}, \tag{10.68}$$

Note that (10.68) can be used to derive the apparent target shift induced by a linear sway seen earlier (10.13).

2. The true along-track position, u, is derived from the apparent position, u_{sway}, using

$$\widehat{u}(u_{\text{sway}}) \approx u_{\text{sway}} + \dot{X}(u_{\text{sway}})x_m. \tag{10.69}$$

3. The data is then re-interpolated from $\chi(k_x, u_{\text{sway}})$ and $\Delta\phi(u_{\text{sway}})$ to obtain $X(u)$ using the calculation of $\widehat{u}(u_{\text{sway}})$ above in a modified wavenumber transform.

The result in (10.69) corresponds to the (geometrical optics based) derivations seen in [Callow et al. 2002a] and [van Dam and Lane 2002b]. The procedure outlined above can be iterated for additional accuracy although this is not required for SAS autofocus. Note that SPGA as described in Section 10.3 implements first-order (linear slope) correction of this effect (SPGA, step 6).

10.3.7   SPGA versus traditional algorithms

SPGA provides a framework which can describe a number of traditional autofocus algorithms.

PGA

SPGA implements PGA if narrow-band and narrow-beam approximations are made and the along-track position of the targets is assumed to be zero ($y_m \approx 0$) in the phase compensation and interpolation (steps 6 & 7) . Under these approximations, the wavenumber transform becomes the along-track Fourier transform that PGA uses. The centre shifting step of PGA is implicitly implemented by SPGA's centring the Fourier transform around the target region of interest. Further work is required on the phase estimation techniques of the more sophisticated PGA variants such as WPGA.

mPGA

SPGA is able to emulate mPGA when using mPGA's target selection strategy and narrow-band, narrow-beam approximations. mPGA implements phase compensation and interpolation the same way as PGA although the along-track position estimate, y_m, for each target is chosen differently. mPGA estimates y_m by setting it equal to the along-track coordinates of each sub-aperture. mPGA thus has more approximations and poorer phase estimation performance than SPGA. The target selection technique mPGA provides can be used with SPGA if desired.

PCA

PCA is implemented by using narrow-band, narrow-beam approximations and using a phase curvature based phase estimation.

2-D PCA

Implemented by using phase curvature based estimation.

PPP

PPP is implemented by allowing SPGA only a single point and using a phase-averaging estimation kernel.

10.3.8   Future work

In the embodiment described here, SPGA has a number of possible limitations. The most severe limitations are caused by a requirement for a statistically homogeneous scene. Further work is required in each of the following areas:

Phase estimators — WPGA's phase estimator is the most accurate of those investigated but relies on spotlight SAR assumptions. Further work is required to use this estimator instead of the ML estimator currently used with SPGA.

Clutter information — The information contained in the sea-floor clutter is useful for autofocus is the data is oversampled in along-track. SPGA currently discards this source of information. However, we expect that SPGA is able to operate on clutter regions with a different patch selection policy. Clutter regions (on average) satisfy the requirement for a statistically homogeneous scene and so are able to be used directly. The signal-to-noise for autofocus is poor so many such scenes are required.

Non-Lambertian scatterers — SPGA assumes that the energy scattered from a target patch is the same from all incidence angles. SPGA relies on beam symmetry for the linear slope estimation so targets that have non-Lambertian scattering profiles can cause incorrect results.

A further assumption that the apparent position of the target should not move with varying incidence angle. Most current autofocus algorithms require apparent target position to remain constant with incidence angle. Targets where that is not the case degrade autofocus performance. For example, the shadowing behind proud targets moves and could cause such degradation; similar problems are likely to be caused by specular reflections on cylindrical targets.

Squinted systems — SPGA will not currently operate on systems that image with a squinted geometry (see above). However, if the squint angle is known, this can be compensated.

Bulk sway estimation — Approximations made in SPGA's derivation require that the sway is small so there is no shift in the stationary phase point in the derivation of the wavenumber transform. The method outlined in Section 10.3.6 lifts these resrictions at the expense of additional computation. A micronavigation technique may prove more efficient for estimating the bulk of the sway.

Platform yaw — A reliance on the spatial Doppler centroid makes the algorithm vulnerable to platform yaws. A yaw will shift the spatial Doppler centroid leading to inaccurate linear slope estimation. SPGA assumes that no yaw is present in the system. Yaw currently represents the major limitation of SGPA. Further work to remove the reliance on yaw-free data is required.

Multiple-receiver autofocus — The SPGA method described here does not imme-
diately lend itself to multiple-receiver autofocus. Currently, like the PGA variant
used by Sutton et al. [2002], a single-receiver blurring model is used for autofo-
cus and multiple-receiver information inferred from the results. Multiple-receiver
autofocus using single-receiver assumptions has potential problems. In partic-
ular, collection redundancy can lead to multiple sways for a given along-track
sampling position. A multiple-receiver blurring model derived along the lines
of the wavenumber transform and the derivations of Chapter 3 would improve
autofocus.

Path prior inclusion — Autofocus procedures benefit if statistical properties of the
motion are known [Fortune et al. 2001b]. The inclusion of this prior information
can prevent unlikely motions from being estimated and can prevent the autofo-
cus result from diverging. The improvements come at the expense of increased
computation in the phase estimation.

Path prior inclusion for SPGA can be made by altering the phase estimation step.
Instead of using a weighted phase estimation kernel, a weighted-with-prior kernel
should be used [Franklin et al. 1998]. This should be investigated at the same
time the WPGA phase kernel is implemented for the SPGA algorithm.

3-D geometry — The derivation of SPGA presented does not include the effects of
slant-plane geometry—it is assumed that everything in the image is in the ground
plane of the sonar. These assumptions restrict SPGA's use slightly as accuracy of
the wavenumber transform decreases in slant-plane situations. Further research
is required to derive the wavenumber transform for slant-plane geometries.

10.4   Range offset

Stripmap imagery is adversely affected by unknown range offsets or unknown sonar
velocities. Either of these effects cause a quadratic defocus in the image. Blurring from
these sources is space-invariant (all of the targets have a similar defocus). Current
space-variant autofocus procedures do not effectively deal with this type of blurring
and ambiguous results occur.

The failure of SPGA and the traditional prominent point/phase gradient methods
stems from a problem with PPP methods. Blurring due to reconstruction errors is
assumed to be the same as that caused by aperture phase errors—quadratic blurring
from reconstruction deficiencies is assumed to come from sway.

The blurring from reconstructing with an incorrect range offset (Δx_0) can be mod-
elled as

$$\widetilde{S}(\omega, k_u) = S(\omega, k_u) \exp\left(j\Delta x_0 \sqrt{4k^2 - k_u^2}\right)\Big|_{k_u = k_y}. \tag{10.70}$$

(a)   (b)

Figure 10.5 Illustration of the range offset problem. Targets simulated at x_0=50 m and reconstructed with x_0=55 m (equivalent to having a +5 m constant sway error). (a) Image reconstructed with incorrect x_0 factor. Note blurring spreads the along-track response slightly. (b) SPGA sway estimate generated from the blurred image. The quadratic blurring each target suffers is mistakenly assumed to be caused by sway. The result is poor autofocus performance. Space-invariant blurring caused by poor reconstruction must be eliminated prior to space-variant autofocus.

A similar model may be derived for a velocity error using the binomial approximation and the Fourier scaling theorem

$$\widetilde{S}(\omega, k_u) \approx S(\omega, k_u) \exp\left(jx_0 \frac{(\alpha - 1)k_u^2}{4k^2} \right)\Bigg|_{k_u=k_y} . \tag{10.71}$$

where the reconstruction velocity is v_s and the actual velocity v_s/α.) Note that the blurring from both models is invariant of target position.

With space-invariant blurring, targets at different positions have the same blurring. Where the problem lies is that a PPP-based algorithm will map the blurring for each target into a sway estimate. This occurs even when there is no sway. A stripmap algorithm uses many targets and estimates the *same* (quadratic) sway for each distinct target. When many incorrect sways are averaged, the overall sway estimate is incorrect. This occurs because PPP (and phase gradient) algorithms operate by sharpening each defocused point in the image[24].

Figure 10.5(a) shows an image focused with an incorrect x_0 parameter (in this case $x_{0_{\text{recon}}} = 55$ m when $x_0 = 50$ m, so $\Delta x_0 = 5$ m). There is a slight along-track blurring of the targets in the image due to the incorrect reconstruction parameters. Note that reconstructing with an inaccurate velocity estimate has a similar effect. Figure 10.5(b) illustrates the effect poorly reconstructed imagery has on autofocus performance. Quadratic sways are estimated local to each target; when combined, very poor sway estimation results.

The problem may be mitigated by using a space-invariant autofocus technique prior

---

[24]Broadly speaking this is the case, they compensate for phase fluctuations in the spatial frequency domain which has the general effect of sharpening.

to stripmap autofocus—such as any of the spotlight autofocus methods, or an altered SPGA (neglecting the y_m term). Kirk and Maloney [1998] describe a space-invariant algorithm for this explicit purpose. However, even after using these methods, ambiguity exists and poor estimates of velocity and focus range can be obtained when sway errors are present [Fortune et al. 2001b]. It is better to have accurate estimates of important reconstruction parameters using micronavigation or other techniques.

10.5   Aperture undersampling

Aperture undersampling can degrade autofocus performance. In general, performance is decreased when alias grating lobes can be seen above clutter. As the undersampling becomes more severe, autofocus performance degrades. Note that echo-correlation algorithms (see Chapter 9) do not operate at all with sample spacings $> D/2$.

The reason that autofocus is affected by along-track undersampling is that aliased grating lobes (alias targets) are selected as prominent scatterers. Autofocus algorithms estimate a path to remove the Fourier phase of these scatterers—i.e., to sharpen them. Alias targets cause incorrect sway estimates to be returned since they are smeared during reconstruction and autofocus mistakenly attempts to remove the smearing.

Alias targets are always present in SAS imagery due to the unusual sampling requirements discussed in Chapter 2. The relative level of the alias targets varies with along-track sample spacing. One measure of this level is known as the along-track ambiguity to signal ratio (AASR). The AASR is -8 dB for $D/2$ [Hawkins 1996, page 125], [Gough and Hawkins 1997]. For a $D/2$ sampled system, each of the two main grating lobes is -14 dB relative to the main-lobe response—i.e., the energy contained in each alias target is 14 dB less than the main-lobe[25].

Figure 10.6 illustrates the spatial frequency situation when $D/2$ sampling is used. Note that the responses shown are the main-lobe response of the combined transmit receive aperture and its first alias. The AASR is calculated by considering the energy of the main-lobe and the alias within the processed bandwidth—from $-2\pi/D$ to $2\pi/D$ in this case.

The difficulty for stripmap autofocus is that alias targets appearing in the image above the surrounding clutter can be selected as prominent scatterers. To demonstrate this effect, a hypothetical autofocus problem with a single $D/2$ sampled target was simulated[26]. No sway was injected and any sway estimate other than zero is due to undersampling-related errors. Figure 10.7(a) shows the reconstructed scene—note the obvious alias targets at ± 10 m. (In a normal autofocus situation the alias targets will be partly or wholly obscured by clutter.) For comparison with Figure 10.6, the spatial

---

[25]The level of the alias target in the image is further reduced due to the alias smearing that occurs during reconstruction.

[26]Simulated system corresponds to the KiwiSAS-II with x_0=50 m.

Figure 10.6   Aliasing in an undersampled system.   Main-lobe and side-lobe response for a $D/2$ sampled system.   The reconstructed bandwidth (without filtering) extends from $-2\pi/D$ to $2\pi/D$. Note only one side-lobe is shown. The along-track ambiguity to signal ratio (AASR) for the illustrated side-lobe is -14 dB.

frequency response of the region around the main-lobe and the alias target at $+10\,\mathrm{m}$ have been illustrated in Figure 10.7(b).   This confirms the AASR at approximately -14 dB for the alias target at $+10\,\mathrm{m}$.

SPGA was used to autofocus the image in Figure 10.7(a).   The algorithm selected the main-lobe and each of the alias targets (in the absence of any other targets). The ping-to-ping sway differences (sway-gradient) estimated by SPGA are shown in Figure 10.8(a). Note that SPGA estimates both an incorrect linear sway and curvature in the presence of alias targets[27]. For comparison, the sway-curvature (estimated using the PCA estimator to prevent phase unwrapping artefacts) is shown in Figure 10.8(b). This shows that an incorrect sway-curvature is also estimated over the region of the alias targets. Thus alias targets cause errors in the SPGA algorithm. These difficulties are more significant when using the phase gradient estimator rather than the phase curvature estimator.

In summary, aperture undersampling causes some degradation in autofocus performance.   The overall performance degradation is minor but can lead to unexpected results in scenes with a large dynamic range or with strong-scattering isolated targets. Enhanced windowing (windowing both the main-lobe and alias) reduces the the extent of the problem as does the use of phase curvature estimation kernels.   Moreover, alias

---

[27]Alias targets have a large apparent linear sway and smaller components of high-order sway caused by reconstruction smearing.

(a)

(b)

Figure 10.7   Simulated system results for $D/2$ sampling. (a) Undersampled image shown in linear grey-scale and clipped to 20 dB below peak. Note grating lobes at ± 10 m and range side-lobes. (b) Measured spatial frequency coverage of main-lobe and grating lobe at $y_n = +10$ m.

(a)

(b)

Figure 10.8   Autofocus difficulty with undersampled imagery—worst case scenario. Single simulated scatterer in centre of scene (0 m) with along-track sampling of D/2. (a) SPGA phase gradient estimate from undersampled target. Note acceptable performance in the middle region of the path, from -5 m to 5 m. The performance is degraded in the grating lobes of the target. Both the gradient and curvature are incorrect. Also note that phase unwrapping has not been employed resulting in λ jumps at $\pm\,8$ m. SPGA performance is severely degraded because both linear slope and higher order sways are estimated incorrectly. (b) Phase curvature estimate of undersampled target using PCA's phase curvature estimator within the SPGA framework. Note that undersampling still causes incorrect curvature estimates in the grating lobe response (although overall sway estimation is not as bad as with SPGA which also predicts linear sways incorrectly).

Parameter	Value	Units
Carrier freq	30	kHz
Bandwidth	20	kHz
Number receivers	1	
Receiver length	0.3	m
Transmitter length	0.3	m

Table 10.1   Simulated system parameters approximately corresponding to KiwiSAS-II [Hayes et al. 2001].

targets can be detected via their unusual spatial Doppler spectrum and rejected if necessary. Selecting additional widely separated targets for autofocus operation further reduces the problem.

10.6   Autofocus results

The following sections detail results of testing the SPGA algorithm on simulated and field collected data.

10.6.1   Simulated data

This section presents results of the SPGA technique on a simulated data set. The data sets were simulated based on a system roughly equivalent to the KiwiSAS-II. Parameters for the system are summarised in Table 10.1. For most tests, the along-track sample spacing was chosen to be $D/4$ and has an AASR of -21 dB (the undersampled system uses $D/2$ sampling and has an AASR of -8 dB).

A data set consisting of a number of point scattering targets was simulated (in the ground-plane) and the sway depicted in Figure 10.9(b) was inserted into the data as a *timing-error* resulting in Figure 10.10(a). The data set used is depicted in Figure 10.9(a). Note that the timing-error assumption is valid for the KiwiSAS-II sonar. The sway was randomly generated and has a correlation-length of approximately 1 m (10-15 sonar pulses) and a peak-to-peak amplitude of 0.2 m. These values were chosen to approximate the distortions expected from the KiwiSAS-II sonar. Note the 2-D blurring apparent in the corrupted image shown in Figure 10.10(b). The blurring is different for each scatterer and in the worst case extends approximately 3 m in along-track and 0.5 m in across-track. Autofocus results for comparative analysis were then generated using the SPGA algorithm to remove the scene blurring. Quantitative convergence measures were obtained by measuring errors in estimating the known sway path.

The convergence rates of a number of SPGA variations are illustrated in the following figures. The RMS sway *gradient* error is plotted for both gradient and curvature

(a)

(b)

Figure 10.9   Simulated data set used in SPGA testing. (a) Ideal uncorrupted image. (b) Corrupting sway.

(a)

(b)

Figure 10.10 Simulated data set used in SPGA testing. (a) Pulse compressed data corrupted using the sway of Figure 10.9(b). (b) Blurred image reconstructed from (a).

estimation kernels versus iterations[28]. The RMS gradient error is a more direct mea-
sure of image blurring than the RMS sway error used in Section 7.3. In addition, the
mean gradient error was removed from the measure to make it comparable between
the curvature and gradient estimators. The initial window size was selected as 1 m in
across-track and 9 m in along-track based on the known blurring extent of the image.
The window was reduced by 40% per iteration (factor = 0.6). Note that allowance
was made for the additional initial box width required for SPGA's along-track target
position estimate. When using a curvature estimator, the initial along-track window
width was chosen to be 1/2 that of the gradient estimator since along-track position
estimation is not needed. Only 25 regions of the image are taken due to its small
extent—this is approximately two regions for each target. Unless otherwise specified,
the Doppler correlation method has been used for along-track position estimation.

Figures 10.11 and 10.12 illustrate SPGA's convergence rates using phase gradient
and phase curvature estimators for various window reduction rates. The window size
of the i^{th} iteration is given by $W_{y_i} = \alpha W_{y_{i-1}}$ for various α between 0.3 and 0.8. Note
that in a scene such as this it is better to base the window width reduction on *measured*
autofocus performance. (Windowing was chosen this way to mimic that used in clutter-
scene autofocus of SAS scenes.) SPGA with a gradient estimator provides a better
first iteration and generally improved convergence performance when compared with
the curvature estimator. The downside is divergence when using the phase gradient
estimator when α is low (< 0.4). When the window size is rapidly decreased, the
along-track position estimation fails leading to algorithm divergence. The effect may
be avoided by ensuring adequate along-track window size.

The apparent initial divergence when using the curvature estimator is a common
problem in early autofocus iterations. Whilst the accuracy of estimating the sway is
poor, the image is sharpened. The curvature estimator effectively autofocuses the image
as two disjoint scenes in the first iteration. Each scene has an incorrect phase gradient
(local linear sway) and although individual targets are well focused (suffering only from
a linear sway) the disjoint autofocus result causes a large error in the performance
metric chosen. Later iterations are able to use the improved separation of the two
targets in the centre of the image to avoid further divergence[29].

For the majority of window reduction rates, the phase gradient estimator offers
faster convergence than the curvature estimator. However, care must be taken to en-
sure the along-track window-width is sufficiently large to prevent autofocus divergence.
The phase gradient estimator should be used in preference to phase curvature based
estimation and should include a method of estimating the necessary window-width (for

---

[28]Only the error over the region where target coverage exists is considered—approximately -15 m to
+15 m along-track.

[29]The centre targets provide the only source of information for the linear sway over the disjoint
scenes. Any autofocus error (residual blurring) in these targets results in incorrect linear trends across
each scene. This is due to the double integration required when using a phase curvature estimator.

Figure 10.11   Convergence rate versus window size reduction rate for SPGA. Window size reduced using $W_{y_i} = \alpha W_{y_{i-1}}$ for various α between 0.3 and 0.8. Note divergence for $\alpha = 0.3$.

preventing divergence).

Figure 10.13 illustrates the differences in SPGA convergence rate when using when using the Doppler correlation and the two Doppler centroiding methods of along-track position estimation when combined with the phase gradient estimator. Note that in this autofocus scenario the Doppler aliasing effects noted in Section 10.3.3 are not significant (which assists the Doppler centroiding method). The overall convergence rate of the correlation and centroid estimators is similar, with the centroid estimation showing better operation in early iterations but slight divergence in later iterations. The modified centroid estimator shows all-round improvements over the other methods of along-track position estimation; it is the best estimator for this scene.

The reason for the improved performance of the modified centroid estimator is due to better estimation of the energy shift occurring in the image. Accuracy in estimating energy shift corresponds directly to accuracy in linear phase (linear sway) estimation. Thus the ability of the modified centroid estimator to account for off-centre energy distribution when multiple targets exist inside the target window improves autofocus accuracy[30].

The relatively poor performance of the Doppler correlation-based methods can probably also be attributed to off-centred energy in the initial target region. In the case of two targets within the target window region, the Doppler spectrum has two

---

[30]Position estimation using standard Doppler centroiding assumes that the scene energy is centred in the target region—an assumption that is not valid in general.

Figure 10.12 Convergence rate versus window size reduction rate for PCA (SPGA with curvature estimator). Window size reduced using $W_{yi} = \alpha W_{yi-1}$ for various α between 0.3 and 0.8. Note large first iteration error. This is caused by the scene being autofocused in two disjoint sections—each with an unknown linear sway (sway-gradient). This leads to a large error in the global sway estimate but accurate estimates for each local section.

distinct peaks (due to superposition)—selecting the maximum point in the correlation chooses one peak or the other rather than the (more accurate in this situation) average of the two. This deficiency of model and the effect of noise on a broad correlation peak make accurate linear phase prediction (and thus accurate target position estimation) difficult to achieve using Doppler correlation.

Based on the results shown in Figure 10.13 position estimation via modified Doppler centroiding is recommended. Where Doppler spectrum aliasing caused by large sways is evident the circular centroiding technique should be employed.

Figure 10.14 shows SPGA convergence properties when operating on $D/2$ sampled data. This sampling rate represents a slight undersampling of the data. Again, the gradient estimator gives improved performance over the curvature estimator which suffers poor apparent performance in the early iterations (see earlier discussion). Note that the target extent is altered and so the initial RMS sway error values are altered in comparison with the other tests presented here.

It is important to note in Figure 10.14 that SPGA still provides autofocus improvement even with $D/2$ sampled data—micronavigation based autofocus techniques are unable to provide any improvement in this situation. Note however that autofocus performance is degraded slightly compared with the results obtained early in the section with $D/4$ sampled data.

Figure 10.13 SPGA convergence using gradient based estimator when using the Doppler centroiding and Doppler correlation along-track position estimators discussed in Section 10.3.3. Note faster convergence in the initial iterations using the centroid based position estimators and overall better convergence with the modified centroid estimator.

Figure 10.14 SPGA convergence on $D/2$ (slightly undersampled) data for gradient and curvature based estimators.

Figure 10.15 Image after SPGA autofocus using gradient estimator.

Figure 10.15 illustrates the autofocus result for the scene using the optimum parameters found in this section. In this case, the result is from nine iterations of SPGA using the gradient estimator with modified Doppler centroiding ($\alpha = 0.6$). The autofocus quality, like the blurring itself, is space-variant with diffraction limited imagery at $u = +10$ m and minor residual blurring at $u = 0$ m. The minor residual blurring is due to the self-clutter caused by neighbouring targets (causing self-noise in the phase estimation). In this case, the adjacent targets around (53,-10) and (50,1) cause small autofocus errors.

Figure 10.16 demonstrates the small autofocus errors at convergence by depicting the injected and estimated sways after nine iterations of SPGA using the gradient estimator. The estimated sway corresponds to that used when reconstructing the image in Figure 10.15.

A constant sway error exists which causes slight autofocus (and image) errors. However, as the constant sway offset error has little effect on the reconstructed image and is difficult to estimate—see Section 10.4—its estimation has been avoided in the SPGA implementation. Other salient points to note in Figure 10.16 are the generally good sway estimation, and slightly larger errors at the edges of the scene where there is no target coverage.

It is possible to see more detail and glean more information from Figure 10.17 which depicts the injected and estimated sway-gradients (obtained by taking first-order differences of the sways in Figure 10.16). (The results from using the curvature

Figure 10.16 Sway estimated using SPGA and a phase gradient estimation kernel. The sway offset has only a minor effect and is not estimated by the algorithm.

estimator with other parameters held is also shown.) As shown in Section 7.3 sway-gradient errors correspond directly to image blurring; this makes the depiction of sway-gradients rather than sway better for interpretation purposes.

Small estimation errors are apparent in Figure 10.17 around u=-10 m and u=0 m (both gradient and sway based estimation). Those errors are caused by nearby targets at u=-10 m and u=0 m causing self-clutter—they are in effect a source of autofocus noise. Other sources of errors occur at the edges of the scene where there is no aperture coverage or target information. Note that the errors visible in Figure 10.17 are not large enough to visibly blur Figure 10.15.

From the injected and estimated sway-gradients shown in Figure 10.17 it is possible to see that gradient estimation performs better than curvature estimation. As noted earlier in this chapter, curvature estimation is unable to estimate the sway-gradient offset (effectively caused by an unknown integration constant). This results in the constant offset visible in the curvature estimated sway-gradient in Figure 10.17. In addition to offset error, the curvature result shows larger errors than the gradient result around the u=-10 m region.

Summary

The testing presented in this section has been performed in the absence of clutter and receiver-noise. These noise sources can limit autofocus performance so only conclusions about an algorithm's structure and approximations can be reliably be drawn. However,

Figure 10.17 Sway-gradients estimated using SPGA. Phase gradient estimation kernel (SPGA) and phase curvature kernel (SPCA) shown.

scene variation (self-clutter and complicated, extended targets in particular) limit autofocus potential more than those noise sources and is significantly more difficult to test. More advanced simulation systems may aid in this type of testing and allow meaningful results to be considered while also accounting for noise effects.

Using a simulated point scatterer target scene in a noise-free situation, SPGA shows clear image improvement and convergence in 3-4 iterations. The phase-gradient based estimator gives better performance than the phase-curvature based estimator (bear in mind there is no yaw in the simulation). The best results for this autofocus scenario were obtained using a window reduction rate of $\alpha = 0.6$ combined with a modified centroid along-track position estimator.

For a scene of point-like scatterers without corrupting noise or clutter SPGA effectively provides autofocus convergence within 3 iterations. With accurate target windowing, autofocus improvement is reliable with the other autofocus parameters affecting the amount of improvement possible. Target window width estimation techniques would improve the robustness of the algorithm.

10.6.2   Field-collected data

The field-collected data for testing the autofocus algorithms was obtained in July 2001 in Sydney Harbour with the KiwiSAS-II. The image used is from a different pass of the same region as shown in the field-collected data shown in Section 9.3.2. Like the

Figure 10.18   Pulse-compressed image of Sydney Harbour scene. Note the large sonar calibration rail at $u = 15$ m. Data is sampled at approximately $D/3$.

data in Section 9.3.2, the image was collected with an along-track sample spacing of approximately $D/3$ and has a shallow grazing angle.

Figures 10.18 and 10.19 correspond to the collected pulse-compressed and reconstructed images respectively. No micronavigation techniques or INS were used in the formation of the images. The images are linear grey-scale and are clipped to -20 dB below the image peak. In addition, the 5 m at each of the range extents of the images have been removed to remove cross-talk and prevent circular convolution artefacts biasing the autofocus results. Salient features of the image include a 20 m long sonar calibration rail from (17,15) to (37,15) made up of small retro-reflectors, an approximately 3 m long 0.5 m diameter pipe from (32,-11) to (32,-14) (reflections from both front and rear wall are visible) and apparent shadowing at (12,-15). Other features are strong multi-path reflections from the pipe, a range-aliased target from (58,5) appearing at (6,5) and visible alias-lobes from mild aperture undersampling. Note that the alias-lobe level is higher than expected for $D/3$ sampling because of spatial Doppler spectrum expansion caused by sway and a constant yaw offset.

A number of features of the initial reconstructed image Figure 10.19 make it difficult to autofocus. In particular, the scene has a large dynamic range with a number of strong-scattering, extended targets. The clutter background is also bland and lacks scene information or contrast changes. In addition, the centre region of the image has very few (if any) scatterers that are above the clutter. This makes autofocus over that region challenging. Furthermore, alias-lobes from the strong-scattering calibra-

Figure 10.19 Reconstructed image corresponding to the pulse-compressed data in Figure 10.18. Note visible alias-lobes appearing in the image. The alias-lobes in the image are not symmetric implying platform yaw. Differences in symmetry throughout the image are caused by yaw varying with along-track position.

tion rail are visible above the clutter in the region causing the problems discussed in Section 10.5. There is also alias-lobe asymmetry caused by a constant platform yaw (beam squint). Interestingly, the relative symmetry in the alias-lobes varies in the image implying a changing platform yaw. The combination of the features noted makes successful autofocus challenging.

The first of the autofocus methods attempted on Figure 10.19 was SPGA with a phase curvature estimator (equivalent to 2-D PCA). Three iterations of the algorithm were run each selecting 200 prominent-scatterers. The initial window size was 0.8 m by 5 m to match the expected blurring. The across-track window size was extended to ensure the independence of prominent-scatterers recommended by Zavattero [1999]. The window size was reduced at each iteration to 0.8 of the size at the previous iteration. Figure 10.20 shows the sway estimate at the end of the third iteration. Note that the sway estimate corresponds to a very large sway that is very unlikely. Figure 10.21 shows the image estimate once the sway of Figure 10.20 has been removed. The SPGA phase curvature focused result of Figure 10.21 does not show any improvement over the original reconstructed image. The autofocused image is significantly worse than the original degraded image.

The reason for the poor autofocus result from the phase curvature estimator is the same as for the generally poor performance of PCA: the double integration of

Figure 10.20 Plot of the estimated sway generated using three iterations of SPGA with the phase curvature estimator. The estimate corresponds to very large sways that are highly unlikely to occur in practice.

Figure 10.21 Autofocus result after applying three iterations of SPGA using the phase curvature estimator (equivalent to 2-D PCA). Autofocus performance is poor with as much as 6 m of sway estimated (this is a highly unlikely sway for the KiwiSAS-II).

Figure 10.22   Plot of the estimated sway double differences (sway curvatures) generated using three iterations of SPGA with the phase curvature estimator. The estimate has a high variance in the regions of little target cover (such as $y = -3\,\mathrm{m}$ to $13\,\mathrm{m}$). Double integration across these regions leads to large errors in the estimated sway response.

phase curvatures. For the image used in this set of results, there is a region in the centre of the image where phase curvature estimates are unreliable (due to only bland clutter in the scene). This can be seen as a region of large curvature variance in the sway double differences (sway-curvature) in Figure 10.22. When double integrating across the region(s) of unreliable data, both the sway offset and the sway-gradient are incorrectly estimated leading to large autofocus errors. The effect was also noted in the previous section with simulated data where the scene becomes disjoint with reasonable autofocus over each section of reliable data. This does not always result in well autofocused sections since the average sway-gradient over each is unknown and can have very large errors. In this case, the unknown sway-gradients over each aperture section cause $\approx 8\,\mathrm{m}$ sway estimates which are clearly erroneous[31]. Unknown sway-gradients, caused by double integration of sway-curvature, cause poor autofocus results in stripmap systems.

SPGA with the phase gradient estimator was the other method used to autofocus the field-collected data. Three iterations of the SPGA algorithm were run with the same settings as used above in combination with a modified-centroid along-track position estimator. In addition, an amplitude weighting was applied in the phase gradient kernel to reduce the effect of the strong-scattering, extended targets. This (admittedly ad-

---

[31]The error is exacerbated by removing the overall linear trend.

Figure 10.23 Sway estimate from three iterations of SPGA using the phase gradient estimation kernel. Note that the linear sway trend is caused by an uncompensated constant platform yaw during imaging. Accounting for the constant yaw would improve the autofocus result.

hoc) approach to prevent autofocus biasing changes the ML phase gradient estimation kernel into a FLOS-PGA style phase estimation kernel. In the results shown, the phase estimator used was

$$\mathbf{Arg}\left\{\int g^{<0.6>}(k_x, k_y) g^{*<0.6>}(k_x, k_y + \Delta k_y)\, \mathrm{d}k_x\right\}$$

where

$$\mathbf{Arg}\left\{\int g(k_x, k_y) g^*(k_x, k_y + \Delta k_y)\, \mathrm{d}k_x\right\}$$

is the original ML kernel. The altered amplitude weighting has only a small effect on the SPGA result and is not necessary with more advanced phase estimation approaches such as WPGA. The altered weighting had no noticeable effect when applied to phase curvature kernel SPCA.

The sway estimate from SPGA is shown in Figure 10.23. Note the linear sway trend. The linear sway estimate is caused by a constant unaccounted-for yaw during collection (beam squint). The yaw degrades autofocus performance and ideally should be estimated, corrected and accounted for within SPGA. Correction via demodulation/reconstruction alone is not enough as SPGA needs the yaw information to map position and spatial-Doppler frequencies via the wavenumber transform.

The SPGA-autofocused image corresponding to the sway of Figure 10.23 is shown in Figure 10.24. SPGA with a phase gradient estimator gives significant image im-

Figure 10.24 Autofocus result after three iterations of SPGA with a phase gradient estimator. (Sway of Figure 10.23 used for motion compensation.) Note general image improvement, lowering of alias-lobes and sharper response of the calibration rail target.

provement. In particular, the alias-lobes are reduced in amplitude (due to less aliasing once sway is compensated) and the calibration rail is in much sharper focus. Residual image blurring is still evident but overall blurring has been reduced significantly without biasing caused by strong-scattering extended targets such as the pipe or the close-range section of the calibration rail. Note that the autofocus performance varies throughout the image, some image regions will be better than others. In particular, the bland region in the centre of the image causes poorer autofocus performance due to alias-lobes and lack of prominent scatterers.

Overall, in a situation with unknown and possibly varying platform yaw, SPGA with the phase gradient estimator performed better than expected. The phase gradient estimator is significantly better than the phase curvature estimator even in the presence of unknown platform yaw. The approach of using an estimate of the along-track target position and using gradient averaging has benefited algorithm performance. Many further methods exist for improving SPGA's performance by improving phase and along-track target position. These improvements will allow better autofocus results.

10.6.3   SAS autofocus testing difficulties

SAS autofocus techniques are notoriously difficult to evaluate [Callow et al. 2001c]. Ideally, field collected data should be used for all autofocus testing. In practice, the motion and phase distortions causing blurring are unable to be measured in the field

preventing quantitative autofocus assessment. Qualitative (and sometimes subjective) assessment is possible by investigating field-collected imagery before and after autofocus. The best looking result is not necessarily the one with the most accurate autofocus. Accurate ground-truth data from a conventional high-frequency side-scan sonar could help but would be of limited benefit since frequency dependent scattering significantly alters the image (see images in [Marx et al. 2000]). Field-collected data from free-towed SAS systems cannot currently be used for quantitative autofocus testing.

Quantitative autofocus testing with field-collected data is possible using the spotlight SAR testing practice of injecting known phase errors into diffraction-limited imagery [Jakowatz et al. 1996]. The availability of the (field-collected) Sandia labs test set [Jakowatz et al. 1996; Morrison 2002; Wahl et al. 1994a] aids in the comparison of spotlight SAR autofocus algorithms. Currently, no diffraction-limited stripmap SAS dataset is available to allow similar testing.

The burden of quantitative autofocus algorithm testing falls to simulated data sets. Simulation allows timely and quantitative autofocus testing and comparison. However, simulation assumptions and inaccuracies can cause misleading results.

Simulation approximations and autofocus

Autofocus performance must be considered carefully when using simulated data sets. Seemingly innocuous assumptions and approximations with regard to imagery can have large impacts on autofocus performance. An example is given by Zavattero [1999] for PGA testing. This work suggests that PGA performance can be underestimated if clutter is not included in the simulation model. Generally, clutter degrades overall autofocus performance but assists early iterations.

Testing with simulated clutter needs to be clearly thought out as the spatial correlation lengths and amplitude probability density functions have effects on autofocus performance. Moreover, occlusion, shadowing, and non-Lambertian scattering have *significant* effects. Very few (if any) open simulation systems can accurately model frequency and position variant phase effects in shadow zones. A simulation suite is currently under development to address these issues [Hunter et al. 2003]. With the interest in applying SAS to mine-detection, and the similarity of mine detection and autofocus algorithms, detailed autofocus testing in simulated mine detection environments is necessary.

The simulations used in this thesis do not attempt to solve the problems noted above. The simulations are based on ray-tracing to point scatterers and are usually constructed in the ground-plane. Where clutter is simulated, it is generated using a point scatterer based simulation of many randomly placed targets. These targets have Rayleigh distributed amplitude to model the expected returns from a coherent speckle surface. Future work in autofocus testing should make use of more sophisticated

simulation techniques.

10.7 Summary

Autofocus is often required to improve imagery suffering unknown path or medium fluctuation errors. A wide-beam, wide-band expression was developed relating sonar sway to the complicated image blurring in SAS images. This expression employs fewer approximations and is more accurate than the blurring models used in PCA and PGA. The model relies on the wavenumber transform which maps spatial-frequency data of a local patch to the pulse-compressed data of that patch. An extension of the model was derived for the case of large sways that cause the wavenumber transform and other blurring models to fail.

Traditional stripmap autofocus techniques were discussed and found to be special cases of the SPGA algorithm with varying approximations. SPGA uses traditional techniques such as scene windowing, phase estimation and interpolation. In addition, the SPGA algorithm employs a novel along-track position estimation that allows phase gradient averaging improving autofocus performance and robustness at the expense of requiring a larger window size at convergence.

Implementation details of the major SPGA components are discussed in detail and it was found that windowing should employ constant factor reduction for SAS autofocus, target selection should ensure sufficient targets are selected for accuracy and aperture coverage, and modified Doppler centroiding should be used for along-track position estimation. Replacement of the ML phase estimation kernel with one similar to the WPGA kernel will remove the requirement for along-track position estimation and improve phase estimation accuracy. Other SPGA implementation details, limiting factors and future improvements were also presented.

Image features causing difficulties for SPGA were discussed with a particular emphasis on aperture undersampling and space-invariant blurring. SPGA performance is slightly degraded on images containing undersampling artefacts. Ambiguous results are possible where space-invariant blurring occurs and accurate parameter estimates from micronavigation should be used to estimate and remove the blurring.

The proposed SPGA algorithm was verified on both simulated and field-collected data. The quantitative simulation-based testing was used to investigate variations of a number of possible autofocus parameters. In all simulation tests SPGA gave image improvement except when the along-track window size was reduced too rapidly. SPGA should be used with phase gradient estimation, modified centroid along-track position estimation, and window reduction rates of 50%–70%. Autofocus testing on field-collected data also showed image improvement using SPGA. SPGA using the phase gradient estimation showed a large performance improvement over phase curvature

based estimation. The double integration of SPCA causes early autofocus divergence and a poor overall result. SPGA provides image improvement in 3–4 iteration on both simulated and field-collected data. Further improvements to the algorithm such as yaw compensation and enhanced phase estimation will provide better autofocus results.

Verification of autofocus results is difficult with a lack of accurate ground truth for field-collected data and inadequate modelling of complicated scenes in simulation. A commonly available test set with ground-truth data would aid quantitative autofocus testing. Further testing on various field-collected scenes is necessary.

Conclusions

The first conclusion drawn is that the improved wide-band system models derived in Chapter 3 should be used when using Fourier-based reconstruction methods. These models have been derived without resorting to the stationary phase method and are more accurate as a result. There is no additional computation involved in reconstruction and the improved modelling results in more accurate imagery. Improved temporal Doppler modelling should be also used in wide-beam, multiple-receiver systems to remove the effects of the stop-and-hop model. The approximation is not normally considered serious for SAS imagery, causing image and spatial Doppler shifting. However, the shifts can have adverse effects on autofocus performance. Compensation requires minor changes to Fourier-based algorithms and adds no additional computational burden for the wavenumber algorithm.

Of the Fourier-based reconstruction methods, the wavenumber algorithm should be used owing to its flexibility and accuracy. Phase errors have been reported when using the (more efficient) chirp-scaling algorithm, suggesting that the wavenumber algorithm is better when autofocus is required (see Section 4.6). When using the wavenumber algorithm, high-order interpolation is generally better than sub-swath reconstruction using low-order interpolation. Efficiency gains may be possible using sub-swath reconstruction with non-uniform sub-swath sizing since the required interpolator order varies with slant-range. The wavenumber algorithm provides an efficient reconstruction technique for data collected from complicated geometries when combined with appropriate motion-compensation techniques.

Compensation for the effects of multiple-receiver geometry is necessary when imaging targets in the near field of the receiver array. Improved wide-beam compensation is possible when using DFT-based along-track interpolation benefiting ultra-wide-beam systems with relatively long receiver arrays. Multiple-receiver reconstruction using the bistatic wavenumber algorithm holds some promise for avoiding the phase-centre approximation but the implementation is not currently as efficient as predicted. More research may reveal methods for improving its efficiency. Bistatic modelling of the multiple-receiver SAS problem gives a useful framework in which the complicated blur-causing yaw and sway motions can be described. This may lead to improved blurring

models for multiple-receiver autofocus algorithms.

The improved wide-beam motion compensation techniques derived in Chapter 6 should be used with any wide-beam multiple-receiver system. Multiple-receiver systems obtain improved imagery with only a small increase in processing cost over traditional narrow-beam methods. Single-receiver systems also derive benefit from wide-beam motion compensation but require more intensive processing. Yaw compensation methods have been improved particularly benefiting wide-beam systems where large relative yaws are present. The yaw compensation improvements require slight modifications to the wavenumber algorithm and require a small increase in computation.

Time-delay estimation techniques used for estimating platform motion should employ amplitude-only with subsequent phase-only correlations rather than full time-series correlations. This results in significant computational savings. Analysis of the Cramér-Rao lower bounds for the estimation show that the accuracy of the estimate is similar to that possible using full correlation. Phase-only correlation is equivalent to the phase-estimation problem in astronomical imaging and SAR autofocus; it also allows more flexibility in selecting target weighting compared with amplitude correlation. The eigenvector phase estimation kernel was found to be equivalent to those used in ultrasound and astronomical imaging and was shown to implement phase-closure. In addition, the improved accuracy of high-order eigenvector kernels at low SNR suggest that an eigenvector phase estimator of order 2–6 would be useful for SAS autofocus. The weighted phase kernel that the WPGA algorithm uses has better low SNR accuracy than the eigenvector kernel but requires modification to allow stripmap operation. Using WPGA for stripmap SAS requires further research. The weighted phase-kernel (and also the QPGA modifications to the eigenvector kernel) also offer the benefit of lessening the autofocus bias caused by strong-scattering, extended targets.

Unknown platform motion should initially be estimated using RPC for multiple-receiver systems and shear average for single receiver systems. These methods provide a good initial estimate of unknown path errors. Where shear average is used it should always have the improved amplitude weightings discussed in Section 9.3.1 rather than the ML weighting. Generally, shear average does not provide enough accuracy for diffraction-limited imagery and subsequent autofocus is needed. The image correlation extensions of RPC should not be used for autofocusing with non-redundant collection: they employ the same assumptions as traditional phase-gradient autofocus methods but without the flexibility, accuracy, or computational efficiency. These problems could be ameliorated in part if complex correlation can be used. This will require similar techniques to those employed by PDA (a correlation-based phase-gradient technique).

Spotlight autofocus algorithms such as PDA and PGA often provide the basis for similar stripmap algorithms. The space-invariant blurring model typical of spotlight algorithms is a special case of the blurring model derived from the stripmap wavenum-

ber transform. Parametric PDA should be chosen instead of map-drift variants since it offers similar performance without requiring iteration. PDA achieves this using the same information as PGA in determining path-estimates although PGA should be employed in preference to PDA or map-drift. In particular, the PGA-like WPGA algorithm should be used since its extensions to PGA provide improved phase estimation accuracy, giving better autofocus performance.

The SPGA algorithm for stripmap autofocus has been proposed and tested on field-collected and simulated data sets. The algorithm is based on traditional autofocus methods but has fewer approximations owing to improved blur modelling. Typically the SPGA algorithm converges in 3-4 iterations although divergence is sometimes seen with additional iterations when the window size becomes too small for reliable along-track position estimation. The likelihood of divergence is scene dependent. Both phase gradient and phase curvature based phase-estimation methods are possible with the use of along-track position estimation but phase gradient estimation gives improved autofocus results on both simulated and field-collected data. The SPGA algorithm is also able to operate on data that has $\geq D/2$ along-track sampling although it has reduced performance in that case.

SPGA is the algorithm most suited to single receiver stripmap autofocus discussed in this thesis. SPGA can be used for bulk sway removal but it is more efficient to use micronavigation or an INS for that purpose. SPGA's use of phase gradient techniques allows additional flexibility, improved performance, and reduced computational cost when compared with other techniques. In addition, the use of phase gradient estimation allows better accuracy and robustness than phase curvature techniques. Future work on the algorithm should initially focus on multiple-receiver autofocus and on improving phase estimation. Other possibilities for improvement are discussed in the following section.

In summary, this thesis offers a number of methods for improving the imagery obtained from stripmap SAS systems: improvements to current modelling, reconstruction, and motion compensation techniques, a comprehensive review of current SAS and SAR autofocus techniques, and the development and testing of the novel SPGA algorithm which holds much promise for SAS/SAR autofocus.

11.1   Recommendations for future research

There are a number of areas in which further research could yield improvements in the imagery obtained from SAS systems. A number of these are listed below:

Bistatic reconstruction improvement — The multiple-receiver wavenumber reconstruction based on the bistatic formulation of the sonar system model holds promise for systems where the phase centre approximation does not hold. Current

implementations require more processing than desirable. A better understanding of the problem may result in efficiency gains.

Autofocus algorithm efficiency — Algorithmic efficiency is sometimes sacrificed for accuracy. In the early iterations of an autofocus algorithm, where accuracy is not as important, computational savings are possible if approximations are made in the autofocus algorithms. Hybrid autofocus schemes do this giving significant benefit to global optimisation style algorithms. More research is needed to determine where approximations can be made without sacrificing performance.

Improved phase estimation — WPGA's phase estimator is better than the ML estimator currently used with the SPGA algorithm but relies on a number of spotlight autofocus assumptions. In addition, the local phase unwrapping technique of WPGA's phase estimator has the potential to remove the requirement for along-track position estimation in SPGA.

Continuous wavelet-transform based autofocus — All of the autofocus methods described in this thesis employ some form of space/spatial-frequency method for estimating the path distortion—most based on short-term Fourier transformation. A wavelet-transform based framework (similar to the Wigner-Ville distribution methods of Brown and Ghiglia [1988], Berizzi and Pinelli [1997] and bispectral estimation methods [Nikias et al. 2000]) may provide additional insight into the autofocus problem.

Additional autofocus priors — All autofocus requires some prior knowledge about the images to be processed. Currently, most of the priors used for autofocus are closely related: contrast and entropy priors implement the same constraints as those based on the average Fourier phase or bispectral phase of an image patch—the so called point scatterer or image sharpness assumption. Additional information based on alternate prior knowledge may provide benefits, particularly if used in an iterative global autofocus framework.

Multiple-receiver autofocus improvement — The stripmap autofocus framework presented in this thesis does not model multiple-receiver systems. Currently, multiple-receiver images are autofocused in the same way as single-receiver images. This neglects some of the complicated interactions possible with multiple-receiver systems. Multiple-receiver systems allow additional prior information (for example that the array is fixed) that is not currently used within the SPGA method.

Target detection investigation — The techniques used in target classification and detection (such as the higher order spectra technique outlined in [Chandran et al.

2002]) have a close relationship with autofocus techniques. In particular, the high-order phase estimators used in autofocus are equivalent to higher order spectral estimation techniques used elsewhere. There is some possibility the techniques can be combined. Further investigation of the relationships between the two fields is necessary.

Advanced simulation testing — Adequately testing autofocus algorithms is difficult. Simulation systems (the most straightforward way of accurately verifying autofocus performance) need to better model effects such as frequency-dependent beam patterning and shadowing—seemingly minor effects that have drastic impact on autofocus performance. In addition, further testing on field-collected data is necessary to show robust performance.

SAS system parameters

This appendix summarises the parameters of some of the SAS systems discussed in this thesis. Transmitter extents are given in terms of an equivalent rectangular transducer—this is not necessarily the same as the transducer extent since it is common practice to defocus a transmitter to mimic a smaller one.

Parameter	Value	Units
Carrier freq	30	kHz
Bandwidth	20	kHz
Number receivers	1	
Receiver length	0.225	m
Transmitter length	0.325	m

Table A.1 KiwiSAS-II parameters [Hawkins 1996; Hayes and Gough 1999].

Parameter	Value	Units
Carrier freq	30/100	kHz
Bandwidth	20	kHz
Number receivers	1	
Receiver length	0.225	m
Transmitter length	0.325	m

Table A.2 KiwiSAS-III (LF/HF) parameters [Hayes et al. 2001].

Parameter	Value	Units
Carrier freq	30/100	kHz
Bandwidth	20	kHz
Number receivers	12	
Receiver length	0.1	m
Transmitter length	0.1	m

Table A.3 KiwiSAS-IV (LF/HF) parameters (under construction).

Parameter	Value	Units
Carrier freq	180	kHz
Bandwidth	30	kHz
Number receivers	32	
Receiver length	0.0625	m
Transmitter length	0.0625	m

Table A.4   US Navy sonar parameters approximately matching specification given by Keeter [2001].

Parameter	Value	Units
Carrier freq	20/180	kHz
Bandwidth	10/30	kHz
Number receivers	14/11	
Receiver length	0.0381/0.0508	m
Transmitter length	0.762/0.0508	m

Table A.5   CSS sonar (LF/HF) parameters [Cook et al. 2001].

Derivation of Weyl's Identity

This section outlines the Fourier decomposition that Soumekh [1994, pp 149-152] performs, i.e., it outlines a proof of the 2-D equivalent of Weyl's identity and gives the wavenumber domain representation of the received field. 3-D derivations follow in similar manner. Chew [1995] gives a similar derivation. This section is intended to replace the Fourier pair derivation presented in Hawkins [1996, Appendix A].

For convenience Weyl's identity is stated initially and the proof of the right-hand side follows. Weyl's identity in 3-D space is given by

$$\frac{\exp\left(-jkr\right)}{4\pi r} = \frac{-j}{8\pi^2} \iint \frac{\exp\left(-j\left|x\right|\sqrt{k^2 - k_y^2 - k_z^2} + jk_y y + jk_z z\right)}{\sqrt{k^2 - k_y^2 - k_z^2}} \, \mathrm{d}k_y \, \mathrm{d}k_z, \quad \text{(B.1)}$$

where $r = \sqrt{x^2 + (y - u)^2 + (z - h)^2}$. The 2-D equivalent is given by

$$-\frac{j}{4} H_0^{(2)}(k\rho) = \frac{-j}{4\pi} \int \frac{\exp\left(-j\left|x\right|\sqrt{k^2 - k_y^2} + jk_y y\right)}{\sqrt{k^2 - k_y^2}} \, \mathrm{d}k_y, \quad \text{(B.2)}$$

where $\rho = \sqrt{x^2 + (y - u)^2}$ and $H_0^{(2)}$ is a Hankel function of the 2nd kind. The left hand sides of (B.1) and (B.2) and their derivations were presented in Section 2.4 and correspond to the system models in 3-D and 2-D respectively; these impulse responses should replace the simple exponential term used in Hawkins [1996].

The right hand side of (B.2) is derived by starting with the free space homogeneous medium Green's equation in 2-D [Morse and Feshbach 1953; Ziomek 1995] (also the starting point for the derivation in Section 2.4)

$$G(k_x, k_y, \omega) = \frac{1}{k_x^2 + k_y^2 - k^2}, \quad \text{(B.3)}$$

and extending to the more general forced form below, known as the forced Helmholtz equation,

$$H(k_x, k_y, \omega) = \frac{F(k_x, k_y, \omega)}{k_x^2 + k_y^2 - k^2}. \quad \text{(B.4)}$$

Taking the Inverse Fourier transform of (B.4) gives,

$$H(x,y,\omega) = \frac{1}{(2\pi)^2} \int_{k_x} \int_{k_y} \frac{F(k_x,k_y,\omega)}{k_x^2 + k_y^2 - k^2} \exp\left(j(k_x x + k_y y)\right) dk_x \, dk_y$$

$$H(x,y,\omega) = \frac{1}{(2\pi)^2} \int_{k_x} \int_{k_y} \frac{F(k_x,k_y,\omega)}{(k_x - k_1)(k_x + k_1)} \exp\left(j(k_x x + k_y y)\right) dk_x \, dk_y, \tag{B.5}$$

where

$$k_1(k,k_y) \equiv \sqrt{k^2 - k_y^2}. \tag{B.6}$$

To simplify (B.5) the residue theorem [Kreyszig 1979] is invoked. To do this, first the integral in k_x is generalised onto a Laplace domain contour integral using[1] $s_x = \alpha + jk_x$. Summarising this for the integration over k_x we get:

$$H(x,y,\omega) = \frac{1}{(2\pi)^2} \int_{k_y} \int_{k_x} \frac{F(k_x,k_y,\omega)}{(k_x - k_1)(k_x + k_1)} \exp\left(j(k_x x + k_y y)\right) dk_x \, dk_y$$

$$= \frac{1}{(2\pi)^2} \int_{k_y} \oint_C \frac{F(s_x,k_y,\omega)}{(s_x - k_1)(s_x + k_1)} \exp\left(s_x x + jk_y y\right) ds_x \, dk_y, \tag{B.7}$$

where the contour C encloses the left-hand Laplace plane. Dividing this contour into two parts we can perform the integral along the k_x axis (what we care about) and an integral on an infinite radius. If the latter integral goes to zero the entire integral is equal to the section along the k_x axis [Kreyszig 1979]. Now if $x < 0$ then the term $\exp(s_x x) \to 0$ as $s_x \to \infty$ and that part of the integral disappears. Evaluating the residue at the left half pole gives:

$$R(x) = \frac{F(s_x,k_y,\omega)}{s_x + k_1} \exp\left(s_x x + jk_y y\right)\Bigg|_{s_x = -k_1}. \tag{B.8}$$

Evaluating the other pole using a contour extending over the right-half Laplace domain for $x > 0$ and evaluating the residues leads to the result below,

$$H(x,y,\omega) = \begin{cases} \frac{-j}{2\pi} \int_{k_y} \left[\frac{F(k_x,k_y,\omega)}{k_x + k_1} \exp\left(j(k_x x + k_y y)\right) \right]_{k_x = -k_1} dk_y, & \text{if } x > 0 \\ \frac{j}{2\pi} \int_{k_y} \left[\frac{F(k_x,k_y,\omega)}{k_x - k_1} \exp\left(j(k_x x + k_y y)\right) \right]_{k_x = k_1} dk_y, & \text{if } x < 0. \end{cases} \tag{B.9}$$

Explicitly expanding the residues, (B.9) becomes:

$$H(x,y,\omega) = \begin{cases} \frac{-j}{2\pi} \int_{k_y} \frac{F(\sqrt{k^2 - k_y^2},k_y,\omega)}{2\sqrt{k^2 - k_y^2}} \exp\left(j(-x\sqrt{k^2 - k_y^2} + k_y y)\right) dk_y, & \text{if } x > 0 \\ \frac{j}{2\pi} \int_{k_y} \frac{-F(-\sqrt{k^2 - k_y^2},k_y,\omega)}{2\sqrt{k^2 - k_y^2}} \exp\left(j(x\sqrt{k^2 - k_y^2} + k_y y)\right) dk_y, & \text{if } x < 0. \end{cases} \tag{B.10}$$

To complete the proof of (B.2) $F(\sqrt{k^2 - k_y^2}, k_y, \omega)$ is set to 1 for all k, k_y, ω to match

---

[1]The α term can be regarded as modelling the absorption in the medium.

the original impulse response and (B.10) may be written as

$$H(x, y, \omega) = \frac{-j}{4\pi} \int \frac{\exp\left(-j\,|x|\,\sqrt{k^2 - k_y^2} + jk_y y\right)}{\sqrt{k^2 - k_y^2}}\, \mathrm{d}k_y. \tag{B.11}$$

To derive the wavenumber domain representation of the field we take the positive[2] x term of (B.10) to get:

$$H(x, y, \omega) = \frac{-j}{4\pi} \int_{k_y} \frac{F(\sqrt{k^2 - k_y^2}, k_y, \omega))\exp(-jx\sqrt{k^2 - k_y^2})}{\sqrt{k^2 - k_y^2}}\, \exp(jk_y y)\, \mathrm{d}k_y, \tag{B.12}$$

which is of the form:

$$H(x, y, \omega) = \frac{-j}{2}\mathcal{F}^{-1}_{(k_y)}\left[\frac{F(\sqrt{k^2 - k_y^2}, k_y, \omega)\exp\left(-jx\sqrt{k^2 - k_y^2}\right)}{\sqrt{k^2 - k_y^2}}\right]. \tag{B.13}$$

Often (B.13) is written without the explicit frequency dependence as

$$h(x, y) = \frac{-j}{2}\mathcal{F}^{-1}_{(k_y)}\left[\frac{f(\sqrt{k^2 - k_y^2}, k_y)\exp\left(-jx\sqrt{k^2 - k_y^2}\right)}{\sqrt{k^2 - k_y^2}}\right]. \tag{B.14}$$

---

[2]As we are treating single-sided imaging.

Time varying gain for Fourier-based reconstruction

In the reconstruction chapters of this thesis TVG was discussed with regard to time-domain based reconstruction and time-domain methods were derived using the 3-D Green's function. The Fourier reconstruction techniques were derived using the 2-D Green's function and require slightly different TVG terms to get correct reconstruction. This section relates the 2-D and 3-D Green's via a delta-sheet model and derives the correct TVG and frequency compensation terms for Fourier-based reconstruction.

Starting with the expression for the raw echo frequency data from Section 4.1

$$
E(\omega, u, h) \approx \iiint \frac{f(x, y, z) P(\omega)}{(4\pi)^2 \left| x^2 + (u - y)^2 + (h - z)^2 \right|} \\
\exp\left(-j2k\sqrt{x^2 + (u - y)^2 + (h - z)^2} \right) \mathrm{d}z \, \mathrm{d}x \, \mathrm{d}y, \quad \text{(C.1)}
$$

a *delta-sheet* scattering approximation (single scattering height for a given (x, y) position) is made such that

$$
f(x, y, z) = f(x, y)\delta(z - Z), \quad \text{(C.2)}
$$

and h is assumed to be zero giving

$$
E(\omega, u) \approx \iint \frac{f(x, y) P(\omega)}{(4\pi)^2 \left| x^2 + (u - y)^2 + Z^2 \right|} \\
\exp\left(-j2k\sqrt{x^2 + (u - y)^2 + Z^2} \right) \mathrm{d}x \, \mathrm{d}y. \quad \text{(C.3)}
$$

At this point the effective scatterer height function Z is a constant (but could be written to vary with scene position if necessary). Rewriting (C.3) in slant range coordinates

$$
x_s = \sqrt{x^2 - Z^2}
$$

(see Chapter 3) gives

$$E(\omega, u) \approx \iint \frac{f(x_s, y) P(\omega)}{(4\pi)^2 |x_s^2 + (u-y)^2|}$$
$$\exp\left(-j2k\sqrt{x_s^2 + (u-y)^2}\right) \frac{x_s}{\sqrt{x_s^2 - Z^2}} \, \mathrm{d}x_s \, \mathrm{d}y. \quad \text{(C.4)}$$

Performing an inverse Fourier transform of (C.4) from $\omega \to t$ gives the echo data as

$$e(t, u) \approx \iint \frac{f(x_s, y) \, p(t - (2/c)\sqrt{x_s^2 + (u-y)^2})}{(4\pi)^2 |x_s^2 + (u-y)^2|} \frac{x_s}{\sqrt{x_s^2 - Z^2}} \, \mathrm{d}x_s \, \mathrm{d}y, \quad \text{(C.5)}$$

which after pulse-compression becomes

$$s(t, u) \approx \iint \frac{f(x_s, y) \, pp(t - (2/c)\sqrt{x_s^2 + (u-y)^2})}{(4\pi)^2 |x_s^2 + (u-y)^2|} \frac{x_s}{\sqrt{x_s^2 - Z^2}} \, \mathrm{d}x_s \, \mathrm{d}y, \quad \text{(C.6)}$$

where $pp(t)$ is the autocorrelation of the transmitted signal. Applying a TVG[1] term of

$$4\pi ct/2 \approx 4\pi \sqrt{x_s^2 + (u-y)^2}$$

to $s(t, u)$ gives

$$s_{\text{tvg}}(t, u) \approx \iint \frac{f(x_s, y) \, pp(t - (2/c)\sqrt{x_s^2 + (u-y)^2})}{4\pi \sqrt{x_s^2 + (u-y)^2}} \frac{x_s}{\sqrt{x_s^2 - Z^2}} \, \mathrm{d}x_s \, \mathrm{d}y, \quad \text{(C.7)}$$

which is

$$S_{\text{tvg}}(\omega, u) \approx \iint \frac{f(x_s, y) \, PP(\omega)}{4\pi \sqrt{x_s^2 + (u-y)^2}} \exp\left(-j2k\sqrt{x_s^2 + (u-y)^2}\right) \frac{x_s}{\sqrt{x_s^2 - Z^2}} \, \mathrm{d}x_s \, \mathrm{d}y$$
$$\text{(C.8)}$$

in the frequency domain. Taking the Fourier transform[2] along u, y (and using an asymptotic expansion of the resulting Hankel function) gives (C.8) as

$$S_{\text{tvg}}(\omega, k_u) \approx \int \frac{f(x_s, k_y) \, PP(\omega)}{4\pi \sqrt{jx_s k_x}} \exp\left(-jk_x x_s\right) \frac{x_s}{\sqrt{x_s^2 - Z^2}} \, \mathrm{d}x_s, \quad \text{(C.9)}$$

$$\approx \frac{PP(\omega)}{4\pi \sqrt{jk_x}} \int f(x_s, k_y) \exp\left(-jk_x x_s\right) \frac{\sqrt{x_s}}{\sqrt{x_s^2 - Z^2}} \, \mathrm{d}x_s, \quad \text{(C.10)}$$

where

$$k_x = \sqrt{k^2 - k_y^2}. \quad \text{(C.11)}$$

The result (and method of TVG application) may be seen using (C.10); absorbing the $\sqrt{x_s} \, \sqrt{x_s^2 - Z^2}$ term into a modified reflectivity, $f_{\text{amp}}(x_s, y)$, allows the pulse

---

[1]The successful application of this term assumes that the autocorrelation of the transmitted signal is short. This is necessary for the application of TVG to mimic the desired RVG.

[2]Using Weyrich's identity [Chew 1995, page 70].

compressed data after TVG to be written as

$$S_{\text{tvg}}(\omega, k_u) \approx \frac{f_{\text{amp}}(k_x, k_y) \, PP(\omega)}{\sqrt{jk_x}}. \tag{C.12}$$

The method of reconstructing scenes using 2-D Fourier domain techniques with a 3-D Green's function is as follows:

1. Apply linear TVG to the pulse-compressed time-domain data.

$$s_{\text{tvg}}(t, u) = s(t, u)2\pi ct. \tag{C.13}$$

2. Reconstruct $f_{\text{amp}}(x, y)$ from s_{tvg} using (for example the wavenumber algorithm)

$$f_{\text{amp}}(k_x, k_y) \approx S_{\text{tvg}}(\omega, k_u)\sqrt{jk_x}, \tag{C.14}$$

where k_x is given by (C.11). Note that the modification to the frequency compensation term ($\sqrt{jk_x}$ instead of k_x) is the same for any of the Fourier-domain methods described in Chapter 4.

3. Recover $f(x, y)$ from $f_{\text{amp}}(x, y)$ by removing the space variant amplitude term via

$$f(x_s, y) = f_{\text{amp}}(x_s, y)\frac{\sqrt{x_s^2 - Z^2}}{\sqrt{x}}. \tag{C.15}$$

The derivations shown in this appendix demonstrate the relationship between 2-D and 3-D Green's functions with regard to Fourier-domain reconstruction techniques. In practice the 3-D Green's function should always be used to model typical SAS scenarios (even in shallow water) and the TVG described here applied in preference to the methods described in Chapters 4 and 5. Those methods require only very minor adjustment to apply the appropriate TVG terms.

Derivation of the wavenumber transform

The wavenumber transform is the basis of the stripmap blurring model used in this thesis and as such is also the basis of the SPGA autofocus method.

D.1 Stationary phase derivation

The stationary phase method is often used to derive the wavenumber algorithm [Bonifant 1999; Hawkins 1996]. In this situation it is required that the sway does not shift the stationary phase point—i.e., the effect of sway is negligible. This approximation is only valid if the sway is small.

This example of the derivation follows the derivation of the wavenumber algorithm in [Hawkins 1996, Appendix A]. Starting with a simplified 2-D spatial-frequency response (from Chapter 3 neglecting spreading losses)

$$E(\omega, k_u) \approx \int a(\omega, u) \iint f(x, y) \exp\left(-j2k\sqrt{x^2 + (u - y)^2} - jk_u u\right) \mathrm{d}x \, \mathrm{d}y \, \mathrm{d}u, \quad \text{(D.1)}$$

with a given amplitude-only filter $a(\omega, u)$. Note that for the point at x_n, y_n given by $f_n(x_n, y_n)$ (D.1) is given by

$$E(\omega, k_u)_n \approx \int a(\omega, u) f_n(x_n, y_n) \exp\left(-j2k\sqrt{x_n^2 + (u - y_n)^2} - jk_u u\right) \mathrm{d}u, \quad \text{(D.2)}$$

The phase and phase-derivatives of (D.1) are,

$$\phi(\omega, u) = -2k\sqrt{x_n^2 + (u - y_n)^2} - k_u u \tag{D.3}$$

$$\dot{\phi}(\omega, u) = \frac{2k(u - y_n)}{\sqrt{x_n^2 + (u - y_n)^2}} - k_u. \tag{D.4}$$

Solving for the stationary phase point

$$u^* = y_n - \frac{k_u x_n}{\sqrt{4k^2 - k_u^2}}, \tag{D.5}$$

which after recognising the wavenumber algorithm's change of variables becomes

$$u^* = y_n - \frac{k_y x_n}{k_x}. \tag{D.6}$$

The expression for the stationary phase point u^* is substituted back into (D.1) to derive the echo signal's wavenumber domain representation

$$E(\omega, k_u) \approx a\left(\omega, y_n - \frac{k_u x_n}{\sqrt{4k^2 - k_u^2}}\right)\sqrt{\frac{\pi x}{jk}} f_n\left(\sqrt{4k^2 - k_u^2}, k_u\right), \tag{D.7}$$

where $f_n(k_x, k_y)$ is the Fourier transform of the small image patch $f_n(x_n, y_n)$.

Equation (D.6) gives an expression relating the spatial-frequencies from a given position in the reconstructed image to the pulse compressed data that generated the result at that position. The relation can be used to express the amplitude-only filter $a(k, u)$ in terms of k_x, k_y. Sway is mapped in a similar fashion although it causes phase terms that shift the stationary phase point. The shift in stationary phase-point caused by large sways invalidates the model given here. The breakdown of the wavenumber transform for large sways is important for stripmap autofocus and improved modelling should be used (see Section 10.1.1 and [Madsen 2001]). Small sways such as those experienced at autofocus convergence have little effect on the wavenumber transform.

D.2   Geometry based derivation

The wavenumber transform can also derived using straightforward geometrical arguments as in Soumekh [1999].

Starting with the angle to target θ

$$\theta = \tan^{-1}\left(\frac{u - y_n}{x_n}\right) \tag{D.8}$$

and noting that in the far-field of the target patch the spatial frequency coverage is described by the same angle

$$\theta = \tan^{-1}\left(\frac{k_y}{k_x}\right). \tag{D.9}$$

The wavenumber transform relation may now be derived as

$$u = y_n - \frac{k_y x_n}{k_x} \tag{D.10}$$

which matches that derived in the previous section. Note that

$$k_x = 2k\cos\theta, \tag{D.11}$$

$$k_y = 2k\sin\theta, \tag{D.12}$$

are equivalent to the use of the wavenumber transform—i.e., only a far field approximation is necessary. The use of the later equations also gives

$$2k = \sqrt{k_x^2 + k_y^2}, \tag{D.13}$$

which is the second half of the wavenumber transform mapping.

Appendix E

Eigenvector phase kernel and phase-closure

The eigenvector method estimates the phase by calculating the eigenvector corresponding to the largest eigenvalue of the data sample covariance matrix. This is derived, maximising the functional below[1]:

$$\mathcal{Q}_1 = \mathbf{v}^H \hat{\mathbf{C}} \mathbf{v}, \tag{E.1}$$

where v is the phase only vector $\begin{bmatrix} 1 & \exp\left(j\Delta\phi_1[l]\right) & \dots & \exp\left(j\Delta\phi_L[l]\right) \end{bmatrix}$ and $\hat{\mathbf{C}}$ is the co-variance matrix of the measured data. For the sake of clarity we define,

$$\alpha_a \equiv g[n, l + a], \tag{E.2}$$

for integer a. Expanding (E.1) for two adjacent sets of samples (pulses) we get,

$$\mathcal{Q}_1 =$$

$$\frac{1}{N}\begin{bmatrix} 1 \\ \exp\left(j\Delta\phi_1[l]\right) \\ \exp\left(j\Delta\phi_2[l]\right) \end{bmatrix}^H \begin{bmatrix} \sum\limits_{n=0}^{N-1}|\alpha_0|^2 & \sum\limits_{n=0}^{N-1}\alpha_0\alpha_1^* & \sum\limits_{n=0}^{N-1}\alpha_0\alpha_2^* \\ \sum\limits_{n=0}^{N-1}\alpha_0^*\alpha_1 & \sum\limits_{n=0}^{N-1}|\alpha_1|^2 & \sum\limits_{n=0}^{N-1}\alpha_1\alpha_2^* \\ \sum\limits_{n=0}^{N-1}\alpha_0^*\alpha_2 & \sum\limits_{n=0}^{N-1}\alpha_1^*\alpha_2 & \sum\limits_{n=0}^{N-1}|\alpha_2|^2 \end{bmatrix} \begin{bmatrix} 1 \\ \exp\left(j\Delta\phi_1[l]\right) \\ \exp\left(j\Delta\phi_2[l]\right) \end{bmatrix}. \tag{E.3}$$

After ignoring the constant amplitude terms, $\sum\limits_{n=0}^{N-1}|\alpha_0|^2, \sum\limits_{n=0}^{N-1}|\alpha_1|^2, \dots$ (E.3) becomes,

$$\mathcal{Q}_2 = \frac{1}{N}\sum_{n=0}^{N-1}\Big\{\alpha_0\alpha_1^*\exp\left(j\Delta\phi_1[l]\right) + \alpha_0^*\alpha_1\exp\left(-j\Delta\phi_1[l]\right)$$

$$+ \alpha_0\alpha_2^*\exp\left(j\Delta\phi_2[l]\right) + \alpha_0^*\alpha_2\exp\left(-j\Delta\phi_2[l]\right)$$

$$+ \alpha_1\alpha_2^*\exp\left(-j\Delta\phi_1[l]\right)\exp\left(j\Delta\phi_2[l]\right) + \alpha_1^*\alpha_2\exp\left(j\Delta\phi_1[l]\right)\exp\left(-j\Delta\phi_2[l]\right)\Big\}$$

$$\tag{E.4}$$

---

[1]A proof exists in [Jakowatz and Wahl 1993].

as the expression to be maximised. Equation (E.4) can now be written in the form

$$\mathcal{Q}_2 = q + q^*, \tag{E.5}$$

where,

$$
\begin{aligned}
q =& \frac{1}{N} \sum_{n=0}^{N-1} [\alpha_0 \alpha_1^* \exp(j\Delta\phi_1[l])] \\
&+ \frac{1}{N} \sum_{n=0}^{N-1} [\alpha_0 \alpha_2^* \exp(j\Delta\phi_2[l])] \\
&+ \frac{1}{N} \sum_{n=0}^{N-1} [\alpha_1 \alpha_2^* \exp(-j\Delta\phi_1[l]) \exp(j\Delta\phi_2[l])].
\end{aligned} \tag{E.6}
$$

Using the relationship $q + q^* = 2\mathrm{Re}\{q\}$, (E.6) becomes,

$$
\begin{aligned}
\mathcal{Q}_2 =& 2\mathrm{Re}\left\{ \frac{1}{N} \sum_{n=0}^{N-1} \left(\alpha_0 \alpha_1^* \exp(j\Delta\phi_1[l])\right) \right\} \\
&+ 2\mathrm{Re}\left\{ \frac{1}{N} \sum_{n=0}^{N-1} \left(\alpha_0 \alpha_2^* \exp(j\Delta\phi_2[l])\right) \right\} \\
&+ 2\mathrm{Re}\left\{ \frac{1}{N} \sum_{n=0}^{N-1} \left(\alpha_1 \alpha_2^* \exp(-j\Delta\phi_1[l]) \exp(j\Delta\phi_2[l])\right) \right\}.
\end{aligned} \tag{E.7}
$$

Converting (E.7) using $\mathrm{Re}\{q\} \equiv |q| \cos(\arg(q))$,

$$
\begin{aligned}
\mathcal{Q}_2 =& 2\left| \frac{1}{N} \sum_{n=0}^{N-1} (\alpha_0 \alpha_1^*) \right| \cos\left(\Delta\phi_1[l] + \arg\left(\frac{1}{N} \sum_{n=0}^{N-1} \alpha_0 \alpha_1^* \right) \right) \\
&+ 2\left| \frac{1}{N} \sum_{n=0}^{N-1} (\alpha_0 \alpha_2^*) \right| \cos\left(\Delta\phi_2[l] + \arg\left(\frac{1}{N} \sum_{n=0}^{N-1} \alpha_0 \alpha_2^* \right) \right) \\
&+ 2\left| \frac{1}{N} \sum_{n=0}^{N-1} (\alpha_1 \alpha_2^*) \right| \cos\left(\Delta\phi_2[l] - \Delta\phi_1[l] + \arg\left(\frac{1}{N} \sum_{n=0}^{N-1} \alpha_1 \alpha_2^* \right) \right).
\end{aligned} \tag{E.8}
$$

By inspection, maximisation of $\mathcal{Q}_2$ occurs when each cos term is maximised[2]. Using the principal solution[3] for $\mathcal{Q}_2$ in an over-determined set of equations we have,

$$\Delta\phi_1[l] = -\arg\left(\frac{1}{N}\sum_{n=0}^{N-1}\alpha_0\alpha_1^*\right),$$

$$\Delta\phi_2[l] = -\arg\left(\frac{1}{N}\sum_{n=0}^{N-1}\alpha_0\alpha_2^*\right),\qquad\text{(E.9)}$$

$$\Delta\phi_2[l] - \Delta\phi_1[l] = -\arg\left(\frac{1}{N}\sum_{n=0}^{N-1}\alpha_1\alpha_2^*\right).$$

Which can be written in matrix form as:

$$\begin{bmatrix} 1 & 0 \\ -1 & 1 \\ 0 & 1 \end{bmatrix}\begin{bmatrix}\Delta\phi_1[l] \\ \Delta\phi_2[l]\end{bmatrix} = \begin{bmatrix} -\arg\left(\frac{1}{N}\sum_{n=0}^{N-1}\alpha_0\alpha_1^*\right) \\ -\arg\left(\frac{1}{N}\sum_{n=0}^{N-1}\alpha_1\alpha_2^*\right) \\ -\arg\left(\frac{1}{N}\sum_{n=0}^{N-1}\alpha_0\alpha_2^*\right) \end{bmatrix}.\qquad\text{(E.10)}$$

Solving for $[\Delta\phi_1[l], \Delta\phi_2[l]]$ in a least squares sense [Franklin et al. 1998] we get,

$$\begin{bmatrix}\Delta\phi_1[l] \\ \Delta\phi_2[l]\end{bmatrix} = \begin{bmatrix}\frac{2}{3} & \frac{1}{3} & -\frac{1}{3} \\ \frac{1}{3} & \frac{2}{3} & \frac{1}{3}\end{bmatrix}\begin{bmatrix} -\arg\left(\frac{1}{N}\sum_{n=0}^{N-1}g[n,l]g^*[n,l+1]\right) \\ -\arg\left(\frac{1}{N}\sum_{n=0}^{N-1}g[n,l+1]g^*[n,l+2]\right) \\ -\arg\left(\frac{1}{N}\sum_{n=0}^{N-1}g[n,l]g^*[n,l+2]\right) \end{bmatrix}.\qquad\text{(E.11)}$$

A weighted LMS solution to (E.10) may be obtained by including a weighting in the equations above.

To estimate the corrupting aperture phase, $\varphi[l]$, the resulting phase estimates are integrated in a more complicated fashion than other methods:

$$\varphi[l] = (\varphi[l-1] + \Delta\phi_1[l]) + (\varphi[l-2] + \Delta\phi_2[l]).\qquad\text{(E.12)}$$

Notice now that there are multiple estimates for $\varphi[l]$ that can be combined to improve performance[4]. The combination of multiple phase differences is a typical phase-closure relation. The multiple phase difference estimates may also be weighted to account for varying noise properties or other anomalies in the collected data. The original

---

[2]Note that this is equivalent to the phase closure relation used in astronomical imaging. [Lohmann et al. 1983; Roddier 1986]

[3]The assumption involved here is that the angles $\Delta\phi_L[l]$ are small. When the angles are large, phase unwrapping can allow this method to be used. Also there are recursive methods for phase estimation not requiring phase unwrapping [Bartelt et al. 1984].

[4]As long as the noise is independent.

eigenvector solution also leads to the bispectrum methods used in astronomical imaging [Ayers et al. 1988] (another phase-closure method [Roddier 1986]). Freeman et al. [1988] describe a number of different methods of weighting bispectrum phases and conclude that minimum variance weighting gives reliable results. Satherly [1994, Chapter 4] provides a summary of phase retrieval techniques in astronomical imaging.

Bispectrum and eigenvector phase estimation equivalence

This appendix outlines the equivalence of using bispectrum-based phase estimation and using a 3-term eigenvector phase estimation.

F.1 Bispectrum kernel

Phase estimation for $\varphi[l]$ from bispectrum data for $j = -1, 1, \ldots$

$$\varphi[l] = \varphi[j] + \varphi[l-j] - \arg\left(\frac{1}{N}\sum_{n=0}^{N-1}\left(gG^*[n,l]gG[n,-j]gG[n,l+j]\right)\right), \qquad \text{(F.1)}$$

rewriting each separate estimate,

$$\varphi[l] =$$
$$\begin{cases} \varphi[-1] - \arg\left(\frac{1}{N}\sum_{n=0}^{N-1}gG[n,-1]\right) + \varphi[l-1] - \arg\left(\frac{1}{N}\sum_{n=0}^{N-1}\left(gG^*[n,l]gG[n,l+1]\right)\right), \\ \varphi[+1] - \arg\left(\frac{1}{N}\sum_{n=0}^{N-1}gG[n,+1]\right) + \varphi[l+1] - \arg\left(\frac{1}{N}\sum_{n=0}^{N-1}\left(gG^*[n,l]gG[n,l-1]\right)\right). \end{cases}$$
$$\text{(F.2)}$$

F.2 Eigenvector kernel

Phase estimation from a 3 term eigen-method estimates uses covariance matrix difference phases

$$\begin{aligned} \Delta\phi_1[l] &= \varphi[l+1] - \varphi[l] &= -\arg\left(\frac{1}{N}\sum_{n=0}^{N-1}g[n,l]g^*[n,l+1]\right), \\ \Delta\phi_2[l] &= \varphi[l+2] - \varphi[l] &= -\arg\left(\frac{1}{N}\sum_{n=0}^{N-1}g[n,l]g^*[n,l+2]\right), \qquad \text{(F.3)} \\ \Delta\phi_1[l+1] &= \varphi[l+2] - \varphi[l+1] &= -\arg\left(\frac{1}{N}\sum_{n=0}^{N-1}g[n,l+1]g^*[n,l+2]\right), \end{aligned}$$

rewriting[1] for the phase estimate $\varphi[l]$,

$$
\varphi[l] = \begin{cases} \varphi[l-1] - \arg\left(\frac{1}{N} \sum_{n=0}^{N-1} gG[n, l-1]gG^*[n, l] \right), \\ \varphi[l+1] - \arg\left(\frac{1}{N} \sum_{n=0}^{N-1} gG[n, l+1]gG^*[n, l] \right). \end{cases}
\tag{F.4}
$$

Clearly (F.4) is the same as (F.2) except for a constant phase factor,

$$
cp = \begin{cases} \varphi[-1] - \arg\left(\frac{1}{N} \sum_{n=0}^{N-1} gG[n, -1] \right), \\ \varphi[+1] - \arg\left(\frac{1}{N} \sum_{n=0}^{N-1} gG[n, +1] \right). \end{cases}
\tag{F.5}
$$

The constant phase factor, cp, shown in (F.5) is trivial to account for. To remove the difference, the phase retrieval algorithm only has to undergo minor modification. The trispectrum result is the same but with a constant phase of different form.

---

[1]But possibly throwing away information at this point.

Fourier transform properties

This appendix summaries useful properties of Fourier transforms used extensively throughout this thesis.

Property	Description
1. Linearity	$af(t) + bg(t) \leftrightarrow aF(\omega) + bG(\omega)$, where a, b are constants
2. Time Scaling	$f(at) \leftrightarrow \dfrac{1}{\|a\|}F\left(\dfrac{\omega}{a}\right)$, where a is a constant
3. Duality	If $f(t) \leftrightarrow F(\omega)$ then $F(t) \leftrightarrow f(-\omega)$
4. Time shifting	$f(t - t_0) \leftrightarrow F(\omega)\exp\left(-j\omega t_0\right)$
5. Frequency shifting	$\exp\left(j\omega_c t\right)f(t) \leftrightarrow F(\omega - \omega_c)$
6. Area under $f(t)$	$\displaystyle\int_{-\infty}^{\infty} f(t)\,\mathrm{d}t = F(0)$
7. Area under $F(\omega)$	$f(0) = \displaystyle\int_{-\infty}^{\infty} F(\omega)\,\mathrm{d}\omega$
8. Differentiation in the time-domain	$\dfrac{\mathrm{d}}{\mathrm{d}t}f(t) \leftrightarrow j\omega F(\omega)$
9. Integration in the time-domain	$\displaystyle\int_{-\infty}^{t} \leftrightarrow \dfrac{1}{j\omega}F(\omega) + \dfrac{F(0)}{2}\delta(\omega)$
10. Conjugate functions	If $f(t) \leftrightarrow F(\omega)$ then $f^*(t) \leftrightarrow F^*(-\omega)$
11. Multiplication in the time-domain	$f(t)g(t) \leftrightarrow \displaystyle\int_{-\infty}^{\infty} F(\lambda)G(\omega - \lambda)\,\mathrm{d}\lambda$
12. Convolution in the time-domain	$\displaystyle\int_{-\infty}^{\infty} f(\tau)g(t - \tau)\,\mathrm{d}\tau \leftrightarrow F(\omega)G(\omega)$
13. Correlation in the time-domain	$\displaystyle\int_{-\infty}^{\infty} f(\tau)g^*(\tau - t)\,\mathrm{d}\tau \leftrightarrow F(\omega)G^*(\omega)$

$f(t)$	$F(\omega)$
$\delta(t)$	1
1	$\delta(\omega)$
$\delta(t - t_0)$	$\exp\left(-j\omega t_0\right)$
$\mathrm{rect}\left(\dfrac{t}{T}\right)$	$T\mathrm{sinc}(\omega T)$
$\mathrm{sinc}(2Bt)$	$\dfrac{1}{2B}\mathrm{rect}\left(\dfrac{\omega}{2B}\right)$
$\exp\left(-\pi t^2\right)$	$\exp\left(-\omega^2/4\pi\right)$
$\exp\left(-\left|a\right|t\right),\ a > 0$	$\dfrac{2a}{a^2 + \omega^2}$
$f(t) \odot_t f(t)$	$F(\omega)F(\omega)$
$f(t) \star_t f(t)$	$F(\omega)F^*(\omega)$

Table G.1   Useful 1-D Fourier transform pairs, compiled from [Bracewell 1986; Haykin 1994].

$f(x, y)$	$f(k_x, k_y)$
$f(x, y - ax)$	$f(k_x + ak_y, k_y)$
$f(x\cos\theta - y\sin\theta, x\sin\theta + y\cos\theta)$	$f(k_x\cos\theta - k_y\sin\theta, k_x\sin\theta + k_y\cos\theta)$

Table G.2   Useful 2-D Fourier transform pairs, compiled from [Bracewell 1986; Haykin 1994].

References

A. E. Adams, O. R. Hinton, M. A. Lawlor, B. S. Sharif, and V. S. Riyait. A synthetic aperture sonar image processing system. *IEE Acoustic Sensing and Imaging Conference*, pages 109–113, Mar. 1993.

A. E. Adams, M. A. Lawlor, V. S. Riyait, O. R. Hinton, and B. S. Sharif. Real-time synthetic aperture sonar processing system. *IEE Proceedings on Radar, Sonar, and Navigation*, 143(3):169–176, June 1996.

E. H. Attia and B. D. Steinberg. Self-cohering large antenna arrays using the spatial correlation properties of radar clutter. *IEEE Transactions on Antennas and Propagation*, 37(1):30–38, Jan. 1989.

S. R. J. Axelsson. Suppressed ambiguity in range by phase-coded waveforms. In *International Geoscience and Remote Sensing Symposium 2001*, volume 5, pages 2001–2005, June 2001.

G. R. Ayers, M. J. Northcott, and J. C. Dainty. Knox-Thompson and triple-correlation imaging through atmospheric turbulence. *Journal of the Optical Society of America A*, 5(7):963–985, July 1988.

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.

S. M. Banks and H. D. Griffiths. Imaging and motion estimation for synthetic aperture sonar based on Fast Factorised Back-Projection. In *Proceedings of the sixth European Conference on Underwater Acoustics (ECUA) 2002*, pages 529–534, June 2002.

S. M. Banks, T. J. Sutton, and H. D. Griffiths. A technique for interferometric synthetic aperture sonar image processing. In *Proc. London Communication Symposium*, UCL, London, UK, Sept. 2001.

S. Barbarossa and A. Scaglione. Autofocusing of SAR images based on the product of high-order ambiguity function. *IEE Proceedings on Radar, Sonar, and Navigation*, 145(5):269–273, Oct. 1998.

H. Bartelt, A. W. Lohmann, and B. Wirnitzer. Phase and amplitude recovery from bispectra. *Applied Optics*, 23(18):3121–3129, Sept. 1984.

R. H. T. Bates and F. M. Cady. Toward true imaging by wideband speckle interferometry. *Optics Communications*, 32(3):365–369, 1980.

V. Behar. Techniques for phase correction in coherent ultrasound imaging systems. *Ultrasonics*, 39:603–610, 2002.

A. Bellettini and M. A. Pinto. Theoretical accuracy of synthetic aperture sonar micronavigation using a displaced phase centre antenna. *IEEE Journal of Oceanic Engineering*, 27(4):780–789, October 2002.

F. Berizzi, G. Corsini, M. Diani, and M. Veltroni. New time-domain clutter-lock algorithm. In *IEE Proceedings on Radar, Sonar, and Navigation*, volume 144, pages 341–347. IEE, Dec. 1997.

F. Berizzi, G. Corsini, F. Gatti, and F. Gini. Cumulant based algorithms for autofocusing in ISAR/SAR systems. In *Proceedings of International Conference on Image Processing (ICIP '96)*, pages 887–890, Lausanne, Switzerland, Sept. 1996. IEEE.

F. Berizzi and G. Pinelli. Maximum-likelihood ISAR image autofocusing technique based on instantaneous frequency estimation. *IEE Proceedings on Radar, Sonar, and Navigation*, 144(5):284–292, Oct. 1997.

D. Billon and M. A. Pinto. Some general consideration for synthetic aperture sonar design. In *Proceedings of the IEEE OCEANS'95 Conference*, volume 3, pages 1665–1670. IEEE, 1995.

D. Blacknell, A. P. Blake, C. J. Oliver, and R. G. White. A comparison of SAR multilook registration and contrast optimisation autofocus algorithms applied to real SAR data. In *International Radar Conference*, pages 363–366, 1992.

W. W. Bonifant, Jnr. Interferometric synthetic aperture sonar processing. Master's thesis, Georgia Institute of Technology, July 1999.

W. W. Bonifant, Jnr., M. A. Richards, and J. H. McClellan. Interferometric height estimation of the seafloor via synthetic aperture sonar in the presence of motion errors. *IEE Proceedings on Radar, Sonar, and Navigation*, 147(6):322–330, Dec. 2000.

R. N. Bracewell. *The Fourier transform and its applications*. McGraw-Hill Book Company, New York, NY, 1986.

W. D. Brown and D. C. Ghiglia. Some methods for reducing propagation-induced phase errors in coherent imaging systems. I. formalism. *Journal of the Optical Society of America A*, 5(6):924–941, June 1988.

W. S. Burdic. *Underwater acoustic system analysis.* Prentice Hall, Englewood Cliffs, NJ, 1984.

L. J. Busse. Three-dimensional imaging using a frequency domain synthetic aperture focusing technique. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 39(2):174–179, Mar. 1992.

C. Cafforio, C. Prati, and F. Rocca. Full resolution focussing of Seasat SAR images in the frequency-wave number domain. *Journal of Robotic Systems*, 12(3):491–510, 1991a.

C. Cafforio, C. Prati, and F. Rocca. SAR data focussing using seismic migration techniques. *IEEE Transactions on Aerospace and Electronic Systems*, 27(2):194–207, Mar. 1991b.

H. J. Callow, P. T. Gough, and M. P. Hayes. Not quite 20,000 leagues under the sea, sonar autofocus in a shallow water reverberant environment. In *Image and Vision Computing New Zealand 2000*, pages 304–309. IVCNZ, November 2000.

H. J. Callow, M. P. Hayes, and P. T. Gough. Advanced wavenumber domain processing for reconstruction of broad-beam multiple-receiver synthetic aperture imagery. In *Image and Vision Computing New Zealand 2001*, pages 51–56. IVCNZ, November 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 *Proceedings of International Geoscience And Remote Sensing Symposium (IGARSS) 2001*, pages 1601–1603. IEEE, July 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*, volume 1, pages 157–162. IEEE, November 2001c.

H. J. Callow, M. P. Hayes, and P. T. Gough. A wavefront sensing method for synthetic aperture sonar autofocus. In *Proceedings of the sixth European Conference on Underwater Acoustics (ECUA) 2002*, pages 231–234, Gdansk, Poland, 2002a.

H. J. Callow, M. P. Hayes, and P. T. Gough. Wavenumber domain reconstruction of SAR/SAS imagery using single transmitter and multiple-receiver geometry. *Electronics Letters*, 38(7):336–337, March 2002b.

T. M. Calloway, C. V. Jakowatz, Jnr., P. A. Thompson, and P. H. Eichel. Comparison of synthetic-aperture radar autofocus techniques: phase gradient versus subaperture. *SPIE-Advanced Signal Processing Algorithms, Architectures, and Implementations II*, 1566:353–364, Dec. 1991.

L. Camp. *Underwater acoustics.* Wiley-Interscience, New York., 1970.

W. G. Carrera, R. S. Goodman, and R. M. Majewski. *Spotlight synthetic aperture radar: signal processing algorithms.* Artech House, 1995.

H. C. Chan and T. S. Yeo. Noniterative quality phase-gradient autofocus (QPGA) algorithm for spotlight SAR imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 36(5):1531–1539, Sept. 1998.

V. Chandran, S. Elgar, and A. Nguyen. Detection of mines in acoustic images using higher order spectral features. *IEEE Journal of Oceanic Engineering*, 27(3):610–618, July 2002.

E. Chang and M. D. Tinkle. Coherence of pulsed signal and implication to synthetic aperture processing. In *OCEANS 2001*, pages 193–201, 2001.

J. Chanussot, A. Hétet, G. le Merrer, and E. Tireau. Multispectral decomposition of synthetic aperture sonar images for speckle reduction. In *Proceedings of the sixth European Conference on Underwater Acoustics (ECUA) 2002*, pages 269–274, Gdansk, Poland, June 2002.

J. Châtillon. Absence of Nyquist criterion for spatial sampling of a strip-map synthetic aperture sonar. In *Proceedings of the fifth European Conference on Underwater Acoustics (ECUA) 2000*, volume 1, pages 401–406, July 2000.

J. Châtillon, A. E. Adams, M. A. Lawlor, and M. E. Zakharia. SAMI: A low-frequency prototype for mapping and imaging of the seabed by means of synthetic aperture. *IEEE Journal of Oceanic Engineering*, 24(1):4–15, Jan. 1999.

J. Châtillon, M. Bouhier, and M. E. Zakharia. Synthetic aperture sonar for seabed imaging: Relative merits of narrow-band and wide-band approaches. *IEEE Journal of Oceanic Engineering*, 17(1):95–105, Jan. 1992.

D. Chevillon, T. Rastello, and D. Vray. A Fourier-based motion estimation approach for wide band synthetic aperture sonar. In *Proceedings of the IEEE OCEANS'98 Conference*, volume 3, pages 1671–1675. IEEE, 1998.

W. C. Chew. *Waves and Fields in Inhomogeneous Media.* IEEE Press, 1995.

Z.-H. Cho, J. P. Jones, and M. Singh. *Foundations of Medical Imaging.* Wiley Interscience, 605 Third Avenue, New York, NY, 1993.

J. Christoff, C. Loggins, and E. Pipkin. Measurement of the temporal phase stability of the medium. *Journal of the Acoustical Society of America*, 71(6):1606–1607, June 1982.

C. E. Cook and M. Bernfeld. *Radar signals: an introduction to theory and application.* Academic Press Inc., 1967.

D. A. Cook, J. T. Christoff, and J. E. Fernandez. Broadbeam multi-aspect synthetic aperture sonar. In *Proceedings of the IEEE OCEANS 2001 Conference*, volume 1, pages 188–192. MTS/IEEE, 2001.

I. G. Cumming, F. H. Wong, and R. K. Raney. A SAR processing algorithm with no interpolation. In *International Geoscience and Remote Sensing Symposium*, volume 1, pages 376–379. IEEE, 1992.

J. C. Curlander and R. N. McDonough. *Synthetic Aperture Radar: Systems and signal processing.* John Wiley & Sons, Inc., 605 Third Avenue, New York, NY, 1996.

L. J. Cutrona. Comparison of sonar system performance achievable using synthetic-aperture techniques with the performance achievable by more conventional means. *Journal of the Acoustical Society of America*, 58(2):336–348, Aug. 1975.

P. de Heering. A synthetic aperture sonar study. Technical report, Huntech-Lapp Systems Limited, Scarborough, Ontario, MIR 3A6, Aug. 1982.

B. Douglas. *Signal processing for synthetic aperture sonar imaging.* PhD thesis, Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, 1993.

B. L. Douglas and H. Lee. Synthetic aperture active sonar imaging. In *IEEE Conference on Acoustics, Speech, and Signal Processing*, volume 3, pages 37–40. IEEE, 1992.

B. L. Douglas and H. Lee. Motion compensation for improved sidescan sonar imaging. In *Proceedings of the IEEE OCEANS'93 Conference*, volume 5, pages 378–383. IEEE, 1993a.

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, Apr. 1993b.

B. L. Douglas, H. Lee, and C. D. Loggins. A calibration technique for sidescan sonar systems based on the second-order statistics of returned signals. In *Proceedings of the IEEE OCEANS'92 Conference*, pages 306–311, Oct. 1992.

Dynamic Technology Incorporated. 2-D PCA description. Presented by Enson Chang in the SAS tutorial at OCEANS 2001, 2001.

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, Jan. 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, Oct. 1989.

S. A. Enright. *Towards Quantitative Computed Tomography*. PhD thesis, Department of Electrical and Electronic Engineering, University of Canterbury, Apr. 1992.

B. G. Ferguson and J. L. Cleary. Estimating the polar distribution of snapping shrimp with a wide-aperture array. In *Proceedings of the Fifth International Symposium on Signal Processing and Its Applications, 1999*, volume 2, pages 843–846. ISSPA, August 1999.

J. R. Fienup. Phase error correction by shear averaging. In *Signal Recovery and synthesis III*, pages 134–137. Optical Society of America, June 1989.

J. R. Fienup. Detecting moving targets in SAR imagery by focusing. *IEEE Transactions on Aerospace and Electronic Systems*, 37(3):794–809, July 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 (IGARSS) 2001*, volume 3, pages 1509–1511. 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*, volume 1, pages 163–169. IEEE, November 2001b.

G. F. Franklin, J. D. Powell, and M. L. Workman. *Digital control of dynamic systems*. Addison Wesley Longman, 1998.

J. D. Freeman, J. C. Christou, F. Roddier, D. W. McCarthy, Jnr., and M. L. Cobb. Application of bispectrum analysis for phase recovery from one-dimensional infrared speckle data. *Journal of the Optical Society of America A*, 5(3):406–415, Mar. 1988.

F. Gatelli, A. M. Guarnieri, F. Parizzi, P. Paswuali, C. Prati, and F. Rocca. The wavenumber shift in SAR interferometry. *IEEE Transactions on Geoscience and Remote Sensing*, 32(4):855–865, July 1994.

D. C. Ghiglia and G. A. Mastin. Two-dimensional phase correction of synthetic-aperture-radar imagery. *Optics Letters*, 14(20):1104–1106, Oct. 1989.

G. A. Gilmour. Synthetic aperture side-looking sonar system. *Journal of the Acoustical Society of America*, 65(2):557, May 1978. Review of U.S. Patent 4,088,978.

E. Gimeno and J. M. Lopez-Sanchez. Near-field 2-D and 3-D radar imaging using a chirp scaling algorithm. In *Proceeding of the International Geoscience and Remote Sensing Symposium 2001*, volume 1, pages 354–356. IGARSS, IEEE Press, July 2001.

J. W. Goodman. *Fourier Optics*. McGraw-Hill, New York, NY, 1968.

J. W. Goodman. Some fundamental properties of speckle. *Journal of the Optical Society of America*, 66(11):1145–1150, Nov. 1976.

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, Jan. 1997.

P. T. Gough and D. W. Hawkins. A short history of synthetic aperture sonar. In *International Geoscience and Remote Sensing Symposium*, volume 2, pages 618–620, 1998.

P. T. Gough and M. P. Hayes. Measurements of the acoustic phase stability in Loch Linnhe, Scotland. *Journal of the Acoustical Society of America*, 86(2):837–839, Aug. 1989a.

P. T. Gough and M. P. Hayes. Test results using a prototype synthetic aperture sonar. *Journal of the Acoustical Society of America*, 86(6):2328–2333, Dec. 1989b.

P. T. Gough, M. P. Hayes, and H. J. Callow. Strip-map path correction using phase matching autofocus. In *Proceedings of the fifth European Conference on Underwater Acoustics (ECUA) 2000*, pages 412–418, July 2000a.

P. T. Gough, M. P. Hayes, and D. R. Wilkinson. An efficient image reconstruction algorithm for a multiple hydrophone array synthetic aperture sonar. In *Proceedings of the fifth European Conference on Underwater Acoustics (ECUA) 2000*, pages 395–400, July 2000b.

P. T. Gough and R. G. Lane. Autofocussing SAR and SAS images using a conjugate gradient search algorithm. In *Proceedings of International Geoscience and Remote Sensing Symposium, 1998*, volume 2, pages 621–623. IEEE, 1998.

P. T. Gough and M. A. Miller. Displaced ping imaging autofocus. Submitted to the IEEE Journal of Oceanic Engineering, 2002.

J. Groen and J. C. Sabel. Interaction between autofocusing and synthetic aperture processing: a study based on simulated and experimental rail data. In *Proceedings of the sixth European Conference on Underwater Acoustics (ECUA) 2002*, pages 235–240, Gdansk, Poland, 2002.

J. Groen, J. C. Sabel, and A. Hétet. Synthetic aperture processing techniques applied to rail experiments with a mine hunting sonar. In *Proceedings of the UDT Europe*, 2001.

P. E. Hagen, R. E. Hansen, K. Gade, and E. Hammerstad. Interferometric synthetic aperture sonar for AUV based mine hunting: the SENSOTEK project. In *Unmanned Systems 2001*, Baltimore, USA, July 2001.

D. W. Hawkins. *Synthetic Aperture Imaging Algorithms: with application to wide bandwidth sonar.* PhD thesis, Department of Electrical and Electronic Engineering, University of Canterbury, Oct. 1996.

D. W. Hawkins and P. T. Gough. 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, Jan. 1997a.

D. W. Hawkins and P. T. Gough. Synthetic aperture imaging using the accelerated chirp scaling algorithm. In *International Geoscience and Remote Sensing Symposium*, volume 1, pages 471–473, Aug. 1997b.

M. P. Hayes. *A CTFM Synthetic Aperture Sonar*. PhD thesis, Department of Electrical and Electronic Engineering, University of Canterbury, Sept. 1989.

M. P. Hayes, P. J. Barclay, P. T. Gough, and H. J. Callow. Test results of a bathymetric synthetic aperture sonar. In *OCEANS 2001*, volume 3, pages 1682–1687. IEEE, November 2001.

M. P. Hayes, H. J. Callow, and P. T. Gough. Strip-map phase gradient autofocus. In D. Kenwright, editor, *Image and vision computing New Zealand 2002*, pages 71–76, Auckland, December 2002.

M. P. Hayes and P. T. Gough. Broad-band synthetic aperture sonar. *IEEE Journal of Oceanic Engineering*, 17(1):80–94, Jan. 1992.

M. P. Hayes and P. T. Gough. Imaging with a multiple-hydrophone synthetic aperture sonar. In *Image Vision Computing New Zealand 1999*, pages 163–168. IVCNZ, 1999.

S. Haykin. *Communication Systems – 3rd ed.* John Wiley & Sons, Inc., 605 Third Avenue, NY, 1994.

S. A. Hovanessian. *Introduction to synthetic array and imaging radars.* Artech House, Dedham, MA, 1980.

A. J. Hunter, M. P. Hayes, and P. T. Gough. Simulation of multiple-receiver, broadband interferometric SAS imagery. In *Oceans 2003*, San Diego, USA, September 2003. IEEE/MTS. CDROM.

B. D. Huxtable and E. M. Geyer. Motion compensation feasibility for high resolution synthetic aperture sonar. In *Proceedings of the IEEE OCEANS'93 Conference*, volume 1, pages I125–I131. IEEE, 1993.

T. Isernia, V. Pascazio, R. Pierri, and G. Schirinzi. Image reconstruction from Fourier transform magnitude with applications to synthetic aperture radar imaging. *Journal of the Optical Society of America*, 13(5):922–934, May 1996a.

T. Isernia, V. Pascazio, R. Pierri, and G. Schirinzi. Synthetic aperture radar imaging from phase corrupted data. *IEE Proceedings on Radar, Sonar, and Navigation*, 142 (4):268–274, Aug. 1996b.

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, Dec. 1993.

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.

F. B. Jensen, W. A. Kuperman, M. B. Porter, and H. Schmidt. *Computational Ocean Acoustics*. American Institute of Physics Press, 500 Sunnyside Boulevard, Woodbury, NY, 1993.

K. A. Johnson. Sonar motion detection using a common phase shift technique. Master's thesis, Department of Electrical and Electronic Engineering, University of Canterbury, 1992.

K. A. Johnson, M. P. Hayes, and P. T. Gough. Estimating sub-wavelength sway of sonar towfish. *IEEE Journal of Oceanic Engineering*, 20(4):258–267, Oct. 1995.

V. A. Kaharl. Sounding out the ocean's secrets. Web page, Feb. 2003. URL http://www.beyonddiscovery.org/content/view.txt.asp?a=219. Adapted from an article written by Victoria A. Kaharl for the National Academy of Sciences' Beyond Discovery: The Path from Research to Human Benefit.

M. Karaman, P. Li, and M. O'Donnell. Synthetic aperture imaging for small scale systems. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 42(3):429–442, May 1995.

H. Keeter. Navy awards Northrop Grumman \$3.5 million synthetic aperture sonar contract. *Defense Daily*, 209, Mar. 2001. Issue 50 (15 March), page numbers not applicable.

D. Kirk and P. Maloney. New autofocus technique for wideband wide-angle synthetic aperture radar. *SPIE-Algorithms for Synthetic Aperture Radar Imagery V*, 3370: 52–61, Sept. 1998.

K. T. Knox and B. J. Thompson. Recovery of images from atmospherically degraded short-exposure photographs. *Astrophysical Journal Letters*, 193:L45–L48, Oct. 1974.

E. Kreyszig. *Advanced Engineering Mathematics*. John Wiley & Sons, Inc., 605 Third Avenue, New York, NY, 1979.

S. Krishnan, K. W. Rigby, and M. O'Donnell. Improved estimation of phase aberration profiles. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 44(3):701–713, May 1997.

R. Lanari. A new method for wide field synthetic aperture radar processing. In *International Geoscience and Remote Sensing Symposium 1994*, volume 4, pages 1961–1963, Aug. 1994.

R. Lanari. A new method for the compensation of the SAR range cell migration based on the chirp Z-transform. *IEEE Transactions on Geoscience and Remote Sensing*, 33:1296–1299, 1995.

R. M. Lerner and R. C. Waag. Wave space interpretation of scattered ultrasound. *Ultrasound in Medicine and Biology*, 14(2):97–102, 1988.

C. J. Leuschen and R. G. Plumb. A matched-filter-based reverse-time migration algorithm for ground-penetrating radar data. *IEEE Transactions on Geoscience and Remote Sensing*, 39(5):929–936, May 2001.

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.

H.-J. Li, G.-T. Huang, and S.-L. Yen. Non-uniformly spaced array imaging. *IEEE Transactions on Aerospace and Electronic Systems*, 41(3):278–286, Mar. 1993.

Z. Liang and P. C. Lauterbur. *Principles of Magnetic Resonance Imaging*. Series in Biomedical Engineering. IEEE Press, 3 Park Avenue, NY, 2000.

S. Lockwood, A. Brown, and H. Lee. Backward propagation image reconstruction techniques for bistatic synthetic-aperture radar imaging systems with circular-aperture configurations. In *Signals, Systems and Computers, 2001. Conference Record of the Thirty-Fifth Asilomar Conference.*, volume 1, pages 110–115, 2001.

A. W. Lohmann, G. Weigelt, and B. Wirnitzer. Speckle masking in astronomy: triple correlation theory and applications. *Applied Optics*, 22(24):4028–4037, Dec. 1983.

J. M. Lopez-Sanchez and J. Fortuny-Guasch. 3-D radar imaging using range migration techniques. *IEEE Transactions on Aerospace and Electronic Systems*, 48(5):728–736, May 2000.

Q. H. Lui, N. Nguyen, and X. Y. Tang. Accurate algorithms for nonuniform fast forward and inverse Fourier transforms and their applications. In *Proceeding of the*

International Geoscience and Remote Sensing Symposium 1998, volume 1, pages 288–290. IGARSS, 1998.

S. N. Madsen. Estimating the Doppler centroid of SAR data. *IEEE Transactions on Aerospace and Electronic Systems*, 25(2):134–140, Mar. 1989.

S. N. Madsen. Motion compensation for ultra wide band SAR. In *Proceeding of the International Geoscience and Remote Sensing Symposium 2001*, pages 1436–1438. IGARSS, IEEE, July 2001.

D. Marx, M. Nelson, E. Chang, W. Gillespie, A. Putney, and K. Warman. An introduction to synthetic aperture sonar. In *IEEE Signal Processing Workshop on Statistical Signal and Array Processing*, pages 717–721, 2000.

A. S. Milman. SAR imaging using the $\omega - k$ migration. *International Journal of Remote Sensing*, 14:1965–1979, 1993.

A. Moreira. Improved multilook techniques applied to SAR and SCANSAR imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 29(4):529–534, July 1991.

A. Moreira, J. Mittermayer, and R. Scheiber. A SAR auto-focus technique based on azimuth scaling. In *Proceeding of the International Geoscience and Remote Sensing Symposium 1997*, volume 4, pages 2028–2030. IGARSS, IEEE Press, Aug. 1997.

R. L. Morrison, Jnr. Entropy based autofocus for synthetic aperture radar. Master's thesis, Graduate College of the University of Illinois at Urbana-Champaign, 2002.

R. L. Morrison, Jnr. and D. C. Munson, Jnr. An experimental study of a new entropy-based SAR autofocus technique. In *Proc. of the IEEE International Conference on Image Processing, Rochester*, volume 2, pages 441–444, Rochester, New York, 2002.

P. M. Morse and H. Feshbach. *Methods of Theoretical physics*, volume 1. McGraw Hill, New York, NY, 1953.

D. G. Muff, A. P. Blake, and A. M. Horne. Spaceborne SAR autofocus. *SPIE-Synthetic Aperture Radar and Passive Microwave Sensing*, 2584:7–17, Nov. 1995.

C. E. Nahum. Autofocusing using multiscale local correlation. *SPIE-SAR Image Analysis, Modeling, and Techniques*, 3497:21–30, Nov. 1998.

G. C. Ng, P. D. Freiburger, W. F. Walker, and G. E. Trahey. A speckle target adaptive imaging technique in the presence of distributed aberrations. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 44(1):140–151, Jan. 1997.

R. O. Nielsen. *Sonar signal processing*. Artech House, Inc., 685 Canton Steet, Norwood, MA, 1991.

C. L. Nikias, V. Z. Marmarelis, and D. C. Shin. Method for automatic focusing of radar or sonar imaging systems using high-order measurements. U.S. Patent 006037892, Jan. 2000.

C. L. Nikias and A. P. Petropulu. *Higher Order Spectra Analysis: a nonlinear signal processing framework*. Prentice Hall, Englewood Cliffs, 1993.

A. V. Oppenheim and R. W. Schafer. *Discrete-time signal processing*. Prentice Hall signal processing series. Prentice Hall, Upper Saddle River, N.J., 2nd edition, 1999.

C. B. Owen and F. Makedon. High Quality Alias Free Image Rotation. Technical Report PCS-TR96-301, Dartmouth College, Computer Science, Hanover, NH, November, 1996. URL ftp://ftp.cs.dartmouth.edu/TR/TR96-301.ps.Z.

J. B. Pat. Synthetic aperture sonar image reconstruction using a multiple-receiver tow-fish. Master's thesis, Department of Electrical and Electronic Engineering, University of Canterbury, Mar. 2000.

J. Pihl, P. Ulriksen, O. Kroling, B. Lovgren, and G. Shippey. MLO classification using an ROV-mounted wideband synthetic-aperture sonar. In *Proceedings of the fifth European Conference on Underwater Acoustics (ECUA) 2000*, volume 1, pages 433–438, July 2000.

E. N. Pilbrow, P. T. Gough, and M. P. Hayes. Dual-transponder precision navigation system for synthetic aperture sonar. In *Proceedings of Electronics New Zealand Conference*, Dunedin, New Zealand, 14-15 November 2002a.

E. N. Pilbrow, P. T. Gough, and M. P. Hayes. Long baseline precision navigation system for synthetic aperture sonar. In *Proceedings of The 11th Australasian Remote Sensing and Photogrammetry Association Conference*, Brisbane, Australia, 3-5 November 2002b.

M. A. Pinto, R. D. Hollet, A. Bellettini, and S. Chapman. Bathymetric imaging with wideband interferometric synthetic aperture sonar. *IEEE Journal of Oceanic Engineering*, 2002. Accepted for publication.

A. Potsis, A. Reigber, J. Mittermayer, A. Moreira, and N. Uzunoglou. Sub-aperture algorithm for motion compensation improvement in wide-beam SAR data processing. *Electronics Letters*, 37(23):1405–1407, Nov. 2001.

A. D. Poularikas, editor. *The transforms and applications handbook*. The electrical engineering handbook series. CRC Press, Boca Raton Fla., 1996.

C. Prati and F. Rocca. Focusing SAR data with time-varying Doppler centroid. *IEEE Transactions on Geoscience and Remote Sensing*, 30(3):550–559, May 1992.

D. Rachlin. Direct estimation of aberrating delay in pulse-echo imaging-systems. *Journal of the Acoustical Society of America*, 88(1):191–197, July 1990.

R. K. Raney. An exact wide field digital imaging algorithm. *Journal of Robotic Systems*, 13(5):991–998, 1992.

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.

R. S. Raven. Electronic stabilization for displaced phase centres systems. U.S. Patent 006037892, Jan. 1981.

A. Reigber. Correction of residual motion errors in airborne SAR interferometry. *Electronics Letters*, 37(17):1083–1084, Aug. 2001.

F. Roddier. Triple correlation as a phase closure technique. *Optics Communications*, 60(3):145–148, Nov. 1986.

K. D. Rolt and H. Schmidt. Azimuthal ambiguities in synthetic aperture sonar and synthetic aperture radar imagery. *IEEE Journal of Oceanic Engineering*, 17(1):73–79, Jan. 1992.

H. Runge and R. Bamler. A novel high precision SAR focusing algorithm based on chirp scaling. In *International Geoscience and Remote Sensing Symposium*, volume 1, pages 372–375, 1992.

B. L. Satherly. *Zero-based ensemble deconvolution and EEG spectral topography*. PhD thesis, Department of Electrical and Electronic Engineering, University of Canterbury, 1994.

A. A. Sawchuk. Space-variant image motion degradation and restoration. *Proceedings of the IEEE*, 60(7):854–861, July 1972.

R. W. Sheriff. Synthetic aperture beamforming with automatic phase compensation for high frequency sonars. In *Proceedings of the 1992 Symposium on Autonomous Underwater Vehicle Technology*, pages 236–245. IEEE, 1992.

G. Shippey, J. Pihl, and M. Jönsson. Autopositioning for wideband synthetic aperture sonar using fast factorised back projection. In *CAD/CAC 2001*, Halifax, Canada, Nov. 2001.

G. Shippey, P. Ulriksen, and Q. Liu. Wideband swath bathymetry, SAS autofocus and underwater navigation fixes: three related problems in echo/image correlation. In *IEE Sonar Signal Processing Conference*, Weymouth, Dec. 1998a. IA/IEE.

G. A. Shippey, P. Ulriksen, and Q. Liu. Quasi-barrowband processing of wideband sonar echoes. In A. Alippi and G. B. Canelli, editors, *Proceedings of the fourth European Conference on Underwater Acoustics (ECUA) 1998*, pages 63–68, Rome, 1998b.

J. M. Silkaitis, B. L. Douglas, and H. Lee. A cascade algorithm for estimating and compensating motion error for synthetic aperture sonar imaging. *IEEE Proceedings ICIP*, 1(1):905–909, Nov. 1994.

C. A. Snarski. Rank one phase error estimation for range-Doppler imaging. *IEEE Transactions on Aerospace and Electronic Systems*, 32(2):676–688, Apr. 1996.

M. Soumekh. Bistatic synthetic aperture radar inversion with application in dynamic object imaging. *IEEE Transactions on Signal Processing*, 39(9):2044–2055, Sept. 1991.

M. Soumekh. A system model and inversion for synthetic aperture radar imaging. *IEEE Transactions on Image Processing*, 1(1):64–76, Jan. 1992.

M. Soumekh. *Fourier Array Imaging.* Prentice Hall, Englewood Cliffs, NJ, 1994.

M. Soumekh. *Synthetic Aperture Radar Signal Processing with MATLAB Algorithms.* John Wiley & Sons, Inc., 1999.

H. Stark and J. W. Woods. *Probability, random processes, and estimation theory for engineers.* Prentice Hall, Englewood Cliffs, N.J., 2nd ed. edition, 1994.

R. H. Stolt. Migration by Fourier transform. *Geophysics*, 43(1):23–48, Feb. 1978.

P. N. Stoyle. Point scatterers in radar imaging. *SPIE-Radar Processing, Technology, and Applications III*, 3462:14–23, Oct. 1998.

T. J. Sutton, S. A. Chapman, and H. D. Griffiths. Robustness and effectiveness of autofocus algorithms applied to diverse seabed environments. In *Proceedings of the fifth European Conference on Underwater Acoustics (ECUA) 2000*, volume 1, pages 407–412, July 2000.

T. J. Sutton, H. D. Griffiths, and J. Robinson. Synthetic aperture sonar imaging using an unmanned underwater vehicle. In *Proc. IOA Conference on Acoustics*, volume 24, Salford, England, Mar. 2002.

D. G. Thompson, J. S. Bates, D. V. Arnold, and D. G. Long. Extending the phase gradient autofocus algorithm for low-altitude stripmap mode SAR. In *Proceeding of the International Geoscience and Remote Sensing Symposium 1999*, volume 1, pages 564–566, 1999.

K. Tomiyasu. Tutorial review of synthetic-aperture-radar (SAR) with applications to imaging of the ocean surface. *Proceedings of the IEEE*, 66(5):563–583, May 1978.

V. Tonard and M. Brussieux. Towards development of autofocusing schemes for phase compensation of synthetic aperture sonars. In *Proceedings of the IEEE OCEANS'97 Conference*, volume 2, pages 803–808, Nova Scotia, Canada, Oct. 1997. IEEE.

G. E. Trahey and L. F. Nock. Synthetic receive aperture imaging with phase correction for motion and for tissue inhomogeneities–Part II: effects of and correction for motion. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 39(4):496–501, July 1992.

G. E. Trahey and S. W. Smith. Properties of acoustical speckle in the presence of phase aberration part I:first order statistics. *Ultrasonic Imaging*, 10:12–28, 1988.

P. Tsakalides and C. L. Nikias. High-resolution autofocus techniques for SAR imaging based on fractional lower-order statistics. *IEE Proceedings on Radar, Sonar, and Navigation*, 148(5):267–276, Oct. 2001.

D. G. Tucker. *Underwater observation using sonar.* Fishing News Books Ltd, London, 1966.

L. Ulander, H. Hellsten, and G. Stenstrom. Synthetic aperture radar processing using Fast Factorised Back-Projection. In *EUSAR 2000, 3^{rd} European Conference on Synthetic Aperture Radar*, pages 753–756, Munich, Germany, 2000.

L. Ulander, H. Hellsten, and G. Stenstrom. Synthetic-aperture radar processing using Fast Factorised Back Projection. submitted to: IEEE Transactions on Aerospace and Electronic Systems, 8 December 2000, 2001.

R. J. Urick. *Principles of underwater sound.* McGraw Hill Book Company, New York, NY, 2nd edition, 1975.

M. A. van Dam and R. G. Lane. Tip/tilt estimation from defocused images. Accepted Journal Optical Society America, 2002a.

M. A. van Dam and R. G. Lane. Wave-front sensing from defocused images using wave-front slopes. Accepted Journal Optical Society America, 2002b.

H. L. Van Trees. *Detection, estimation and Modulation theory*, volume 1. John Wiley and Sons, 1968.

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 1994a.

D. E. Wahl, C. V. Jakowatz, Jnr., D. C. Ghiglia, and P. H. Eichel. Relationships between autofocus methods for SAR and self-survey techniques for SONAR. *SPIE-Applications of Digital Image Processing XIV*, 1567:32–40, Dec. 1991.

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, Oct. 1994b.

L. Wang, A. Bellettini, S. Fioravanti, S. Chapman, D. R. Bugler, Y. Perrot, and A. Hétet. InSAS'00: interferometric SAS and INS aided SAS imaging. In *Proceedings of the IEEE OCEANS 2001 Conference*, volume 1, pages 179–187. MTS/IEEE, 2001.

S. Wang and X. Huang. Autofocus techniques for reducing phase errors in UWB-SAR. In *Aerospace and Electronics Conference, 1997. NAECON 1997*, volume 2, pages 14–18. IEEE, July 1997.

D. W. Warner, D. C. Ghiglia, A. FitzGerrell, and J. Beaver. Two-dimensional phase gradient autofocus. *SPIE-Image Reconstruction from Incomplete Data*, 4123:162–173, 2000.

D. R. Wilkinson. Efficient image reconstruction techniques for a multiple-receiver synthetic aperture sonar. Master's thesis, Department of Electrical and Electronic Engineering, University of Canterbury, April 2001.

R. Williams. Creating an acoustic synthetic aperture in the ocean. *Journal of the Acoustical Society of America*, 60(1):60–73, July 1976.

L. Xi, L. Guosui, and J. Ni. Autofocusing of ISAR images based on entropy minimization. *IEEE Transactions on Aerospace and Electronic Systems*, 35:1240–1252, Oct. 1999.

S. Xiao, D. Munson, Jnr., S. Basu, and Y. Bresler. An $n^2 \log n$ back-projection algorithm for SAR image formation. In *Conference Record of the Thirty-Fourth Asilomar Conference on Signals Systems and Computers 2000*, volume 1, pages 3–7, 2000.

I. Yamaguchi. A technique of multiaperture transmitting-receiving on synthetic aperture sonar. *Electronics and Communication in Japan*, 82(3):66–73, 1999.

I. Yamaguchi and J. Kato. Evaluation of a synthetic aperture sonar with a multi-aperture transducer. In *Underwater Technology*, pages 487–492, 1998.

W. Ye, T. S. Yeo, and Z. Bao. An optimal method for phase error estimation in SAR imagery. In *Proceeding of the International Geoscience and Remote Sensing Symposium 2001*, volume 2, pages 1165–1167, July 1998.

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, Sept. 1999.

P. Zavattero. Distributed target SAR image de-blurring using phase gradient autofocus. In *The record of the Radar Conference*, pages 246–249. IEEE, 1999.

L. J. Ziomek. *Fundamentals of Acoustic Field Theory and Space-Time Signal Processing.* CRC Press, 1995.

Index