

# **Synthetic Aperture Sonar**

## **Motion Estimation and Compensation**

A Thesis  
Presented to  
The Academic Faculty

By

Daniel A. Cook

In Partial Fulfillment  
of the Requirements for the Degree  
Master of Science  
in  
Electrical and Computer Engineering



School of Electrical and Computer Engineering  
Georgia Institute of Technology  
May 2007

# **Synthetic Aperture Sonar**

## **Motion Estimation and Compensation**

Approved by:

Dr. Mark A. Richards, Advisor  
*Adjunct Professor, School of ECE*  
*Georgia Institute of Technology*

Dr. James H. McClellan  
*Professor, School of ECE*  
*Georgia Institute of Technology*

Dr. Anthony J. Yezzi  
*Associate Professor, School of ECE*  
*Georgia Institute of Technology*

Dr. Hao-Min Zhou  
*Assistant Professor, School of Mathematics*  
*Georgia Institute of Technology*

Date Approved: April 2007

## DEDICATION

*To my daughter Lillian. Learning brings me a great deal of pleasure in life,  
but caring for you gives me purpose in living.*

## ACKNOWLEDGMENTS

I consider myself blessed indeed to have the opportunity to work on synthetic aperture sonar imaging. It is a privilege to have a part in the development of SAS technology as it matures. Each time new imagery is collected, I cannot help but experience the awe and excitement that Verne wrote into his characters Captain Nemo and Professor Aronnax. Writing this thesis has been a pleasure, and I could not have done it without the kindness and help from many others.

The opportunity to carry out the research documented here was afforded by my employer, the Naval Surface Warfare Center in Panama City, Florida, and by the Office of Naval Research through their sponsorship of the synthetic aperture sonar program. I would also like to acknowledge several individuals at NSWC-PC. I am grateful to Dr. Kerry Commander for taking the gamble of hiring me and giving me the opportunity to work on SAS. I would like to thank Jim Christoff and Jose Fernandez for their supervision, support, and encouragement of my professional development. Both of these men are worthy of the title “mentor” in the noblest sense of the word. Their vast knowledge and patience offered an ideal environment in which to learn SAS. I also owe much to **DR. QUYEN HUYNH**, whose influence (and personality) deserve nothing less than boldface capital letters, and to Dr. John Lathrop for many enjoyable discussions regarding all aspects of SAS imaging.

My colleague Dan Brown deserves significant recognition. Much of the content of this thesis stems from our work together at NSWC-PC on SAS processing algorithms. If it were permissible, I would have listed Dan as coauthor. Dr. John Stroud, also of NSWC-PC, deserves recognition for singlehandedly writing real-time SAS processing code that was successfully demonstrated on board an AUV in December of 2006. His work is a historical first and is certainly a milestone for SAS technology.

On the academic side, I would like to recognize my friend and thesis advisor, Dr. Mark Richards. This document wouldn’t exist without his help, interest, and encouragement.



The assistance of Professors James McClellan, Anthony Yezzi, and Haomin Zhou is also greatly appreciated. Each has left their mark on my work. The guidance of Marilou Mycko in the ECE Graduate Affairs Office was, needless to say, instrumental to the success of this endeavor.

One of the best parts of my job is having the opportunity to collaborate with members of the SAS community around the world. It is truly a joy to work with people who share a passion for underwater imaging. I would especially like to recognize Prof. Peter Gough and Dr. Michael Hayes of the Acoustics Research Group at the University of Canterbury in Christchurch, New Zealand. Their proofreading and suggestions have improved the quality of this document. I would also like to recognize the SAS group at the NATO Undersea Research Centre in La Spezia, Italy. I have had numerous fruitful conversations with Mr. Andrea Bellettini and Dr. Hans Groen.

SAS imaging would not be possible without a sensor and a vehicle to carry it. It is a true pleasure to work with the hardware provided by the Applied Research Laboratory at the Pennsylvania State University. The sensors produced by the ARL/PSU group consistently give fantastic imagery and are a big part of the reason I have so much fun on the job. Of course, the vehicle is the other half of the hardware equation, and the data shown in this work was collected using vehicles provided by Woods Hole Oceanographic Institution and by Bluefin Robotics. These are amazing robots with the potential to greatly expand man's ability to explore the ocean.

Saving the best for last, I wish to thank my wife Sarah. I could not have devoted the time and energy necessary to complete this work without her love, sacrifice, support, and longsuffering.

# TABLE OF CONTENTS

<b>ACKNOWLEDGMENTS</b> . . . . .	iv
<b>LIST OF TABLES</b> . . . . .	ix
<b>LIST OF FIGURES</b> . . . . .	x
<b>LIST OF SYMBOLS</b> . . . . .	xvii
<b>CHAPTER 1 INTRODUCTION</b> . . . . .	1
1.1 Overview and Organization . . . . .	1
1.2 Real Aperture Sidescan Sonar . . . . .	2
1.2.1 Range Resolution . . . . .	4
1.2.2 Cross-Range Resolution . . . . .	5
1.2.3 Area Coverage Rate . . . . .	6
1.2.4 Frequency . . . . .	7
1.3 Sidescan Synthetic Aperture Sonar . . . . .	8
1.3.1 Stripmap-Mode Synthetic Aperture Data Collection . . . . .	9
1.3.2 Range Resolution . . . . .	11
1.3.3 Cross-Range Resolution . . . . .	11
1.3.4 Area Coverage Rate . . . . .	12
1.3.5 Frequency . . . . .	13
1.4 Vernier Sonar Arrays . . . . .	15
1.5 Practical Concerns for Successful SAS Imaging . . . . .	16
1.6 Display of Synthetic Aperture Sonar Imagery . . . . .	18
1.7 Comments . . . . .	23
<b>CHAPTER 2 OVERVIEW OF SYNTHETIC APERTURE SONAR IMAGE RE- CONSTRUCTION</b> . . . . .	25
2.1 Introduction . . . . .	25
2.2 Coordinate System . . . . .	25
2.3 The Wave Equation and the Green's Function . . . . .	25
2.3.1 Time-Independent Green's Function . . . . .	26
2.3.2 Time-Dependent Green's Function . . . . .	28
2.4 Radiation from a Rectangular Source . . . . .	29
2.5 Modeling the Reflection from the Sea Floor . . . . .	31
2.6 Spatiotemporal Domain Image Reconstruction . . . . .	33
2.7 Fourier-Based Image Reconstruction . . . . .	35
2.8 Comments . . . . .	38
2.8.1 Reconstruction Expense . . . . .	38
2.8.2 Spatial Sampling Concerns . . . . .	40

<b>CHAPTER 3</b>	<b>EFFECTS OF UNCOMPENSATED PLATFORM MOTION . .</b>	<b>41</b>
3.1	Introduction . . . . .	41
3.2	Origin of Stripmap Phase Error Effects . . . . .	42
3.3	Phase Error Analysis . . . . .	44
3.3.1	The Principle of Stationary Phase . . . . .	45
3.3.2	Linear Phase Error . . . . .	45
3.3.3	Quadratic Phase Error . . . . .	46
3.3.4	Cubic Phase Error . . . . .	50
3.3.5	Sinusoidal Phase Error . . . . .	51
3.3.6	Sawtooth Phase Error . . . . .	53
3.3.7	Random Normally-Distributed Phase Error . . . . .	57
3.4	Results and Discussion . . . . .	62
3.4.1	Stripmap vs. Spotlight Phase Errors . . . . .	64
3.4.2	Spatial Frequency Windowing Effects . . . . .	67
3.4.3	Moving Scatterers . . . . .	67
3.5	Comments . . . . .	69
<b>CHAPTER 4</b>	<b>THE REDUNDANT PHASE CENTER TECHNIQUE . . . . .</b>	<b>70</b>
4.1	Introduction . . . . .	70
4.2	Modeling the RPC Time Delays . . . . .	71
4.3	Short-Time Correlation-Based Time Delay Estimation . . . . .	75
4.3.1	Algorithms for Coarse Time Delay Estimation . . . . .	82
4.4	Comments . . . . .	84
<b>CHAPTER 5</b>	<b>MOTION ESTIMATION TECHNIQUES . . . . .</b>	<b>87</b>
5.1	Introduction . . . . .	87
5.2	Defining the Nominal Sensor Trajectory . . . . .	87
5.3	IMU-Based Motion Estimation . . . . .	89
5.4	Slant Plane Redundant Phase Centers . . . . .	90
5.5	Nonlinear Least Squares . . . . .	94
5.6	Simplified Nonlinear Least Squares . . . . .	99
5.7	Dual-Sided Closed-Form Solution . . . . .	104
5.8	Estimating Platform Surge Using RPC . . . . .	111
5.8.1	Array Correlation Matrix Analysis . . . . .	111
5.8.2	Exhaustive Surge Error Estimation . . . . .	113
5.8.3	Surge Estimation via Eigendecomposition . . . . .	115
5.9	Non-RPC Data-Driven Motion Estimation Techniques . . . . .	118
5.10	Comments . . . . .	118
<b>CHAPTER 6</b>	<b>MOTION COMPENSATION . . . . .</b>	<b>120</b>
6.1	Introduction . . . . .	120
6.2	The Range-Varying Nature of Motion Compensation . . . . .	120
6.3	Range Compensation Using a Single Delay . . . . .	122
6.4	Range Compensation Using Variable Delays . . . . .	123
6.5	Crab Compensation . . . . .	123

6.6	Compensation for Angular Motion . . . . .	125
6.7	Cross-Range Compensation . . . . .	126
6.8	Examples of Range-Varying Motion Compensation . . . . .	126
6.9	Comments . . . . .	129
<b>CHAPTER 7</b>	<b>CONCLUSION . . . . .</b>	<b>132</b>
7.1	Example SAS Imagery . . . . .	133
7.2	Approaches to Synthetic Aperture Imaging . . . . .	135
7.3	Summary of Contributions . . . . .	138
7.4	Final Comments . . . . .	139
<b>REFERENCES</b>	<b>. . . . .</b>	<b>141</b>

## LIST OF TABLES

Table 2.1	Conventions used to describe platform motion. The terms for unwanted motions are shown in italics. . . . .	26
Table 4.1	Symbols used to describe the delay observed by a pair of overlapping receivers used by the redundant phase center technique. . . . .	71
Table 5.1	Intended RPC channel pairing for an $N$ -channel array with $M = N - \tilde{N}$ overlaps. . . . .	113

## LIST OF FIGURES

Figure 1.1	The diagram above shows the important features of sidescan sonar data collection. The most common way to image a large region of the sea floor is to tow or drive the sonar in a serpentine pattern similar to the pattern used for mowing a lawn. The sonar thus images strips of the sea floor that may be used to form a composite image. The range of the sonar determines the allowable pitch or spacing between successive legs, $P_{\text{leg}}$ . Also shown is the blind region, known as the holiday, that exists around the nadir point. For complete coverage, the holiday must be imaged by adjusting $P_{\text{leg}}$ or by using a special downward-looking sensor known as a <i>gap filler</i> . . . . .	3
Figure 1.2	The figure above contains both port and starboard sides of a sidescan SAS image. The sandy sea floor transitions from a smooth to a rippled surface rather abruptly. The holiday is visible as the region extending from -5 m to +5 m range. . . . .	4
Figure 1.3	SAS data collection is similar to the raster-style collection shown in Figure 1.1. An important difference is that SAS relies on wide physical beams to insonify the sea floor for many pings which are then coherently processed to form a synthetic aperture having a virtual beam whose cross-range footprint is independent of range. . . . .	9
Figure 1.4	Simulated SAS data and reconstruction. The upper figure (a) shows the hyperbolic signal histories that result from observing a line of ideal point scatterers. The nulls of the sensor's directivity pattern are clearly visible, and the main lobe of the most distant hyperbola just fills the cross-range extent of the figure. At any range, the length of the synthetic aperture is generally taken to be equal to the -3 dB width of the main lobe. The lower figure (b) demonstrates that the resulting image resolution is independent of range. . . . .	10
Figure 1.5	Attenuation of sound in sea water expressed in decibels per meter. The dashed lines represent the three individual phenomena contributing to the overall absorption loss shown by the solid line. (After Clay and Medwin, [3]) . . . . .	14
Figure 1.6	The diagram above illustrates the phase center concept. At long ranges, a bistatic source and receiver pair behaves as a single equivalent sensor that is located midway between the physical sensor locations. The location of this virtual sensor is called the phase center of the source/receiver pair. The signal observed by the phase center is assumed to be the same as that observed by the physical pair of sensors, with the exception of a delay equal to $\tau_{\text{PC}} = (R_{\text{TX}} + R_{\text{RX}} - 2R_{\text{PC}})/c$ . . . . .	16

Figure 1.7	A vernier array is constructed using the phase center concept illustrated in Figure 1.6. The apparent length of the resulting phase center array (or virtual array) is exactly half that of the physical array. Consequently, the advance per ping is determined by the length of the phase center array. The diagram above shows this reasoning applied to an array of six receivers, with the trailing two elements overlapping for the purpose of motion estimation using redundant phase centers. . . .	17
Figure 1.8	The photo above depicts the Small Synthetic Aperture Minehunter system developed for the Office of Naval Research by the Naval Surface Warfare Center in Panama City, Florida [18]. The SAS (the white section) was constructed by the Applied Research Laboratory of the Pennsylvania State University. The vehicle is a Remus 600 supplied by the Woods Hole Oceanographic Institution. (Photo courtesy of Mr. Daniel Brown.) . . . . .	18
Figure 1.9	The upper image (a) shows an example SAS image displayed without any form of normalization. A small number of bright pixels forces the majority of the image to be too dark. The lower image (b) is the result of applying the normalization procedure described in this section to the upper image. . . . .	20
Figure 1.10	The upper plot (a) shows the gamma distribution PDF estimated using the pixel magnitude values from Figure 1.9(a). Also shown is the estimate of the image magnitude PDF computed using the histogram. The horizontal axis is truncated to emphasize the shape of the PDF. The lower plot (b) shows the histogram after modification using the algorithm described in this section. The horizontal axis in this plot is not truncated. The gamma PDF from the upper plot is included for comparison, and its values are given by the right-hand vertical axis. Note the large concentration of dark pixels relative to bright prior to adjusting the histogram. . . . .	21
Figure 2.1	Right-handed coordinate system used for motion estimation and compensation. The array yaw is referenced to the positive $y$ -axis, as this is the direction of forward motion. . . . .	27
Figure 2.2	Rectangular distribution of point sources used to model a sonar projector. The source is located at the origin and lies in the $y$ - $z$ plane. The width and height are $D$ and $H$ , respectively. . . . .	29
Figure 3.1	(a) Comparison of the actual quadratic phase error (with $v x_0/D = 0.88$ ) spectrum $G(k_y)$ and the approximation given by (3.8) which was derived using the principle of stationary phase. (b) Example of an ideal PSR and the PSR corrupted by a quadratic phase error. . . . .	48

Figure 3.2	(a) Comparison of the actual cubic phase error spectrum spectrum $G(k_y)$ and the approximation given by (3.18) which was derived using the principle of stationary phase. (b) Example of an ideal PSR and the PSR corrupted by a cubic phase error with $\lambda/D = 0.31$ and $\rho L_{SA}^2 = 0.035$ . . . . .	52
Figure 3.3	(a) Comparison of the actual sinusoidal phase error spectrum spectrum $G(k_y)$ and the analytical expression given by (3.22). Note that the actual (or experimental) $G(k_y)$ was computed discretely, and it therefore does not consist of delta functions as does the continuous result. (b) Example of an ideal PSR and the PSR corrupted by a sinusoidal phase error with $k_0\beta = 0.63$ and $\gamma/B_{k_y} = 0.069$ . . . . .	54
Figure 3.4	(a) Comparison of the actual sawtooth phase error spectrum spectrum $G(k_y)$ and the analytical expression given by (3.25). For this example, the parameters are $D = 0.4$ , $\theta_{crab} = 2^\circ$ , and $N = 6$ . (b) Example of an ideal PSR and the PSR corrupted by a sawtooth phase error with the same parameters given in Figure 3.4(a). . . . .	56
Figure 3.5	Expected value of the real and imaginary components of (3.27) evaluated for values of $\sigma_\kappa^2$ ranging from 0.01 to 100. . . . .	58
Figure 3.6	(a) Example of an ideal PSR and the PSR corrupted by a normally-distributed white random phase error whose standard deviation is one quarter of a wavelength. (b) Surface plot depicting the ratio of the actual peak to background ratio (in decibels) as a function of $\mu_g$ and the synthetic aperture length. . . . .	61
Figure 3.7	Examples of the effects of the various types of phase errors discussed in Section 3.3: (a) no error; (b) quadratic error; (c) cubic error; (d) sinusoidal error; (e) sawtooth error; (f) random error. . . . .	63
Figure 3.8	The plots above show the result of applying the single-frequency results to a finite-bandwidth signal. Shown are the sinusoidal (a) and sawtooth (b) phase errors for a signal with a center frequency to bandwidth ratio of four. These plots are the finite-bandwidth equivalent of Figures 3.3(b) and 3.4(b) where the ideal PSR is omitted and the horizontal axis is enlarged by a factor of two. . . . .	65
Figure 3.9	Synthetic aperture PSR as influenced by various degrees of quadratic phase error. The upper plot has no weighting applied, while the lower was created using a Hamming window. The horizontal axis is normalized length, where the interval $[-0.5, 0.5]$ corresponds to the ideal cross range resolution, $D/2$ . Also, each PSR is displayed relative to its own peak value in order to better illustrate the relative -3 dB widths of the mainlobes. . . . .	66



Figure 4.1	Schematic of two consecutive SAS pings with the array in motion. The physical arrays are shown. A similar diagram of the phase center arrays would show the expected overlap of elements 10 and 11. The transmitter is colocated with the fifth element of the receiver array. . . . .	72
Figure 4.2	Flowchart of the time delay estimation process. Signals 1 and 2 are typically much longer than the window length. The process shown in the diagram is repeated many times, and the output is a vector of time delays, one for each position of the local short-time window. . . . .	77
Figure 4.3	(a) Time delays estimated using RPC (diamonds) plotted along with the closed-form description of the RPC delay (solid line) found using the nonlinear least squares technique of Section 5.6. The RPC delays shown correspond to $\rho \geq 0.85$ . Delays with $\rho < 0.85$ are not shown and were rejected for the least-squares computation. (b) The lower plot shows the result of subtracting the least-squares delay from the RPC observations (black diamonds, left-hand scale). Also shown are the correlation coefficients for the RPC delays (circles, right-hand scale). Note that all delays are expressed as distances. . . . .	80
Figure 4.4	RPC correlation map (a) and the corresponding SAS image (b). The dark holes in the correlation map are caused by schools of fish that are visible as blurred patches in the image. Also note the bright wishbone-shaped object in the image. The strong scattering from this object causes an increase in the local RPC correlation between 45 m and 50 m range. . . . .	85
Figure 4.5	Cramér-Rao lower bound on the time delay estimates (4.19) for a SAS with a 30 kHz bandwidth and a center frequency of 120 kHz. The curves represent correlation coefficients ranging from 0.5 (upper curve) to 0.9 (lower curve) in increments of 0.1. . . . .	86
Figure 5.1	The nominal trajectory (NT) is defined on a scene-by-scene basis. The y-axis and the yaw relative to it are defined by the mean track over the sea floor. . . . .	88
Figure 5.2	The diagram above, after Pinto and Bellettini [62], shows the DPCA sway and yaw for two consecutive pings. The shaded box indicates the overlapping portion of the SAS phase center array used for motion estimation. The DPCA sway, $\gamma$ , is referenced to the middle of this region. . . . .	92
Figure 5.3	The diagram above, after Pinto and Bellettini [62], shows the trajectory of the SAS phase center array in the slant (or image) plane. The first array is shown as being parallel to the x-axis to indicate that the integrated trajectory estimate given by (5.10) is referenced to the first ping. . . . .	93

Figure 5.4	The example above demonstrates the effectiveness of the nonlinear least squares (NLLS) motion estimation. The upper image (a) was created without motion compensation, and the lower image (b) was compensated using the NLLS estimate. The scene depicts scattered rocks and a pair of lobster traps on the sea floor. A rope is visible connected to the trap in the lower right quadrant of the image. . . . .	98
Figure 5.5	The left diagram (a) shows two consecutive pings of a SAS array with two RPC channel. The figure on the right (b) shows a cross-section in the $x$ - $z$ plane of a pair of redundant phase centers. When the stop-and-hop assumption is made, $\Delta x$ and $\Delta z$ do not vary during the ping reception time. . . . .	100
Figure 5.6	The plot above demonstrates the typical performance of the nonlinear least squares technique. The broken line shows the measured RPC time delays as a function of range, and diamonds indicate the points used to estimate $[\Delta x \ \Delta z]^T$ via nonlinear least squares. These points correspond to delays for which the correlation coefficient was equal to, or greater than, 0.9. The solid line shows the time delay as a function of range as computed using the $[\Delta x \ \Delta z]^T$ estimated by nonlinear least squares. . . . .	101
Figure 5.7	The inter-ping heave ( $\Delta z$ ) estimate is integrated to yield an estimate of the altitude history for a segment of SAS12 data (solid line). The broken line shows the altitude measured by the on-board navigation system. . . . .	102
Figure 5.8	SAS image before (a) and after (b) motion compensation via simplified nonlinear least squares. The scene depicts the edge of a region of sand ripples on the sea floor. (Figure reproduced from [77].) . . . . .	103
Figure 5.9	Linear regression used to estimate the sway, heave, yaw, and pitch using the $[\Delta x \ \Delta z]^T$ estimates for ten pairs of redundant phase centers. The data shown in the plot contains an outlier which was excluded in the computation. The circles represent the points $[\Delta x \ \Delta z]^T$ in 3D space, while the crosses indicate the projections into the $x$ - $y$ and $y$ - $z$ planes. . . . .	104
Figure 5.10	Vehicle yaw (a) and pitch (b) estimated using the simplified nonlinear least squares RPC technique. The attitude estimates from the AUV's inertial measurement unit (IMU) are shown for reference. Both estimates are unsmoothed. . . . .	105
Figure 5.11	SAS image before (a) and after (b) motion compensation showing a trap lying on the bottom of Boston harbor in Massachusetts. . . . .	106

Figure 5.12	The diagram above illustrates the process of computing the time-delay correction needed to map a physical array of receivers into an array of phase centers located along the centerline of the AUV. . . . .	107
Figure 5.13	A dual-sided RPC configuration can be used to derive a closed-form solution to the problem of estimating the ping-to-ping sway and heave.	108
Figure 5.14	The dual-sided closed-form solution was used to compensate the image above. The data was collected by the SAS21 system sponsored by the Office of Naval Research [79, 80]. The uncompensated image (a) is on the left, and the compensated image (b) is on the right. The image shows a small depression in a sandy sea floor that is otherwise covered with sand ripples. Such holes have a number of causes, and a notable example is the crater feeding behavior of dolphins [81]. (Figure reproduced from [80].) . . . . .	110
Figure 5.15	Comparison of the speed over ground (SOG) estimated computed using the exhaustive correlation and eigendecomposition methods. The speed estimate computed by the AUV's on-board navigation system is also shown for reference. . . . .	117
Figure 6.1	The upper figure (a) shows the geometry used to describe the range varying time delay correction. The drawing is not to scale in the sense that a far-field assumption is made in which $ \delta  \ll  \mathbf{R} $ implying $\mathbf{R}$ is nearly parallel to $\mathbf{R}'$ and thus $\mathbf{R} \approx \mathbf{R}' + \delta \cdot \mathbf{e}$ . This simplification leads to the idea that the compensation delay for a particular range is given by the projection of $\delta$ onto a given $\mathbf{R}$ . Using this idea, the total variation over the range swath of the time delay correction is found to be proportional to $\delta \cdot (\mathbf{e}_{\text{far}} - \mathbf{e}_{\text{near}})$ as shown in the lower figure (b) (in which $\mathbf{R}'$ is omitted). The maximum variation occurs when $\delta \parallel (\mathbf{e}_{\text{far}} - \mathbf{e}_{\text{near}})$ .	121
Figure 6.2	Crab occurs when an AUV is forced to place its angle of attack into a cross current in order to maintain a desired track over the sea floor. Crab compensation can be performed in two ways. The first is to beam steer the physical arrays by the angle $\theta_{\text{crab}}$ in order to align them with the nominal trajectory (NT) defined by the vehicle's track over the sea floor. The second approach is to define another compensation baseline $NT_{\text{sway}}$ making an angle $\theta_{\text{crab}}$ with the vehicle track and then treat the crab as sway. . . . .	124
Figure 6.3	Helical motion with a diameter of 0.25 m at a mean altitude of 7 m. . . . .	127

Figure 6.4	The upper figure (a) shows the residual motion compensation error visualized as a surface. Note that the residual error at a range of 45 m is zero, because this is the range at which the bulk compensation was performed. The lower plot (b) shows the residual error at ranges of 15 m and 80 m. To analyze the effect of the residual compensation error on the PSR, the error must be considered only over the spatial extent of the synthetic aperture. . . . .	128
Figure 6.5	Short ranges are especially susceptible to the defocus caused by compensating for the motion using a single delay for all ranges. The top image (a) shows an example of this effect, while the bottom image (b) resulted from applying a range-varying compensation scheme. The dark vertical band appearing at a range of 13 m is the first null of the vertical beam pattern. . . . .	130
Figure 6.6	A taxonomy of motion compensation schemes. . . . .	131
Figure 7.1	The photos above show a pair of World War II vintage aircraft. The left photo (a) depicts a Fairey Firefly, a British carrier-based fighter. The right photo (b) shows a Curtiss SB2C Helldiver, an American dive bomber. The Helldiver was responsible for sinking more tonnage of shipping than any other aircraft during World War II. SAS imagery of these aircraft is shown in Figure 7.2. . . . .	133
Figure 7.2	The top image (a) shows a SAS image of a Fairey Firefly lying on the sea floor in Jervis Bay, Australia. It crashed as a result of a midair collision on November 27, 1956. Today, the wreckage is popular with scuba divers. The bottom image (b) is of a Curtiss SB2C Helldiver in Buzzards Bay, Massachusetts. It was ditched by Naval reservist Ensign John L. Hagerman after the aircraft's engine lost oil pressure and stalled during a training mission on July 7, 1947. . . . .	134
Figure 7.3	The image above shows the World War II era tugboat <i>Vittoria</i> sitting upright on the sea floor off the coast of Italy. . . . .	135
Figure 7.4	The image above shows a pair of underwater pipelines off the coast of Italy. The larger of the two is clearly not functional because many of its segments have become disconnected. . . . .	136

## LIST OF SYMBOLS

$\alpha$	Azimuthal angle, page 30
$\alpha$	Constant of proportionality of a linear phase error across a synthetic aperture, page 45
$*_t$	Convolution with respect to time, page 31
$*_{k_y}$	Convolution with respect to spatial frequency $k_y$ , page 44
$\beta$	Constant of proportionality describing the amplitude of a sinusoidal phase error across a synthetic aperture, page 51
$\beta$	Elevation angle, page 30
$\chi$	Fractional error in the estimate of the forward velocity, page 49
$\Delta\tau$	Time delay observed by a redundant phase center pair of receivers, page 73
$\Delta k$	Change in wavenumber due to a change in propagation speed $\Delta c$ , page 47
$\delta$	Dirac delta function, page 33
$\delta$	Kronecker delta function, page 96
$\delta_x$	Range resolution, page 4
$\delta_y$	Cross-range (along-track) resolution, page 5
$\epsilon$	Unwanted platform motion, page 43
$\epsilon_R$	Motion compensation error for scatterers off boresight, page 131
$\gamma$	Constant of proportionality describing the frequency of a sinusoidal phase error across a synthetic aperture, page 51

$\gamma$	Differential (ping-to-ping) sway measured in the slant place using one or more redundant phase centers, page 92
$\kappa$	Normally-distributed random phase error across a synthetic aperture with variance $\sigma_\kappa^2$ , page 57
$\lambda$	Wavelength, page 5
$\mathcal{F}$	Fourier transform operator, page 83
$\Phi$	Rotation matrix, page 71
$\Psi$	Rotation matrix, page 73
$\mu_g$	Mean of random phase error across a synthetic aperture, page 57
$\nabla$	Gradient operator, $\nabla = \partial/\partial x + \partial/\partial y + \partial/\partial z$ , page 26
$\nabla^2 D$	Hessian matrix of $D$ , page 96
$\nu$	Constant of proportionality of a quadratic phase error across a synthetic aperture, page 47
$\nu$	Signal-to-noise ratio, page 79
$\odot_y$	Correlation with respect to the cross-range variable $y$ , page 43
$\omega$	Radian frequency, page 26
$\phi$	Roll (rotation about the $y$ axis), page 25
$\Psi$	Amplitude of a field variable described by the wave equation, page 26
$\psi$	Differential (ping-to-ping) yaw measured in the slant plane using two or more redundant phase centers, page 92
$\psi$	Pitch (rotation about the $x$ axis), page 25

$\rho$	Constant of proportionality of a cubic phase error across a synthetic aperture, page 50
$\rho$	Correlation coefficient, page 76
$\rho$	Range, page 36
$\tau^n$	Time from ping transmission to reception by the $n^{\text{th}}$ array element (that is, the time of flight), page 71
$\tau_{\text{PC}}$	Time delay necessary to make a bistatic transmit/receive pair appear as a monostatic virtual element located at its phase center, page 15
$\theta$	Yaw (rotation about the $z$ axis), page 25
$\theta_{3\text{dB}}$	Angle between the -3 dB points of the transducer directivity pattern, page 5
$\tilde{N}$	Forward advance between pings expressed as a multiple of the spacing between array phase centers, page 112
$\xi$	Sawtooth function, page 53
$A$	Amplitude of a sawtooth phase error across a synthetic aperture, page 53
$A$	Envelope of a complex-valued acoustic signal, page 29
$A_c$	Autocorrelation of a transmitted signal, page 32
$B_w$	Temporal bandwidth, page 4
$B_{k_y}$	Spatial bandwidth, page 34
$c$	Acoustic wave propagation speed (approximately 1500 m/s in sea water), page 4
$D$	Horizontal extent of a piston transducer, page 5

$D$	Residual function representing the error between the true and estimated values of a quantity, page 95
$D_{\text{ping}}$	Forward distance traveled between pings, page 6
$E$	Expected value operator, page 59
$E$	Received echo signal, page 31
$E_c$	Pulse-compressed received echo signal, page 32
$f$	Amplitude of a radiating acoustic source, page 28
$f$	Cross-range synthetic aperture point response function, page 44
$f$	Reflectivity of a point scatterer, page 31
$G$	Green's function, page 28
$g$	Phase error due to the unwanted motion $\epsilon$ , page 43
$H$	Vertical extent of a piston transducer, page 28
$H_0^{(2)}$	Zeroth-order Hankel function of the second kind, page 36
$I$	Image magnitude, page 19
$J$	Jacobian matrix, page 95
$K$	Number of statistically independent receivers in an array, page 92
$k$	Wavenumber, or spatial frequency ( $k = \omega/c$ ), page 26
$L$	Curve in three-dimensional space, page 34
$L_{\text{SA}}$	Synthetic aperture length, page 43
$L_{\text{win}}$	Spatial extent of correlation window $T$ , page 79



$M$	Number of receivers in the array used as redundant phase centers, page 113
$N$	Number of receivers (hydrophones) in an array, page 15
$P$	Period of a sawtooth phase error across a synthetic aperture, page 53
$P$	Signal emitted from an acoustic point source, page 29
$p$	Ping index, page 70
$p$	Transmitted temporal signal, page 76
$P_{\text{leg}}$	Pitch (spacing) between parallel tracks used for mapping the sea floor, page 3
$R$	Correlation matrix, page 116
$R$	Range, page 5
$s$	Cross-range synthetic aperture signal history, page 43
$s$	Received temporal signal, page 76
$S^{-1}$	Inverse Stolt mapping, $k_x = \sqrt{4k^2 - k_y^2}$ , page 37
$T$	Magnitude transition point used in altering the histogram of an image, page 22
$T$	Temporal duration of correlation window, page 79
$t$	Time, page 26
$v$	Scalar velocity, page 6
$W$	Spatiotemporal response of a transducer, page 31
$W$	Weighting matrix, page 97
$x$	Range, page 9
$y$	Cross-range (along-track), page 9

$z$	Altitude, page 9
$\delta$	Vector pointing from the nominal platform trajectory to the actual location of a given phase center, page 120
$\omega$	Angular velocity vector, $\omega = [\dot{\psi} \ \dot{\phi} \ \dot{\theta}]^T$ , page 71
$C$	Correlation matrix, page 112
$d$	Position of the phase center corresponding to a given projector/receiver pair in an array, page 101
$d_n$	Vector from vehicle center to the $n^{\text{th}}$ array element, page 71
$e$	Unit vector, page 67
$p$	Projection of a given array phase center location onto the centerline of the host vehicle, page 108
$R_n$	Vector at time $t$ from the $n^{\text{th}}$ array element to a given scatterer, page 71
$r_n$	Unit vector in the direction of $R_n(t)$ , page 71
$S$	Matrix whose columns are the observed signals from adjacent receivers, page 115
$s$	Scatterer location, page 71
$v$	Velocity vector, page 67
$w$	Vector of weights used to estimate cross-track motion, page 97
$w$	Vector of weights used to estimate platform surge, page 116
$x$	Position in three dimensional $(x, y, z)$ space, page 26
$x$	Vector representing the ping-to-ping displacement in the $x$ - $z$ plane resulting in the time delay observed by a redundant phase center pair, page 99

$\mathcal{L}$	Helmholtz operator, $\mathcal{L} = (\nabla^2 + k^2)$ , page 26
AASR	Along-track ambiguity to signal ratio, page 13
ACR	Area coverage rate, page 6
APP	Advance per ping, page 111
AUV	Autonomous underwater vehicle, page 1
CPA	Closest point of approach, page 9
DPCA	Displaced phase center antenna, page 90
DVL	Doppler velocity log, page 17
GPS	Global positioning system, page 17
IMU	Inertial measurement unit, page 18
LFM	Linear frequency modulation, page 5
LORAN	Long range navigation, page 17
MTI	Moving target indication, page 90
NLLS	Nonlinear least squares, page 94
NT	Nominal trajectory, page 87
PDF	Probability density function, page 19
PGLR	Peak to grating lobe ratio, page 13
PSP	Principle of stationary phase, page 45
PSR	Point scatterer response, page 4
RAS	Real aperture sonar, page 8

RPC	Redundant phase center, page 2
SAR	Synthetic aperture radar, page 8
SAS	Synthetic aperture sonar, page 1
SOG	Speed over ground, page 17
SPRPC	Slant plane redundant phase center, page 90

# CHAPTER 1

## INTRODUCTION

### 1.1 Overview and Organization

The purpose of this thesis is to provide the reader with a fundamental understanding of the problem of estimating and compensating for the unwanted motion experienced by a synthetic aperture sonar (SAS). Several techniques are presented for solving this problem, and an emphasis is placed on the practical aspects of implementing them. The schemes presented have their relative merits; some are arguably better than others for a given application. They have been proven to be effective by testing them extensively on field data acquired by SAS systems carried by autonomous underwater vehicles (AUVs). Some of the algorithms described in this thesis were used successfully in the world's first demonstration of fully-autonomous real-time SAS imaging software carried aboard an AUV.

This document is organized as follows:

- Chapter 1 provides an overview of the field of synthetic aperture sonar imaging. The purpose of this chapter is to give some necessary background and to provide the context in which the body of this work rests. The fundamental results regarding SAS imaging are given, but most are not derived. Rather, qualitative arguments are given along with references to more complete discussions.
- Chapter 2 gives an introduction to the two dominant types of SAS reconstruction techniques: the spatiotemporal and Fourier-domain approaches. Both are derived in detail, and care is taken to point out the important assumptions and approximations found in these derivations. This chapter may be skipped by readers already well-versed in SAS imaging.
- Chapter 3 presents an analysis of the ways in which various categories of platform motion can degrade SAS imagery.


- Chapter 4 describes the redundant phase center (RPC) technique for SAS motion estimation, which is the foundation for much of Chapter 5.
- Chapter 5 is the focus of this thesis, as it discusses the problem of using the RPC data to measure the motion of the platform carrying the sonar. Several techniques are presented, and their relative merits are discussed.
- Chapter 6 discusses several ways to compensate the SAS data once the array motion has been estimated. Of particular interest are the circumstances under which certain simplifications can be employed that result in lower computational expense.
- Chapter 7 presents sample results and contains some final comments on the work as well as suggestions for building upon it.

## **1.2 Real Aperture Sidescan Sonar**

Before introducing the idea of synthetic aperture sonar, it is helpful to first describe real aperture sidescan sonar imaging. A sidescan sonar consists of a transducer moving in a nominally straight line through the water. At regular time intervals, the sensor transmits a signal and records the received echoes. These echoes are then stacked sequentially to construct a raster scan image of the sea floor as depicted in Figures 1.1 and 1.2. The sensor used for sidescan imaging is frequently towed, with the tow cable being used for data uplink. The sidescan sonar might also be mounted to a ship's hull, and the application most relevant to the present work is sidescan sonar carried by an autonomous underwater vehicle.

The operating characteristics of a sidescan sonar can be summarized using the following quantities:


- **Range Resolution:** The range resolution is often described as the distance between the -3 dB points of the main lobe of the transmitted signal as measured along the



**Figure 1.1.** The diagram above shows the important features of sidescan sonar data collection. The most common way to image a large region of the sea floor is to tow or drive the sonar in a serpentine pattern similar to the pattern used for mowing a lawn. The sonar thus images strips of the sea floor that may be used to form a composite image. The range of the sonar determines the allowable pitch or spacing between successive legs,  $P_{leg}$ . Also shown is the blind region, known as the holiday, that exists around the nadir point. For complete coverage, the holiday must be imaged by adjusting  $P_{leg}$  or by using a special downward-looking sensor known as a *gap filler*.

range axis. If the transmitted signal requires compression, or matched filtering, then the resolution refers to the compressed signal.

- **Cross-Range Resolution:** The cross-range resolution is given by the distance between the -3 dB points of the main lobe of the transmitted signal as measured along the cross-range axis.
- **Area Coverage Rate:** The area coverage rate indicates the amount of area per unit time that may be imaged with a given sonar.
- **Frequency:** The frequency of a sonar gives a rough indication of its maximum range



**Figure 1.2.** The figure above contains both port and starboard sides of a sidescan SAS image. The sandy sea floor transitions from a smooth to a rippled surface rather abruptly. The holiday is visible as the region extending from -5 m to +5 m range.

capability, as the attenuation of sound in the water increases with frequency. The frequency also indicates the acoustic wavelength which, in turn, indicates the lower limit on the size of objects that may be imaged.

The range and cross-range resolution give the user an idea of the sonar's point scatterer response (PSR). That is to say, an ideal point scatterer (with no finite spatial extent) will appear to have the dimensions given by the range and cross-range resolution. The following sections give a brief discussion of each of the parameters listed above. For a detailed survey of real aperture imaging sonar, see the texts by Fish and Carr [1, 2].

### 1.2.1 Range Resolution

The range resolution is simply proportional to the inverse of the bandwidth,  $B_w$ , of the transmitted signal:  $\delta x = c/2B_w$ . These signals might be tone-bursts or they might be



extended-time signals such as linearly-swept frequency modulated (LFM) pulses. A tone-burst requires no added processing, but the signal amplitude is limited by the power that can be put through the projector. This limit might be imposed by the hardware itself through the electronics or the transducer, or by nature in the form of cavitation. Cavitation occurs when the pressure of the rarefaction portion of the acoustic wave drops below the vapor pressure of the surrounding water, effectively boiling it [3].

Extended-time signals such as the LFM require an additional step known as pulse compression or matched filtering [3, 4, 5] in order to achieve the desired resolution. The fact that these signals are extended in time allows the instantaneous power to be reduced. The pulse compression step yields a high signal-to-noise ratio that might otherwise be physically impossible to obtain using a tone burst.

### 1.2.2 Cross-Range Resolution

The cross-range resolution (also referred to as azimuth resolution or along-track resolution) of a real aperture sonar is given by the beamwidth of the transducer. The exact beam pattern of a rectangular sensor of length  $D$  is derived in Section 2.4. For the present, it suffices to use the approximate formula  $\theta_{3\text{dB}} \approx \lambda/D$  where  $\lambda$  is the wavelength. The approximation worsens for large beamwidths, but it is sufficiently accurate for almost all practical use. The cross-range resolution is derived by idealizing the sensor directivity pattern as being equal to zero outside the -3 dB beamwidth and constant within it. Thus, any scatterer whose angle relative to the sensor boresight is greater than  $\pm\theta_{3\text{dB}}/2$  will be invisible. Furthermore, if one or more scatterers (all assumed to be at the same range) fall within the -3 dB beam, they are indistinguishable as separate objects and appear to have cross-range extent equal to  $\delta_y = R\theta_{3\text{dB}}$ . In other words, the cross-range resolution is given by the arc length subtending the sonar's -3 dB beamwidth at a given range  $R$ . Thus, the cross-range resolution of a real aperture sonar worsens linearly with range.

In order to achieve high resolution at long ranges,  $\theta_{3\text{dB}}$  must be reduced. This reduction can be accomplished in two ways. First, the frequency can be increased resulting in

shorter wavelengths. The frequency cannot be increased arbitrarily, however, as absorption eventually limits the useful range. Furthermore, the frequency range may be constrained in some instances by the desire to avoid interfering with other acoustic devices. The second way to improve resolution is to increase  $D$ . But again,  $D$  cannot be arbitrarily large as the physical size of a transducer is often limited by cost, space and power requirements, and manufacturing technology.

### 1.2.3 Area Coverage Rate

Sidescan sonar is often used to survey an area completely or to locate bottom features within a given region. Hence, the area coverage rate (ACR) is an important metric when designing a sonar or when choosing one for use in a given application. These surveys are generally carried out by driving the sonar back and forth across the desired area in a fashion reminiscent of the pattern used when mowing grass (see Figure 1.1). The area coverage rate is the product of the range swath width,  $(R_{\max} - R_{\min})$ , and the sonar's forward velocity,  $v$ . For example, a sonar with a 30 m range swath moving at 1.5 m/s has an ACR of 45 m<sup>2</sup>/sec. Area coverage rate is most commonly expressed as the area imaged per hour, in which case the example yields an ACR of 0.16 km<sup>2</sup>/hr.

It is common to neglect the gap, or invisible region, occurring around nadir implying that  $R_{\min} = 0$ . In this case, determining the ACR amounts to determining the maximum range and forward velocity of the sonar. The maximum range is limited by acoustic attenuation and by the distance the sound is allowed to propagate before reception is terminated and the next ping transmitted. The propagation time is, in turn, driven by the desired spacing of the pings (that is, the cross-range spacing of the lines making up the raster scan image). If the sonar ping spacing is  $D_{\text{ping}}$  then the maximum range is:

$$R_{\max} = \frac{D_{\text{ping}} c}{2v}. \quad (1.1)$$

This equation says that the ping spacing divided by the forward velocity gives the time between pings. The ping time is then multiplied by the sound speed to give the distance

covered by the sound propagation. Finally, this distance is divided by two to account for the fact that the sound has to travel out and back from scatterers in the water. The maximum range is inversely proportional to the forward velocity. Reducing  $v$  increases  $R_{\max}$ , but there is often a practical lower limit on how slowly a body can move through the water and remain stable enough to produce an acceptable image. The ACR is obtained by multiplying  $R_{\max}$  in (1.1) by the forward velocity,  $v$ :

$$ACR_{\text{RAS}} = R_{\max} v = \frac{D_{\text{ping}} c}{2}. \quad (1.2)$$

Equation (1.2) seems to say that the ACR can be increased easily by widening the distance between pings. However, if  $D_{\text{ping}}$  exceeds the cross-range resolution  $\delta_y$ , then the coverage of the sea floor will be incomplete by virtue of the fact that holidays will have been introduced into the imagery in the cross-range dimension.

#### 1.2.4 Frequency

A sonar's frequency is important for several reasons. First, it determines the amount of attenuation (expressed as decibels of loss per unit distance) experienced by the sonar's transmitted signal. Attenuation loss caused by the absorption of acoustic energy in seawater is described in [3]. This topic is also discussed further in Section 1.3.5.


The second factor influenced by the frequency is the appearance of objects in the sonar image. It is, in principle, possible to construct a sensor with bandwidth and aperture large enough to achieve any desired resolution,  $\delta_x$  by  $\delta_y$ . However, the system resolution alone does not determine how features appear in an image. Resolution does limit the sensor's ability to perceive them, though. The appearance of a given feature is determined by its roughness relative to the sensing wavelength. For example, at the wavelengths associated with human vision (380-780 nm), most surfaces are rough and scatter light in all directions. Those that are not, such as mirrors and other highly-polished surfaces, convey information to the viewer through their reflected light in a different manner. For example, a rough sphere held against a black background and illuminated by a single light source is visually

sensed to be spherical because of the gradations of light and shadow created by its shape. Most of its surface scatters some amount of light to the observer's eye. A perfectly smooth mirrored sphere, however, would appear visually as a bright point because only a small region on its surface will reflect light into the observer's eye. The consequences of smooth (specular) reflection and rough (diffuse) scattering must be taken into account when dealing with imaging sonar. It is important to be aware of the sonar's wavelength relative to the roughness scale of the sea floor and any objects lying on it. A good introductory treatment of the overall implications of wavelength on microwave imaging systems is found in Chapter 4 of Steinberg and Subbaram [6]. This discussion applies equally well to sonar imaging.

Thirdly, the frequency can also be used to determine if the sound will be able to penetrate a given material. For example, high frequencies are not effective at penetrating typical sea floor sediments. Depending on the sediment type, lower frequencies can be used for detection and imaging below the surface of the sea floor. There is a large body of literature on the subject of subsurface object detection using sonar.

### **1.3 Sidescan Synthetic Aperture Sonar**

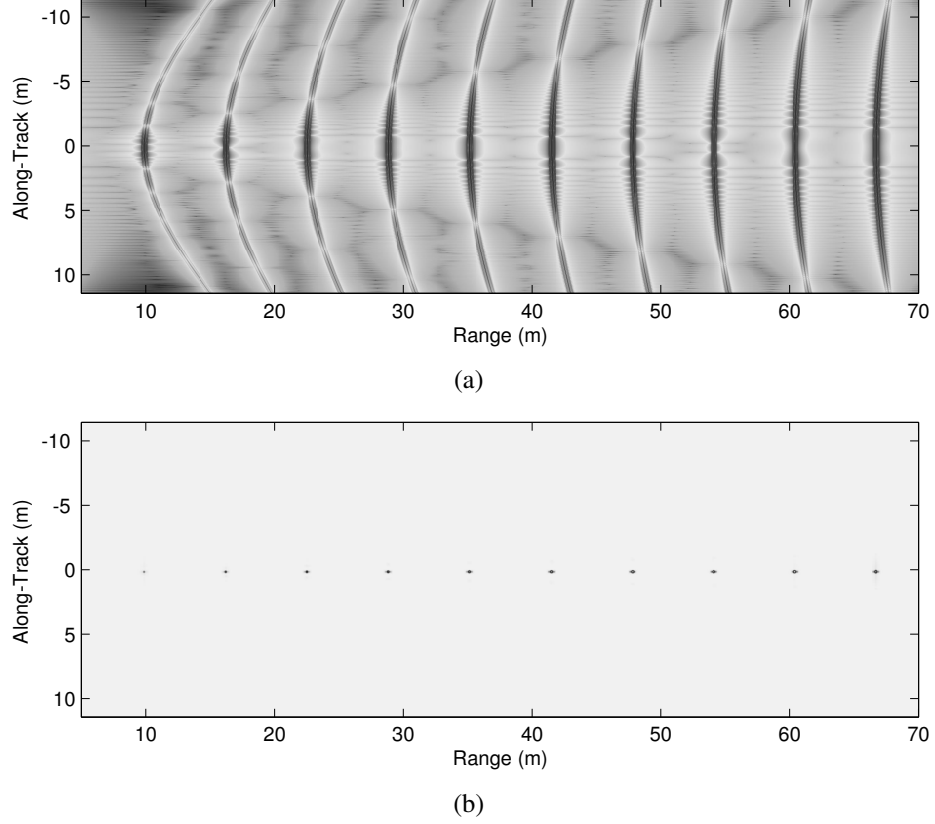
Perhaps the best way to summarize SAS compared to real aperture sonar (RAS) is to say that SAS allows the user to trade hardware complexity for software complexity. Synthetic aperture sonar is the underwater acoustic counterpart to stripmap-mode synthetic aperture radar (SAR), and the SAS community is the beneficiary of a rich body of literature generated by SAR researchers. Synthetic aperture imaging in general is closely related to other technologies such as medical tomography and seismic imaging. A survey of this family tree is beyond the scope of the present work, although a number of connections are noted by Jakowatz [7] and Hawkins [8].



**Figure 1.3.** SAS data collection is similar to the raster-style collection shown in Figure 1.1. An important difference is that SAS relies on wide physical beams to insonify the sea floor for many pings which are then coherently processed to form a synthetic aperture having a virtual beam whose cross-range footprint is independent of range.

### 1.3.1 Stripmap-Mode Synthetic Aperture Data Collection

Just as with RAS, SAS imaging is conducted using sequences of transmissions and receptions, or pings. The SAS moves through the water in the  $y$  direction at an altitude  $z$  off the sea floor. The signal is sent out along the  $x$ , or ground range, axis. The data collection for a SAS (see Figure 1.3) proceeds in much the same fashion as for a RAS. However, the time series from each ping are stored for further processing, as opposed to being immediately accumulated in an image. Consider the behavior of the observed signal as the SAS passes by a single ideal point reflector. The scattered signal first appears at some range and then moves closer until the sonar reaches the closest point of approach (CPA). Once the sonar passes the CPA, the echo migrates back out in range. The locus of points traced out by the reflected signal is a hyperbola, as shown in Figure 1.4. The length of the hyperbola is infinite, but it is nearly always considered to be equal to the -3 dB (that is, the half power)



**Figure 1.4. Simulated SAS data and reconstruction.** The upper figure (a) shows the hyperbolic signal histories that result from observing a line of ideal point scatterers. The nulls of the sensor’s directivity pattern are clearly visible, and the main lobe of the most distant hyperbola just fills the cross-range extent of the figure. At any range, the length of the synthetic aperture is generally taken to be equal to the -3 dB width of the main lobe. The lower figure (b) demonstrates that the resulting image resolution is independent of range.

beamwidth of the transducer. A SAS image of the point scatterer is created by integrating the signal along this single hyperbola. Similarly, an entire SAS image is formed by integrating over all the possible hyperbolae in the scene.

The cross-range extent of the hyperbolae depends on the beamwidth of the transducer and is approximately equal to  $R\theta_{3\text{dB}}$ . Thus, objects at longer ranges have longer hyperbolic signal histories. As a result, the number of pings required to focus a single point scatterer increases with range. For this reason SAS is said to have a constant angle of integration, as opposed to a constant integration time. Figure 1.4 shows a row of point scatterers imaged using SAS processing as well as the magnitude of the raw echo data used to create the image. Note that the cross-range resolution is independent of range. The real aperture

cross-range resolution at a range of 70 m would be seven times worse than at 10 m.

The process described above is known as stripmap mode synthetic aperture data collection. This thesis concentrates solely on the stripmap modality, as it is by far the most commonly encountered for SAS. There are other non-stripmap modes (see [9, 10, 11] for example), but these are not explicitly considered here. Nevertheless, many of the results presented might be usefully extended to other imaging modalities.

The operational properties of SAS can be compared to RAS using the four parameters listed in the previous section:

### 1.3.2 Range Resolution

The range resolution of RAS and SAS is identical,  $\delta x = c/2B_w$ , because it only depends on the bandwidth of the transmitted signal.

### 1.3.3 Cross-Range Resolution

The cross-range (or along-track) resolution behavior is arguably the defining difference between real and synthetic aperture sonar. The cross-range resolution of stripmap SAS is constant with range and is given by:

$$\delta_{y,\text{SAS}} = \frac{D}{2}, \quad (1.3)$$

where  $D$  is the horizontal length of the projector or receiver, whichever is larger [12]. In order to achieve this resolution, the received echoes are coherently integrated over the synthetic aperture length, which is equal to  $R\theta_{3\text{dB}}$ , the distance between the -3 dB points of the beam at range  $R$ . The relationship in (1.3) is typically derived using linear (1D) or rectangular (2D) transducers, but it is valid for any transducer whose directivity pattern satisfies the condition  $\theta_{3\text{dB}} \approx \lambda/D$ .

The fact that SAS cross-range resolution is equal to  $D/2$  leads to a remarkable conclusion: small sensors actually yield better resolution. Of course, there are penalties for making  $D$  very small. Small sensors generally have poorer gain than larger ones. Another concern is the computational expense of forming the image, as decreasing  $D$  results in a

larger beamwidth and, consequently, a longer synthetic aperture. The practical difficulty of addressing the motion compensation problem over a large synthetic aperture is another concern. Decreasing  $D$  without increasing the number of array elements also reduces the area coverage rate. Nevertheless, it is clear that SAS overcomes one of the major design challenges for a RAS. A synthetic aperture sonar enables one to achieve high cross-range resolution using a small sensor. For comparison, one can compute the length required for a real aperture sonar to equal a given SAS cross-range resolution (assuming the same frequency in both cases):

$$\begin{aligned}\delta_{y,\text{RAS}} &= \delta_{y,\text{SAS}} \\ &= \frac{R\lambda}{D_{\text{RAS}}}.\end{aligned}\tag{1.4}$$

Solving for the real aperture length  $D_{\text{RAS}}$  gives:

$$\begin{aligned}D_{\text{RAS}} &= \frac{R\lambda}{\delta_{y,\text{SAS}}} \\ &= \frac{2R\lambda}{D_{\text{SAS}}} \\ &= 2L_{\text{SA}}.\end{aligned}\tag{1.5}$$

It can be seen that the real aperture required to achieve a cross-range resolution of  $\delta_{y,\text{SAS}}$  would have to be equal to twice the length of the synthetic aperture.

Yet another important property of SAS is that the cross-range resolution is independent of frequency. One thus has a great deal of flexibility when choosing the operational band of a SAS. SAS permits the use of low frequencies for high-resolution imaging (subject to the considerations mentioned in Section 1.2.4 and in [6]). Low frequencies have the added benefit of propagating with less attenuation. Consequently, for a given cross-range resolution, a SAS will generally far outperform a RAS in terms of range capability.

### 1.3.4 Area Coverage Rate

The ACR of SAS is more constrained than for RAS in the sense that there exists a strict cross-range sampling requirement for SAS. At a minimum, a synthetic aperture sonar must



ping (sample spatially or transmit and receive) at intervals no greater than  $D/2$ . The details of cross-range sampling requirements are addressed by several authors. For example, Rolt [13] and Hawkins [8] both discuss the sampling constraints in the context of stripmap SAS and provide references to other sources in the SAR literature. The cross-range sampling affects the overall contrast, or dynamic range, of the SAS image. This can be quantified using the peak to grating lobe ratio (PGLR) and the along-track ambiguity to signal ratio (AASR) [8, 14]. The aliasing caused by the spatial undersampling of the synthetic aperture is also of concern for interferometric applications. It is a form of noise that lessens the coherence and therefore adversely affects the phase estimation required to measure bathymetry [15].

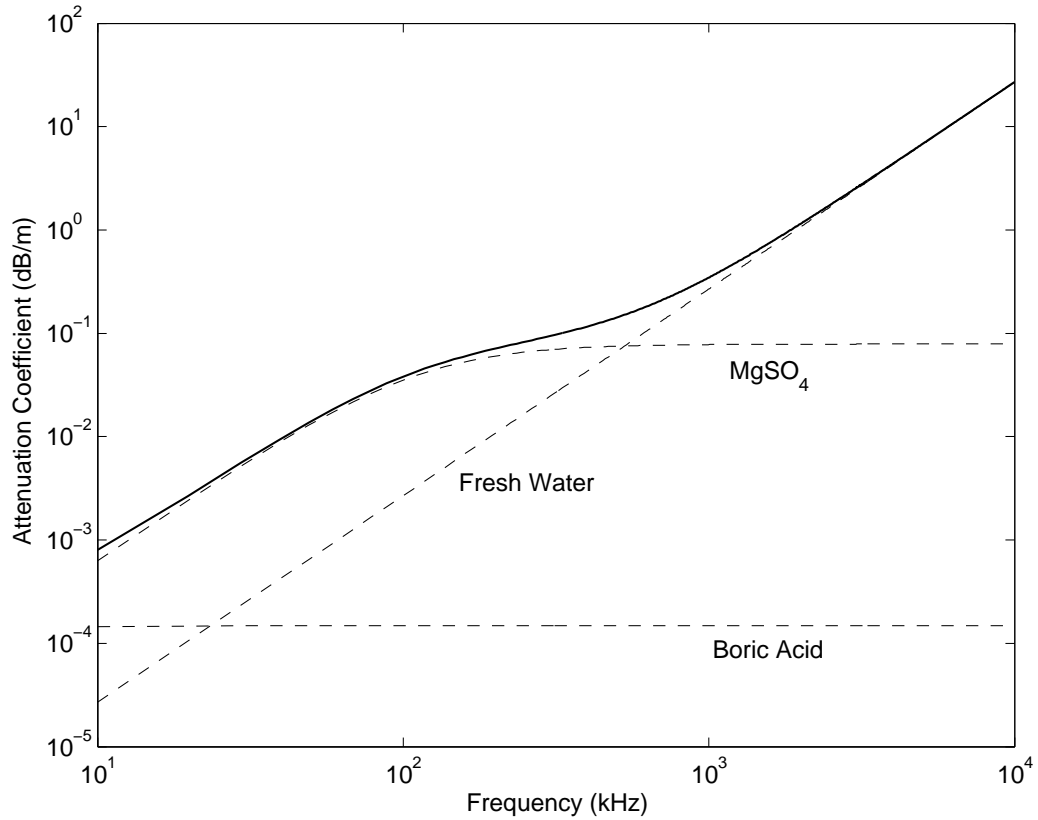
The equation for SAS area coverage rate can be found by multiplying Equation (1.2) by the forward velocity  $v$ :

$$ACR_{SAS} \leq R_{\max} v = \frac{Dc}{2}. \quad (1.6)$$

The inequality is used in order to account for the case in which the SAS samples more frequently than  $D/2$ . Interestingly, the ACR for SAS is a fixed number. It is also generally a low number for high-resolution systems. The low ACR of SAS appears to be a shortcoming, but it is easy to overcome by using arrays of receivers in conjunction with a single transmitter. This scheme can also be applied to RAS and is treated in Section 1.4 below.

### 1.3.5 Frequency

The previous comments regarding the choice of frequency for RAS apply to SAS as well. With respect to underwater imaging, one of the greatest advantages of SAS is that it can operate at much lower frequencies than RAS for any given cross-range resolution. Higher frequencies attenuate far more quickly in seawater, thus severely limiting the useful range of the RAS. Figure 1.5 shows the attenuation per meter as a function of frequency as described by Clay and Medwin [3]. In fresh water, the attenuation increases approximately as the square of the frequency. In sea water the absorption is augmented by the phenomenon of molecular relaxation, which is the rearrangement of certain molecules in response to the



**Figure 1.5. Attenuation of sound in sea water expressed in decibels per meter. The dashed lines represent the three individual phenomena contributing to the overall absorption loss shown by the solid line. (After Clay and Medwin, [3])**


acoustic pressure fluctuation. Boric acid and magnesium sulfate are the compounds present in sea water that exhibit relaxation.

To illustrate the importance of this effect, consider a SAS having 2 cm cross-range resolution and a center frequency of 100 kHz. A commercially-available RAS with similar performance (because of the cost/difficulty of building large transducers) needs to operate around 1 MHz. The absorption of sound energy by seawater at 1 MHz is approximately nine times larger than it is at 100kHz (0.35 dB/m vs. 0.038 dB/m). The 1 MHz real aperture sonar might match and even exceed the SAS resolution at short range. However, its range is limited to approximately 30 m. In contrast, SAS at 100 kHz can produce high-resolution images at long ranges (well in excess of 200 m)—a feat impossible to achieve with RAS.

## 1.4 Vernier Sonar Arrays

As discussed previously, the area coverage rate is severely limited when the distance between consecutive pings must be small. For SAS, the ACR is especially poor because of the spatial sampling constraint. Nevertheless, there is a simple solution: the number of receivers can be increased from one to an arbitrary number,  $N$  [16]. For a single transmission this array of receivers collects multiple cross-range samples simultaneously. When the sonar moves forward and sends its next transmission into the water, it can advance  $N$  times farther than it could if only a single receiver were employed. As a result, the ACR is increased by a factor of  $N$ . This configuration is often referred to as a *vernier array*, and it is virtually always employed when constructing a SAS to be used for at-sea imaging. It is less common for laboratory or tank experiments because the sensor is often moved by rail, and the ACR is usually of no concern.

The key concept for understanding the use of vernier arrays is the notion of the phase center. A phase center is simply defined as the point at which a signal appears to originate. Most transducers, antennas, *etc.*, radiate from an area with an appreciable physical dimension. For example, the face of a flat rectangular hydrophone may be many wavelengths across at its operational frequency. When observed from a distance, the phase center of this hydrophone is at the center of its face (assuming the face moves uniformly) and is used to specify its location if it were to be treated as a point source. For a vernier array, in which each receiver forms a bistatic pair with the transmitter, the phase center is the point located midway between the transmitter and receiver. Thus, at a distance, the physical vernier array can be treated as a virtual phase center array that is exactly half as long and consists of  $N$  monostatic transmit/receive elements. The phase center and vernier array concepts are illustrated in Figures 1.6 and 1.7. It can be seen from the figure that the illusion of the phase center array is not complete until a small delay is removed from the data:  $\tau_{PC} = (R_{TX} + R_{RX} - 2R_{PC})/c$ . The nature of this correction is clear from the geometry, and it is discussed in detail by Bonifant [17].

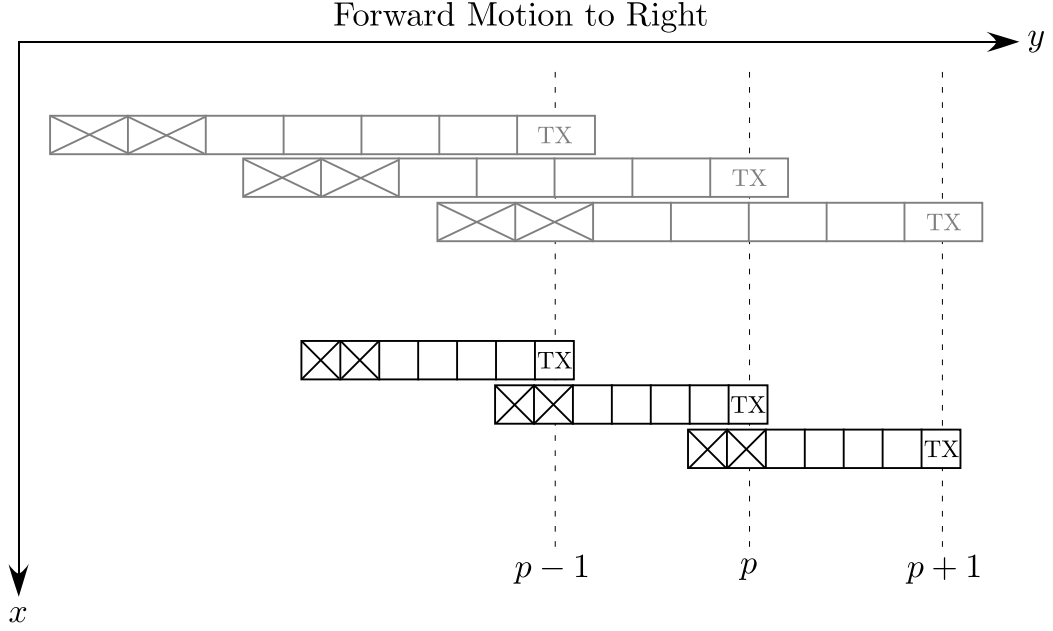


**Figure 1.6.** The diagram above illustrates the phase center concept. At long ranges, a bistatic source and receiver pair behaves as a single equivalent sensor that is located midway between the physical sensor locations. The location of this virtual sensor is called the phase center of the source/receiver pair. The signal observed by the phase center is assumed to be the same as that observed by the physical pair of sensors, with the exception of a delay equal to  $\tau_{PC} = (R_{TX} + R_{RX} - 2R_{PC})/c$ .

The idea of replacing the actual bistatic configuration with a monostatic virtual transmit/receive pair is critical for correct implementation of Fourier-based image reconstruction techniques. These are derived using the monostatic assumption, so the observed data must be conditioned to appear as if it were actually collected monostatically. This phase center correction is less important for spatiotemporal domain reconstruction. While it may be used in order to simplify or improve the efficiency of the reconstruction algorithms, the reconstruction method itself is perfectly capable of accommodating the bistatic geometry.

## 1.5 Practical Concerns for Successful SAS Imaging


It is desirable for the SAS sensor to ping at uniform predetermined spatial intervals. Otherwise, there is a risk of undersampling in the cross-range dimension. Also, efficient SAS image reconstruction algorithms assume regular sampling in both range and cross-range. If uniform sampling is not achieved, it is sometimes possible to resample the data onto a regularly-spaced grid. This procedure is described in Chapters 5 and 6.



**Figure 1.7.** A vernier array is constructed using the phase center concept illustrated in Figure 1.6. The apparent length of the resulting phase center array (or virtual array) is exactly half that of the physical array. Consequently, the advance per ping is determined by the length of the phase center array. The diagram above shows this reasoning applied to an array of six receivers, with the trailing two elements overlapping for the purpose of motion estimation using redundant phase centers.

The ping timing is typically accomplished using the AUV's on-board estimate of its speed over ground (SOG), which in turn is usually estimated using a device known as a Doppler velocity log (DVL). A DVL operates by sending sound waves in orthogonal directions down toward the sea floor and observing the Doppler frequency shift of the backscattered signal. The measured Doppler shift in the orthogonal directions is then used to estimate velocity. DVLs vary in accuracy, and their speed estimates are often fed through a smoothing filter prior to being used to trigger the SAS. These navigation/control filters combine the available information from the on-board motion sensors to compute the best estimate of the AUV's current speed and position.

AUVs travel slowly, and undersea navigation is challenging due to the fact that radio waves do not propagate underwater, rendering useless technologies like LORAN and GPS. For the most part, the velocity/position of an AUV is known from combining the



**Figure 1.8.** The photo above depicts the Small Synthetic Aperture Minehunter system developed for the Office of Naval Research by the Naval Surface Warfare Center in Panama City, Florida [18]. The SAS (the white section) was constructed by the Applied Research Laboratory of the Pennsylvania State University. The vehicle is a Remus 600 supplied by the Woods Hole Oceanographic Institution. (Photo courtesy of Mr. Daniel Brown.)

DVL output with the output of an inertial measurement unit (IMU). In general, the rotational components of motion are easier to measure with sufficient accuracy, while the translational components are more difficult. For this reason, the techniques presented in Chapter 5 largely assume that the vehicle rotations are known and that the translations are to be estimated using the SAS data itself. An example of an AUV-based SAS is shown in Figure 1.8 (also see [18]).

## **1.6 Display of Synthetic Aperture Sonar Imagery**

Imagery generated by synthetic aperture sonar is generally characterized as having a high dynamic range. In other words, a single scene usually contains both very bright and very dark pixels (a magnitude ratio of 1,000,000:1, or 60 dB, is not unusual). The human eye also has a large dynamic range of approximately 90 dB, being able to function both in starlight and sunlight. However, the eye can only perceive about 30 dB of dynamic range at any given moment. Because of this, it is often challenging to display SAS imagery in a manner that allows the human eye to observe both bright and dark details simultaneously. This problem is compounded by the fact that many common display devices, such

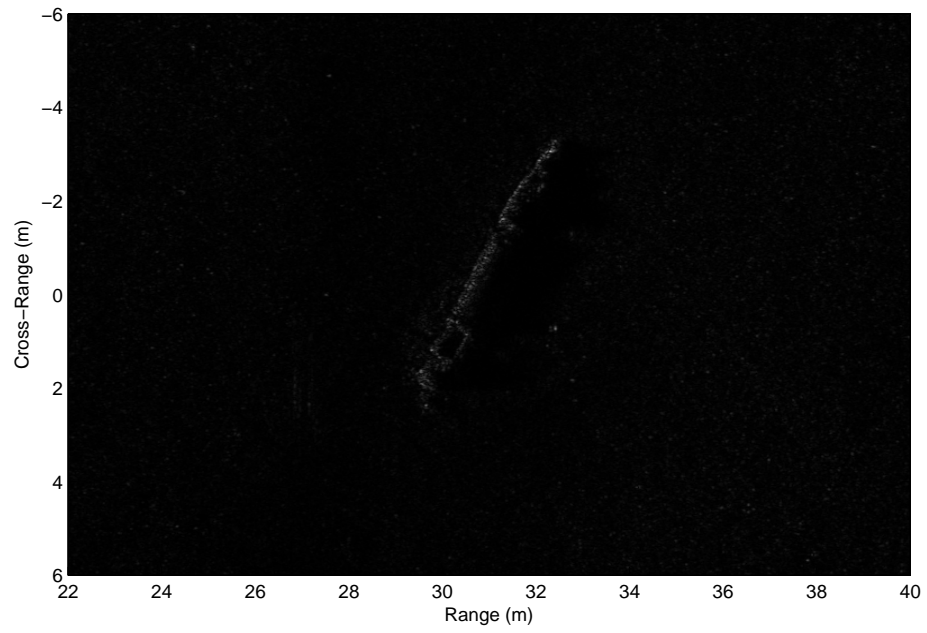
as monitors and printers, have a dynamic range that is less than 30 dB.

There are any number of techniques used to compress high dynamic range data to make viewing and interpretation easier. These range from simply taking the logarithm of the image intensity to state-of-the-art techniques based on partial differential equations (see Fattal *et al.* [19], for example). However, sonar imagery is somewhat unique in the sense that it is often viewed in massive quantities by human operators subject to boredom and fatigue. For this reason, making the output pleasing to the eye is as important as accurately representing the information it contains.


SAS imagery displayed without regard to its dynamic range often appears to be dark. This is due to the character of the statistical distribution of the amplitude of the scene reflectivity as represented by the image pixel values. Figure 1.10 shows a typical example of such an image before and after applying the dynamic range modification technique described in this section. Meanwhile, Figure 1.10(a) shows the probability density function (PDF) estimated using the histogram of this image. The pixel magnitudes are normalized to fall within the range  $[0, 1]$  prior to the computation. Note that the vast majority of the pixels are tightly clustered around the value of 0.01 and that the horizontal axis only covers the range  $[0, 0.07]$ . Also shown is a curve describing a gamma distribution PDF that was generated using parameters estimated from the SAS image.

The technique employed in this thesis to deal with the large dynamic range of SAS imagery is to expand the histogram so that the pixel magnitudes are spread over a wider range of the display gamut. A simple algorithm is used:

1. Estimate the shape and scale parameters of the gamma distribution PDF that best fits the pixel amplitude distribution found in a given scene.
2. Define the tails of this distribution as beginning at  $1/1000^{\text{th}}$  of the peak value of the best-fit gamma PDF. These locations (pixel magnitudes) are denoted as  $I_{\text{lower}}$  and  $I_{\text{upper}}$ .




(a)




(b)

**Figure 1.9.** The upper image (a) shows an example SAS image displayed without any form of normalization. A small number of bright pixels forces the majority of the image to be too dark. The lower image (b) is the result of applying the normalization procedure described in this section to the upper image.





(a)



(b)

**Figure 1.10.** The upper plot (a) shows the gamma distribution PDF estimated using the pixel magnitude values from Figure 1.9(a). Also shown is the estimate of the image magnitude PDF computed using the histogram. The horizontal axis is truncated to emphasize the shape of the PDF. The lower plot (b) shows the histogram after modification using the algorithm described in this section. The horizontal axis in this plot is not truncated. The gamma PDF from the upper plot is included for comparison, and its values are given by the right-hand vertical axis. Note the large concentration of dark pixels relative to bright prior to adjusting the histogram.

3. Set all the values in the image which are less than  $I_{\text{lower}}$  to be equal to  $I_{\text{lower}}$ .
4. Set all the values in the image which are greater than  $I_{\text{upper}}$  to be equal to  $I_{\text{upper}}$ .
5. Rescale the pixel values to range from 0 to 1. In other words,  $I_{\text{lower}}$  and  $I_{\text{upper}}$  now correspond to 0 and 1 respectively.

This scheme tends to produce solid patches of saturated pixels if a bright extended target is present. The effect can be visually unappealing and can result in the loss of important detail. Instead of clipping all the pixels in the upper tail to a single value, they can be allocated to a subset of the display gamut. This is accomplished using the following simple modification to the normalization algorithm:

4. Assign a transition point  $T$  between the values of 0 and 1. All values less than  $I_{\text{upper}}$  are then rescaled to cover the range  $[0, T]$  instead of  $[0, 1]$  as before.
5. Rescale all the values in the image which are greater than  $I_{\text{upper}}$  according to:

$$I = T + \frac{I - I_{\text{upper}}}{1 - I_{\text{upper}}}(1 - T).$$

By selecting  $T$ , the user can assign a portion of the display gamut to the very bright highlights without totally clipping them or allowing them to suppress the apparent brightness of the majority of the image content. The resulting highlights retain some of their brightness gradation. An example application is shown in Figure 1.10(b) in which  $T = 0.8$ . This plot shows that a cluster of bright pixels occurs at  $I = 0.8$ , but that there are still brighter pixels with magnitudes ranging up to 1. Also shown in the plot is the gamma PDF used to fit the original data in Figure 1.10(a). The effect of expanding the histogram is dramatic, as shown in Figure 1.10(b). This scheme was used to normalize all the SAS imagery in this document, with the exception of Figure 3.7 which is plotted on a decibel scale, and Figures 5.8 and 5.14 which were reproduced directly from earlier publications.

The approach described above is simple and appears to yield reasonable image quality over a wide range of output devices. However, no claims are made regarding its optimality.

The gamma distribution was chosen because it fits the data well and because reasonably accurate closed-form estimates exist for computing the best-fit parameters. There is also some justification in the literature for this choice, as high-resolution low grazing angle SAR imagery is documented to be well-modeled by gamma, Rayleigh, and Weibull distributions [20, 21, 22, 23].

## 1.7 Comments

The sea is restless and always in motion. This aspect of its nature is beloved by poets, but not so endearing to engineers attempting to image the sea floor. For SAS it is particularly bothersome, as the signal processing required for image formation becomes simpler and more computationally efficient if the sonar travels a perfectly straight line through the water. The techniques for using the SAS data to help measure the unwanted motion and the subsequent compensation are the subject of this thesis. These techniques are described in detail for practical implementation, and an analysis of the effects of unwanted platform motion is given.

The basics of SAS image reconstruction are covered in Chapter 2. This discussion, however, is less practical and is instead aimed at highlighting the assumptions and simplifications involved in deriving the most important reconstruction methods. Hawkins's 1996 Ph.D. dissertation [8] gives a thorough explanation of how to actually implement a number of SAS reconstruction techniques. It also primarily deals with stripmap imaging, which is the most common form of SAS. The synthetic aperture radar (SAR) community has produced a number of excellent texts having broad applicability. Some of the more recent and well-known of these are by Carrara *et al.* [24], Jakowatz *et al.* [7], and Soumekh [25, 12].

A final note regarding terminology is in order. Motion estimation and compensation are sometimes referred to as *autofocus* in the SAS literature since the image is corrected using its own data. However, the SAR literature has a long history of using the term *autofocus* to refer to postprocessing done on the reconstructed image that brings it into sharper

focus. This sort of autofocus relies on the image data for finding the desired correction, and the two types of autofocus have little in common. The term *micronavigation* is also widespread in the SAS literature as a way to refer to motion estimation and compensation. This term is fitting in that typical AUV on-board navigation systems are too inaccurate to use for high-resolution SAS imaging. In the present work the term autofocus is used in its traditional SAR meaning. Meanwhile, the term micronavigation is generally avoided because it is useful to discuss motion estimation and compensation separately, and the term micronavigation tends to lessen the distinction between these two concepts.

## CHAPTER 2

### OVERVIEW OF SYNTHETIC APERTURE SONAR IMAGE RECONSTRUCTION

#### 2.1 Introduction

Although the subject of this thesis is motion estimation and compensation for synthetic aperture sonar, it is useful to provide a brief overview of the synthetic aperture imaging model. The early discussion in this chapter follows a combination of that found in [26] and Section 6.6 of [27]. The purpose here is not to give a rigorous mathematical development of synthetic aperture imaging. Rather, an outline of such a development is presented, and it is shown exactly where the major simplifying assumptions are made. With this knowledge, the limitations of the traditional SAS imaging model become clear.

#### 2.2 Coordinate System

The coordinate system used for the motion estimation is shown in Figure 2.1. It is an ordinary right-handed set of axes except that the rotations about the  $z$ -axis are referenced to the  $y$ -axis, not the  $x$ -axis. This is done so that a yaw angle of zero corresponds to the array being parallel to the along-track direction. The notation for translations and rotations is summarized in Table 2.1. The name for each kind of motion is shown in normal type. The italicized type indicates the terms used for error, or unwanted motion, in the given direction.

#### 2.3 The Wave Equation and the Green's Function

This section develops a model for the signal observed at a point in space caused by a time-dependent source of finite spatial extent. The development begins with the nonhomogeneous wave equation:

$$\nabla^2 \Psi(t, \mathbf{x}) - \frac{1}{c^2} \frac{\partial^2 \Psi(t, \mathbf{x})}{\partial t^2} = -f(t, \mathbf{x}) \quad (2.1)$$

**Table 2.1. Conventions used to describe platform motion. The terms for unwanted motions are shown in italics.**

$x$	range ( <i>sway</i> )
$y$	cross-range ( <i>surge</i> )
$z$	altitude ( <i>heave</i> )
$\psi$	pitch ( <i>pitch</i> )
$\phi$	roll ( <i>roll</i> )
$\theta$	heading ( <i>yaw</i> )


in which  $\Psi(t, \mathbf{x})$  represents the amplitude of a field variable,  $c$  is the speed of propagation (assumed constant), and  $f(t, \mathbf{x})$  is a source function. For acoustic waves, it is customary for  $\Psi$  to be the excess pressure relative to the ambient pressure in the quiescent medium. In sensing applications  $\Psi$  often represents the voltage at the output of the transducer, which is usually assumed to be proportional to the field variable being measured.

### 2.3.1 Time-Independent Green's Function

In the derivation of linear acoustics, all of the field equations are linear and have time-independent coefficients [28]. Thus, it is possible for the field variables to oscillate with the same frequency, but not necessarily the same phase, everywhere in the medium [28]. This suggests that (2.1) should be analyzed by considering each frequency separately. This is done by taking the Fourier transform of (2.1) with respect to time,

$$(\nabla^2 + k^2)\Psi(\omega, \mathbf{x}) = -f(\omega, \mathbf{x}), \quad (2.2)$$

where  $k = \omega/c$  is the wavenumber which has units of radians per meter. Equation (2.2) is known as the nonhomogeneous Helmholtz equation. The advantage of the Helmholtz equation is that the number of explicit unknowns has been reduced by one; the time dependence for each Fourier component is known to be an oscillation with frequency  $\omega$ . The problem is thus reduced to a purely spatial one. The left side of (2.2) is a form of Sturm-Liouville operator on  $\Psi$ , which is denoted as  $\mathcal{L} = (\nabla^2 + k^2)$ . As such, it belongs to a class of operators whose inverse exists and is represented as an integral operator whose kernel is known



**Figure 2.1. Right-handed coordinate system used for motion estimation and compensation. The array yaw is referenced to the positive  $y$ -axis, as this is the direction of forward motion.**

as the Green's function [29]. In the present context, the Green's function corresponding to  $\mathcal{L}$  is defined to be the solution of the point source nonhomogeneous equation:

$$\mathcal{L}G(\mathbf{x}, \mathbf{x}') = -\delta(\mathbf{x} - \mathbf{x}') \quad (2.3)$$

subject to the same boundary condition as (2.1). This is the condition of outward-propagating waves in free space, also known as the *Sommerfeld radiation condition*, under which the required Green's function is

$$G(\mathbf{x}; \mathbf{x}') = \frac{\exp\{-jkR\}}{4\pi R}, \quad (2.4)$$

where  $R = |\mathbf{x} - \mathbf{x}'|$  is simply the distance between the source at  $\mathbf{x}'$  and observer at  $\mathbf{x}$ . Knowledge of the Green's function allows one to proceed toward the objective of calculating the field due to an arbitrary source. This field can be shown to be the convolution of the Green's function with the source distribution [30],

$$\begin{aligned} \Psi(\omega, \mathbf{x}) &= \int f(\omega, \mathbf{x}') G(\mathbf{x}; \mathbf{x}') d^3 \mathbf{x}' \\ &= \int f(\omega, \mathbf{x}') \frac{\exp\{-jk|\mathbf{x} - \mathbf{x}'|\}}{4\pi|\mathbf{x} - \mathbf{x}'|} d^3 \mathbf{x}', \end{aligned}$$

in which  $\mathbf{x}$  is the point of observation and the integral is carried out over all points  $\mathbf{x}'$  for which the source exists. Thus, the Green's function serves to weight the influence of  $f$  at the observer location. For this reason, the Green's function is also commonly referred to as the spatial factor of the operator  $\mathcal{L}$ .

### 2.3.2 Time-Dependent Green's Function

Next, the derivation above is repeated with the goal of constructing the time-dependent Green's function that satisfies

$$\left(\nabla^2 - \frac{1}{c^2} \frac{\partial^2}{\partial t^2}\right) G(t, \mathbf{x}; t', \mathbf{x}') = -\delta(\mathbf{x} - \mathbf{x}') \delta(t - t'). \quad (2.5)$$

In this case, the nonhomogeneous Helmholtz equation (2.2) with this new source term becomes:

$$(\nabla^2 + k^2) \Psi(\mathbf{x}, \omega) = -\delta(\mathbf{x} - \mathbf{x}') \exp\{-j\omega t'\}, \quad (2.6)$$

and the solution is  $G(\mathbf{x}; \mathbf{x}') \exp\{-j\omega t'\}$  [27]. Taking the inverse Fourier transform gives the desired time-dependent Green's function:


$$\begin{aligned} G(t, \mathbf{x}; t', \mathbf{x}') &= \frac{1}{2\pi} \int \frac{\exp\{-jkR\}}{4\pi R} \exp\{j\omega(t - t')\} d\omega \\ &= \frac{1}{2\pi} \int \frac{\exp\{j[-R/c + (t - t')]\omega\}}{4\pi R} d\omega \\ &= \frac{\delta\left(t - \left[t' + \frac{|\mathbf{x} - \mathbf{x}'|}{c}\right]\right)}{4\pi|\mathbf{x} - \mathbf{x}'|}. \end{aligned} \quad (2.7)$$

This result may then be used to describe the time-dependent solution of the linear wave equation in free space for any source function  $f(t', \mathbf{x}')$ :

$$\begin{aligned} \Psi(\mathbf{x}, t) &= \int \int f(\mathbf{x}', t') G(\mathbf{x}, t; \mathbf{x}', t') d^3\mathbf{x}' dt' \\ &= \int \int f(\mathbf{x}', t') \frac{\delta\left(t - \left[t' + \frac{|\mathbf{x} - \mathbf{x}'|}{c}\right]\right)}{4\pi|\mathbf{x} - \mathbf{x}'|} d^3\mathbf{x}' dt'. \end{aligned} \quad (2.8)$$

The next section applies this result to obtain the time-dependent field due to a rectangular source.





**Figure 2.2.** Rectangular distribution of point sources used to model a sonar projector. The source is located at the origin and lies in the  $y$ - $z$  plane. The width and height are  $D$  and  $H$ , respectively.

## 2.4 Radiation from a Rectangular Source

In the following, the signal emitted from a point source is modeled as a complex sinusoid modulated by a slowly-varying envelope function:  $P(t) = A(t) \exp j\omega t$ . When the point source is placed at  $\mathbf{x}$ , the resulting field  $\Psi(t, \mathbf{x} - \mathbf{x}')$  satisfies the following wave equation:

$$\left( \nabla^2 - \frac{1}{c^2} \frac{\partial^2}{\partial t^2} \right) \Psi(t, \mathbf{x}; t', \mathbf{x}') = -P(t) \delta(\mathbf{x} - \mathbf{x}'). \quad (2.9)$$

The solution is then calculated from the spatiotemporal convolution (2.8) above:

$$\Psi(t, \mathbf{x}) = \frac{A\left(t - \frac{|\mathbf{x} - \mathbf{x}'|}{c}\right)}{4\pi|\mathbf{x} - \mathbf{x}'|} \exp \left\{ j\omega \left( t - \frac{|\mathbf{x} - \mathbf{x}'|}{c} \right) \right\}. \quad (2.10)$$

This is the field observed at  $\mathbf{x}$  due to the prescribed point source at  $\mathbf{x}'$ . Next consider a continuous rectangular distribution of identical point sources with horizontal length  $D$  and height  $H$  that is centered at the origin and is parallel to the  $y$ - $z$  plane (see Figure 2.2). As above, a given location on the rectangular source is denoted as  $\mathbf{x}'$ , and points of observation are designated by  $\mathbf{x}$ . Unit vectors are indicated using the tilde. Furthermore, the following

far-field approximation is made:

$$\begin{aligned} |\mathbf{x} - \mathbf{x}'| &= |\mathbf{x}| - \tilde{\mathbf{x}} \cdot \mathbf{x}' + O\left(\frac{D^2}{|\mathbf{x}|}\right) \\ &\approx |\mathbf{x}| - \tilde{\mathbf{x}} \cdot \mathbf{x}', \end{aligned} \quad (2.11)$$

in which the distance from the source to the observer is much larger than the characteristic dimension of the source (that is,  $|\mathbf{x}| \gg D$ ). This approximation is common in acoustics and optics, and it is also used in Chapter 6 in a different context. Here it leads to a closed-form solution for the field  $\Psi$  due to the distributed source. Next, insert the expansion (2.11) into (2.10) and use the zeroth-order expansion of (2.11) for the amplitude term because it is assumed to be slowly varying and because  $|\mathbf{x}| \gg |\tilde{\mathbf{x}} \cdot \mathbf{x}'|$ . The first-order expansion, however, is used for the phase of the complex exponential. The result is

$$\begin{aligned} \Psi(t, \mathbf{x}) &\approx \frac{A\left(t - \frac{|\mathbf{x} - \mathbf{x}'|}{c}\right)}{4\pi|\mathbf{x} - \mathbf{x}'|} \exp\left\{j\omega\left(t - \frac{|\mathbf{x} - \mathbf{x}'|}{c}\right) + jk(\tilde{\mathbf{x}} \cdot \mathbf{x}')\right\} \\ &= \frac{P\left(t - \frac{|\mathbf{x} - \mathbf{x}'|}{c}\right)}{4\pi|\mathbf{x} - \mathbf{x}'|} \exp\{jk(\tilde{\mathbf{x}} \cdot \mathbf{x}')\}. \end{aligned} \quad (2.12)$$

With this approximation, (2.8) can be evaluated to obtain a closed-form solution for the field far away from the rectangular point source distribution:

$$\begin{aligned} \Psi(t, \mathbf{x}) &\approx \frac{P\left(t - \frac{|\mathbf{x} - \mathbf{x}'|}{c}\right)}{4\pi|\mathbf{x} - \mathbf{x}'|} \int_{-D/2}^{+D/2} \int_{-H/2}^{H/2} \exp\{jk(\tilde{\mathbf{x}} \cdot \mathbf{x}')\} dx dy \\ &= \frac{P\left(t - \frac{|\mathbf{x} - \mathbf{x}'|}{c}\right)}{4\pi|\mathbf{x} - \mathbf{x}'|} \int_{-D/2}^{+D/2} \exp\{jk \cos \alpha\} dx \int_{-H/2}^{+H/2} \exp\{jk \cos \beta\} dy \\ &= \frac{P\left(t - \frac{|\mathbf{x} - \mathbf{x}'|}{c}\right)}{4\pi|\mathbf{x} - \mathbf{x}'|} \text{sinc}\left(\frac{kD}{2} \cos \alpha\right) \text{sinc}\left(\frac{kH}{2} \cos \beta\right) \\ &= \frac{A\left(t - \frac{|\mathbf{x} - \mathbf{x}'|}{c}\right) \exp\left\{j\omega\left(t - \frac{|\mathbf{x} - \mathbf{x}'|}{c}\right)\right\}}{4\pi|\mathbf{x} - \mathbf{x}'|} \text{sinc}\left(\frac{kD}{2} \cos \alpha\right) \text{sinc}\left(\frac{kH}{2} \cos \beta\right), \end{aligned} \quad (2.13)$$

in which  $\alpha$  is the azimuthal angle (about the  $z$ -axis) and  $\beta$  is the elevation angle (about the  $y$ -axis). These angles are referenced to boresight, which is the outward normal centered on the face of the rectangular source.

The solution in (2.13) is for a single frequency  $\omega$ . Thus, more complicated solutions can be constructed by integrating over all the appropriately-weighted frequencies in a desired

pulse. Noting that the complex exponential term in (2.13) can be written in two parts, and carrying out the integration with respect to  $\omega$  yields the time-space domain solution for a given transmitted signal observed at  $\mathbf{x}$

$$S(t, \mathbf{x} - \mathbf{x}') = \frac{A\left(t - \frac{|\mathbf{x} - \mathbf{x}'|}{c}\right)}{4\pi|\mathbf{x} - \mathbf{x}'|} *_t W(t, \mathbf{x} - \mathbf{x}') \quad (2.14)$$

in which the term  $W(t, \mathbf{x} - \mathbf{x}')$  represents the spatiotemporal response of the rectangular source and  $*_t$  indicates convolution with respect to time.

## 2.5 Modeling the Reflection from the Sea Floor

This section presents the traditional model for the scattered signal used in the mathematical development of the standard synthetic aperture image reconstruction algorithms. This model employs the following assumptions (see Section II.2.2 of [17]): the sound velocity is constant throughout the medium (for example, multipath and medium inhomogeneity are ignored, and the signal travels a straight path), the sea floor is modeled as a surface continuum of ideal omnidirectional point scatterers, there are no occlusions, and the amplitude of the sound scattered by an object is much smaller than that of the incident field. The last is the so-called *Born approximation* or *small scattering approximation* [28, 27]. This assumption is important because it implies that multiple scattering is neglected. Lastly, the *stop and hop* or *stop-start* condition is imposed in which the sensor is considered to remain effectively motionless during the time of flight of a single ping.

Under the above assumptions, the echo signal received at a point  $\mathbf{x}$  by the sensor used for both transmission and reception is the sum of all the signals scattered by the individual point reflectors  $f(\mathbf{x}')$  on the sea floor weighted by the sensor's two-way directivity pattern  $W(t, \mathbf{x} - \mathbf{x}')$ :

$$E(t, \mathbf{x}) = \int f(\mathbf{x}') \frac{A\left(t - \frac{2|\mathbf{x} - \mathbf{x}'|}{c}\right)}{(4\pi)^2|\mathbf{x} - \mathbf{x}'|^2} *_t W(t, \mathbf{x} - \mathbf{x}') d^3\mathbf{x}'. \quad (2.15)$$

The temporal Fourier transform of the received signal is given by:

$$E(\omega, \mathbf{x}) = \int \frac{f(\mathbf{x}')A(\omega)W(\omega, \mathbf{x} - \mathbf{x}')}{(4\pi)^2|\mathbf{x} - \mathbf{x}'|^2} \exp\{-j2k\sqrt{\mathbf{x} - \mathbf{x}'}\} d^3\mathbf{x}'. \quad (2.16)$$

Common simplifications to these expressions are to neglect the  $1/(4\pi R)^2$  spreading term ( $1/4\pi R$  out to the scatterer and another  $1/4\pi R$  for the reflected signal going back to the sensor) and the directivity pattern  $W$ . The former can be compensated using simple time-varying gain, and the latter can be deconvolved in the frequency domain as set forth in Chapter 4 of [8]. Note that  $W$  in the present context represents the weighting for both transmission and reception. In (2.14),  $W$  indicated the transmit weighting alone.

The resolution in the range direction is unaffected by synthetic aperture processing. It is inversely proportional to the bandwidth of the transmitted pulse. High-bandwidth signals can be generated using very short tone bursts. However, the peak power requirements are large and there are physical limits, such as cavitation, that place an upper bound on what is achievable with these signals. A common workaround is to transmit a lower-power signal with large time-bandwidth product (such as linearly-swept FM chirps). Upon reception, such signals undergo *pulse compression* or *matched filtering* to achieve good range resolution. The operation of pulse compression is equivalent to computing the cross-correlation between the received and transmitted signal. This is represented in the frequency domain by multiplying the received signal with the conjugate of the transmitted signal spectrum:

$$E(\omega, \mathbf{x}) = |A(\omega)|^2 \int \frac{f(\mathbf{x}') W(\omega, \mathbf{x} - \mathbf{x}')}{(4\pi)^2 |\mathbf{x} - \mathbf{x}'|^2} \exp\{-j2k \sqrt{\mathbf{x} - \mathbf{x}'}\} d^3 \mathbf{x}'. \quad (2.17)$$

Taking the inverse temporal Fourier transform gives the pulse-compressed model:

$$E_c(t, \mathbf{x}) = \int f(\mathbf{x}') \frac{A_c\left(t - \frac{2|\mathbf{x} - \mathbf{x}'|}{c}\right)}{(4\pi)^2 |\mathbf{x} - \mathbf{x}'|^2} *_t W(t, \mathbf{x} - \mathbf{x}') d^3 \mathbf{x}', \quad (2.18)$$

in which  $A_c(t)$  simply represents the autocorrelation of the transmitted pulse. The pulse is often idealized as having a flat spectrum in which case  $A_c(t)$  is a sinc function whose main lobe -3 dB width is  $c/2B_w$  where  $B_w$  is the bandwidth. If the transmit pulse is a tone burst which requires no pulse compression, then (2.15) is the appropriate model.

At this point it is customary to absorb all the amplitude terms into the reflectivity term  $f(\mathbf{x}')$ . In other words, the amplitude terms are generally ignored in most derivations. They

are retained for a while longer, however, as doing so helps to point out where the simplifying assumptions are made in developing the reconstruction techniques. In practice, it is rarely worth the complication and expense of accounting for the amplitude terms; synthetic aperture imaging is almost completely a problem of properly accounting for the phase.

## 2.6 Spatiotemporal Domain Image Reconstruction

The classical approach to synthetic aperture image reconstruction is presented next. Here is employed the term *spatiotemporal domain* reconstruction to distinguish this approach from that in the Fourier domain, but the term *time delay and sum* is also common as is *backprojection*. The scheme is straightforward. Each point in the scene being imaged corresponds to a particular locus of echo returns in the observed data. In order to compute the value of a single point in the reconstructed image, all that is required is to integrate the data along this locus after multiplying by the complex conjugate of the expected locus as given by the system model. This operation has the form of an inner product, and the reconstructed image can be thought of as resulting from a spatially-varying correlation operation. To obtain the equation for the spatiotemporal reconstruction, consider the ideal locus for a single unit-amplitude point scatterer position at  $\hat{\mathbf{x}}$ . The hat notation is used to indicate either the location at which the reflectivity is estimated or the estimated reflectivity itself. This is simply (2.15) into which has been substituted  $\delta\left(t - \frac{2|\mathbf{x}-\hat{\mathbf{x}}|}{c}\right)$ , an infinite-bandwidth transmitted pulse, for  $A\left(t - \frac{2|\mathbf{x}-\hat{\mathbf{x}}|}{c}\right)$  and  $f(\hat{\mathbf{x}})$  has been set equal to one:

$$E_{\delta}(t, \mathbf{x}, \hat{\mathbf{x}}) = \frac{\delta\left(t - \frac{2|\mathbf{x}-\hat{\mathbf{x}}|}{c}\right)}{(4\pi)^2|\mathbf{x} - \hat{\mathbf{x}}|^2} W(t, \mathbf{x} - \hat{\mathbf{x}}). \quad (2.19)$$

The kernel for the image reconstruction integral, denoted as  $\overline{E}_{\delta}(t, \mathbf{x}, \hat{\mathbf{x}})$ , is formed from (2.19) by taking its reciprocal. Observe that if any of the quantities are complex, the result is the reciprocal amplitude multiplied by a complex exponential with negative phase. Next, the desired inner product is formed that enables the computation of an estimate  $\hat{f}(\hat{\mathbf{x}})$  of the

reflectivity at the point  $\hat{\mathbf{x}}$  on the sea floor:

$$\begin{aligned}\hat{f}(\hat{\mathbf{x}}) &= \int \bar{E}_\delta(t, \mathbf{x}, \hat{\mathbf{x}}) E_c(t, \mathbf{x}) d^3 \mathbf{x} dt \\ &= \int \bar{E}_\delta(t, \mathbf{x}, \hat{\mathbf{x}}) \int f(\mathbf{x}') \frac{A_c \left( t - \frac{2|\mathbf{x}-\mathbf{x}'|}{c} \right)}{(4\pi)^2 |\mathbf{x}-\mathbf{x}'|^2} *_t W(t, \mathbf{x}-\mathbf{x}') d^3 \mathbf{x}' d^3 \mathbf{x} dt.\end{aligned}\quad (2.20)$$

The next step is to simplify this expression by exploiting the sifting property of the delta function when it appears inside an integral. The function  $\bar{E}_\delta(t, \mathbf{x}, \hat{\mathbf{x}})$  is nonzero only along some path in 3D space, which is denoted as  $L$ . The symbol  $d\ell$ , where  $\ell = \ell(t, \mathbf{x})$ , indicates a differential element along the curve  $L$ . Making this substitution gives:

$$\hat{f}(\hat{\mathbf{x}}) = \int \bar{E}_\delta(\ell) E_c(\ell) d\ell. \quad (2.21)$$

For the usual ideal case in which the sensor trajectory is a straight line, the locus  $L$  is a hyperbola.

The proceeding development of the spatiotemporal image reconstruction is general with respect to the trajectory that may be realized by the sensor. Normally, the sensor is assumed to fly in a straight line. If the sensor boresight is perpendicular to the collection path, then the synthetic aperture length  $L$  is determined by the sensor's beamwidth. The usual practice is to set  $L$  equal to the -3 dB beamwidth. Then for a rectangular aperture,  $L = R\theta_{3\text{dB}} = R\lambda/D$ , where  $R$  is the range under consideration. This is known as stripmap mode imaging. The cross-range (or along-track) resolution can be determined as follows: the cross-range spatial bandwidth (see [8]) is given by

$$\begin{aligned}B_{k_y} &= 2k \sin \theta_{3\text{dB}} \\ &\approx 2k \frac{\lambda}{D} \\ &= \frac{4\pi}{D}.\end{aligned}\quad (2.22)$$

The resolution  $\Delta y$  is inversely proportional to bandwidth:  $\Delta y = 2\pi/B_{k_y} = D/2$ . Here is one of the primary advantages of synthetic aperture processing. The cross-range resolution is independent of range. This is because the processing uses the same spatial bandwidth regardless of range (that is, it integrates over a constant angle).

## 2.7 Fourier-Based Image Reconstruction

The most widely-used Fourier-based image reconstruction technique is derived in this section. It is usually called the *wavenumber algorithm*, but is also commonly referred to as the  $\omega$ - $k$ , *range migration*, and *seismic migration algorithm*. The wavenumber algorithm first found its way into the synthetic aperture literature in the late 1980s [31, 32, 33, 34]. Prior to that, it was used in the geophysics community to form images of the Earth's interior structure [35, 36].

Up to this point, it has not been necessary to make assumptions about the nature of the geometry of the synthetic aperture. Four key simplifications are now made in order to render the mathematics tractable for the derivation. First, the sea floor is modeled as a plane, or delta-sheet, of scatterers of the type described above located at  $z = 0$  and centered on the  $x$ - $y$  origin. Secondly, the sensor trajectory is modeled as a perfectly straight line with constant altitude  $z_0$  and constant negative offset  $-x_0$  from the  $x$ -axis. It is thus written as  $(x, y, z) = (-x_0, y, z_0)$ . Thirdly, the sensor is considered to be stationary during transmission and reception at each position  $y$  (that is, the start-stop assumption is employed). This simplifies the geometry of the problem. The fourth item is to change from two-way spherical ( $1/R^2$ ) spreading to one-way ( $1/R$ ) spreading (or to simply ignore the spreading terms altogether). The one-way spreading assumption is made in order to make the problem compatible with (2.26).

The third and fourth simplifications can be thought of as consequences of the *exploding source* model in which the field of point scatterers is considered to emit the observed signals directly instead of reflecting them after being emitted from the true source. This model, described by Claerbout [36], involves treating the sound speed as  $c/2$  instead of  $c$ . It also implies the use of the Born approximation for scattered fields, as it disallows the possibility of multiple scattering. Strictly speaking, the exploding source model exactly describes the problem solved in this section: The simplifications above are made in order to make synthetic aperture reconstruction fit the solution given below.

The derivation begins with (2.17) and follows by applying the above assumptions. First replace  $f(x, y, z)$  with  $f(x, y)\delta(z)$ , and substitute  $[(-x_0 - x') + (y - y') + z_0]$  for  $(x - x')$  to get:

$$E(\omega, y) = |A(\omega)|^2 \iint \frac{f(x', y') \exp \left\{ -j2k \sqrt{(-x_0 - x')^2 + (y - y')^2 + z_0^2} \right\}}{(4\pi)^2 \sqrt{|(-x_0 - x')^2 + (y - y')^2 + z_0^2|}} dx' dy'. \quad (2.23)$$

Looking at the left hand side of this equation, it is clear that the problem is two-dimensional.

To simplify further, the following change of variables is made:

$$x'_s = -x_0 - \sqrt{(-x_0 - x')^2 + z_0^2}, \quad (2.24)$$

which represents the transformation from ground range and altitude to slant range. The model then becomes:

$$E(\omega, y) = |A(\omega)|^2 \iint \frac{f(x'_s, y') \exp \left\{ -j2k \sqrt{(-x_0 - x'_s)^2 + (y - y')^2} \right\}}{(4\pi)^2 \sqrt{|(-x_0 - x'_s)^2 + (y - y')^2|}} dx'_s dy'. \quad (2.25)$$

Since there has been a change of variables under the integral, there should be an extra term representing the Jacobian of the transformation. This is an amplitude term only, and it is neglected in order to proceed with the derivation. At this stage, a 2-D result is employed whose 3-D equivalent is known as Weyl's identity. It expresses the 2-D Green's function as an expansion in terms of plane waves:

$$\begin{aligned} \frac{-j}{4\pi} \int \frac{\exp \left\{ -j|x| \sqrt{4k^2 - k_y^2} + jk_y y \right\}}{\sqrt{4k^2 + k_y^2}} dk_y &= -\frac{j}{4} H_0^{(2)}(2k\rho) \\ &\approx \frac{1}{\sqrt{j16\pi k\rho}} \exp \{-j2k\rho\}, \end{aligned} \quad (2.26)$$

where  $H_0^{(2)}(2k\rho)$  is the zeroth-order Hankel function of the second kind and the distance between source and observer is  $\rho = \sqrt{(-x_0 - x_s)^2 + (y - y')^2}$ . The second line of (2.26) applies in the limit as  $k\rho \rightarrow \infty$ , which is a suitable approximation for practical imaging



frequencies and geometries. Making the required substitution gives:

$$E(\omega, y) = \frac{1}{(4\pi)^2} \sqrt{\frac{k}{j\pi}} |A(\omega)|^2 \int \int f(x'_s, y') \int \frac{\exp \left\{ -j(x_0 + x'_s) \sqrt{4k^2 - k_y^2} + jk_y(y - y') \right\}}{\sqrt{4k^2 + k_y^2}} dk_y dx'_s dy', \quad (2.27)$$

which can be rearranged to give:

$$E(\omega, y) = \frac{1}{(4\pi)^2} \sqrt{\frac{k}{j\pi}} |A(\omega)|^2 \int \frac{\exp \left\{ -jx_0 \sqrt{4k^2 - k_y^2} + jk_y y \right\}}{\sqrt{4k^2 + k_y^2}} \int \int f(x'_s, y') \exp \left\{ -jx'_s \sqrt{4k^2 - k_y^2} - jk_y y' \right\} dx'_s dy' dk_y. \quad (2.28)$$

The inner double integral over  $(x'_s, y')$  is found to be the 2-D spatial Fourier transform of the scene reflectivity:

$$E(\omega, y) = \frac{1}{(4\pi)^2} \sqrt{\frac{k}{j\pi}} |A(\omega)|^2 \int \frac{\exp \left\{ -jx_0 \sqrt{4k^2 - k_y^2} + jk_y y \right\}}{\sqrt{4k^2 + k_y^2}} F \left( \sqrt{4k^2 - k_y^2}, k_y \right) dk_y. \quad (2.29)$$

The Fourier transform of both sides can be taken with respect to  $y$  leaving:

$$E(\omega, k_y) = \frac{1}{(4\pi)^2} \sqrt{\frac{k}{j\pi}} |A(\omega)|^2 \frac{\exp \left\{ -jx_0 \sqrt{4k^2 - k_y^2} \right\}}{\sqrt{4k^2 + k_y^2}} F \left( \sqrt{4k^2 - k_y^2}, k_y \right). \quad (2.30)$$

This equation can, at last, be solved explicitly for the wavenumber spectrum in terms of the 2-D Fourier transform of the collected data:

$$F(k_x, k_y) = \mathcal{S}^{-1} \left\{ \frac{(4\pi)^2}{|A(\omega)|^2} \sqrt{\frac{j\pi}{k}} \sqrt{4k^2 + k_y^2} \exp \left\{ jx_0 \sqrt{4k^2 - k_y^2} \right\} E(\omega, k_y) \right\}, \quad (2.31)$$

where the following nonlinear mapping is defined:  $\mathcal{S}^{-1}$  as  $k_x = \sqrt{4k^2 - k_y^2}$ . This mapping is required in order to obtain the samples of the left side of (2.31) which are located on a regularly-spaced rectangular grid (and thus amenable to the inverse FFT). This mapping

is known as *inverse Stolt mapping* and is implemented in practice via interpolation in the frequency domain. Prior to inverse Stolt mapping the right hand side of (2.31), the samples of the observed data reside on an annular grid whose radius is  $2k_0$ , which is the wavenumber at the center (or carrier) frequency of the temporal signal.

In (2.31) the transmitted spectrum is explicitly deconvolved in the Fourier domain. This frequency-domain division operation cannot be relied upon to be stable. It is generally ignored in most derivations: a typical scenario is to assume a flat spectrum and ignore the  $|A(\omega)|^2$  altogether. For wideband systems however, the transmitted spectrum may not be sufficiently flat over the entire bandwidth. In this case, division by  $|A(\omega)|^2$  could be replaced by a spectral whitening operation if desired.

In practice, there are subtleties of the implementation of the wavenumber algorithm. These details are not relevant to the mathematical development but pertain mainly to the fact that efficient signal processing is done using baseband data. The detailed implementation of the wavenumber algorithm is given by Hawkins [8]. Another good discussion of the wavenumber algorithm in the SAS literature is presented by Callow [37].

## **2.8 Comments**

### **2.8.1 Reconstruction Expense**

Strictly speaking, the spatiotemporal method of image reconstruction is probably the best way to create a SAS image. Its only real drawback is that it is computationally demanding. By comparison, the wavenumber algorithm is very inexpensive. Its efficiency is nothing short of revolutionary, as it allows large quantities of imagery to be produced in a reasonable time using ordinary desktop computers. In the absence of any aberrations (motion error, medium instability, *etc.*), the wavenumber and spatiotemporal schemes produce essentially identical imagery. Neither is superior from a mathematical point of view, and both preserve phase which is useful for interferometric applications.

A simple analogy for the distinction between these two algorithms is the relationship between the discrete Fourier transform (DFT) and the fast Fourier transform (FFT). The

efficiency of the wavenumber algorithm is orders of magnitude better than spatiotemporal backprojection. Also, just as the FFT is not an approximation to the DFT, the wavenumber algorithm is not an approximation to spatiotemporal backprojection. Rather, the limitation to both the FFT and to wavenumber processing is the fact that the input data must be sampled in a particular fashion. For the FFT, this means uniform spacing of a one-dimensional signal. The wavenumber algorithm requires uniform sampling in two dimensions. It also requires that the data be collected along a perfectly straight sensor trajectory. This rarely happens, so one is forced to compensate the data ahead of time to make it appear as if it was collected in this way. This reality is the motivation for the material in Chapters 5 and 6.

From a practical standpoint, it is easy to build any required corrections, such as motion compensation, into the spatiotemporal scheme as each pixel is reconstructed explicitly from its own unique phase and amplitude history. On the other hand, the wavenumber algorithm is inflexible. With the exception of a few relatively minor effects (examples are given by Callow [37]), it is impossible to alter the algorithm to compensate for most aberrations. Thus, the typical approach is to adjust the raw synthetic aperture data before presenting it to the wavenumber algorithm. The drawback to this is that there are limits to the type/amount of adjustment the data can tolerate. Assuming the vehicle carrying the sensor is stable and the environmental conditions are not severe, wavenumber-based imaging performs as well in terms of image quality as the spatiotemporal approach.

There are other image reconstruction techniques in addition to the two mentioned above. These are detailed in the literature and generally fall into two categories. The first group consists of those conceived prior to the wavenumber algorithm. These are often based on some sort of approximation resulting in inferior imagery with no real computational advantage (see Chapter 4 of [8]). The second group of algorithms postdate the wavenumber method and are often aimed at achieving the flexibility of spatiotemporal backprojection while preserving the computational efficiency of wavenumber reconstruction. An example of this group is fast factorized backprojection [38, 39, 40, 41].

## 2.8.2 Spatial Sampling Concerns

A discussion of synthetic aperture sonar cross-range sampling and its implications is given by Hawkins [8] and earlier by Rolt [13]. It is common to design systems with  $D/2$  sampling, as this is generally considered to be the minimum spatial sampling rate that yields good quality imagery. The  $D/2$  sampling rate is somewhat analogous to the Nyquist rate for temporal signals. There is a serious distinction to be made, however: The Nyquist rate is predicated on the existence of a bandlimited signal. The cross-range spatial signal of synthetic aperture imaging is not bandlimited; it only approximates this condition. Thus, the  $D/2$  sampling rate causes spatial aliasing in even the best circumstances. In the presence of motion errors, the array can become yawed causing the directivity pattern to further amplify spatial frequencies that are already aliased by  $D/2$  sampling. Thus, constructing a SAS system to sample at  $D/3$  or  $D/4$  would enable it to produce imagery that is more tolerant to errors in the array pointing angle.

This discussion of the effects of spatial sampling implies that the projector and receivers have the usual sinc-like directivity pattern associated with a rectangular aperture. To be truly spatially bandlimited, either the projector or the receiver elements would have to possess no sidelobes. Parametric sources [42, 3, 43] have this property, and they have been used in experimental SAS configurations for generating narrow beams at low frequencies (for example, [44, 45, 46, 47]). If the projector were designed to provide much less illumination at angles outside the -3 dB beamwidth, then the  $D/2$  sampling would be more suitable. Indeed, this approach is probably preferable to oversampling the array because the cost associated with the improved projector would be expected to be less than the cost of adding more receivers and the required electronics. The parametric source would be useful in this respect, but these devices generally have low source levels and high power consumption and are thus inefficient.

## CHAPTER 3

### EFFECTS OF UNCOMPENSATED PLATFORM MOTION

#### 3.1 Introduction

It is often of interest to consider how uncompensated platform motion can degrade the ideal point scatterer response (PSR). Such knowledge can be used to shape the design of an AUV-based SAS system. (A discussion of this problem in the context of real aperture sonar is given by Zehner [48].) Also, knowing how certain types of motion affect a SAS image can reduce the time spent in troubleshooting estimation and compensation schemes. In the field of spotlight mode SAR, the effects of phase errors across the synthetic aperture are relatively easy to describe and are well-documented (see, for example, Chapter 5 of Carrara *et al.* [24]). However, the counterpart problem for the stripmap mode is less well-developed in the literature. Some coverage exists in the SAR literature [49, 50, 51, 52, 53, 54], but it is not as thorough as the treatment of spotlight errors. The goal of this chapter is to extend the analysis of phase error effects for stripmap mode synthetic aperture imaging.

From a qualitative standpoint, the ability to associate a given type of phase error with a characteristic degradation of the PSR is useful in diagnosing image quality problems when troubleshooting SAS hardware and processing algorithms. For example, high-frequency random phase errors result in an overall loss of contrast (that is, a reduction in peak to average sidelobe ratio), but they do not degrade the resolution ( $-3$  dB width of the PSR mainlobe). On the other hand, periodic errors tend to create replica, or ghost, copies of the PSR. Furthermore, the quantitative analysis of phase error effects is useful in the design process. An error budget can be constructed that allows the designer to allocate resources in ways aimed at reducing the most serious forms of degradation. Similar error budgets can be used for developing the processing algorithms, tuning any relevant vehicle control software, and for optimizing the operational procedures for deploying the SAS.

### 3.2 Origin of Stripmap Phase Error Effects

For SAS, the primary source of phase errors is unwanted platform motion. Even when the motion has been estimated and compensated, there are inevitably residual phase errors in the data. From a hardware standpoint, phase errors might be introduced by a lack of proper calibration of the hydrophone array or by timing problems in the ping triggering or in the data acquisition system. Another important source of phase errors is the range-varying nature of the motion compensation as discussed in Chapter 6. The analysis presented below is general and applies to all stripmap phase errors regardless of origin.

Unwanted motion causes delays in the observed signal. If these delays exceed a range resolution cell, then the synthetic aperture focusing will generally be poor. For this reason, the discussion is restricted to the simpler case of motion errors whose magnitude is less than a resolution cell. Such errors can be modeled as an additional phase term in the observed signal. The analysis is furthermore restricted to a single point scatterer located at  $(x, y) = (x_0, 0)$  which then creates a planar geometry defined by the nominal sensor trajectory (prescribed here to coincide with the  $y$ -axis) and the scatterer of interest. Within this plane, cross-range focusing can be thought of as a spatially-varying two-dimensional correlation operation (see Section 2.6). Here the problem will be treated as one-dimensional by assuming that the curvature of the hyperbolic range migration of the received signal is small or has been removed. Then, the cross-range focusing operation reduces to correlating the observed signal with the ideal signal. This operation is given by:

$$\begin{aligned} f(y) &= \exp\{-j2kR'(y)\} \odot_y \exp\{-j2kR(y)\} \\ &= \int_{-\frac{L_{SA}}{2}}^{\frac{L_{SA}}{2}} \exp\{-j2kR'(u)\} \\ &\quad \cdot \exp\{+j2kR(u-y)\} du, \end{aligned} \tag{3.1}$$

where

$$\begin{aligned} R(y) &= \sqrt{x_0^2 + y^2} \\ R'(y) &= \sqrt{(x_0 - \epsilon(y))^2 + y^2}, \end{aligned}$$

$L_{\text{SA}}$  is the length of the synthetic aperture, and  $\odot_y$  represents correlation with respect to the variable  $y$ . The quantity  $R$  is the ideal range (that is,  $R$  is referenced to the nominal trajectory),  $R'$  is the actual range from the sensor to the scatterer at cross-range position  $y = 0$ , and  $\epsilon$  is the motion error. Also, the directivity, or beam pattern, has been ignored. In keeping with the small magnitude error assumption,  $R'$  can be expanded about the point  $\epsilon(y) = 0$  as follows:

$$\begin{aligned} R' &\approx R + \left. \frac{dR'}{d\epsilon} \right|_{\epsilon=0} \epsilon \\ &\approx R - \left( 1 - \frac{y^2}{2x^2} \right) \epsilon \\ &\approx R - \epsilon, \end{aligned} \tag{3.2}$$

where the last simplification is based on the condition that  $y \ll x$ . This is equivalent to assuming that the sensor has a narrow beam, in which case the width of a given illuminated interval of  $y$  would be small compared to its range  $x$ . This also implies that the delay induced by the motion error is the same for all scatterers in the beam. Note also that the analysis is carried out only for a single frequency  $k_0$ . This monochromatic representation is a suitable approximation for relatively narrowband systems. The extension to finite (nonzero) bandwidth signals is addressed in Section 3.4.

Equation (3.2) allows the observed phase history from the right hand side of (3.1) to be written as:

$$\exp \{-j2k_0 R'(y)\} = \exp \{-j2k_0 (R(y) - \epsilon(y))\}. \tag{3.3}$$

This indicates that small motion errors simply introduce a phase distortion into the observed signal. The significance of this result is that it can be used to make quantitative statements about the behavior of the cross-range PSR as a function of the error  $\epsilon(y)$ . Note that (3.3) is the product of the ideal and error signals. For convenience, the phase-error function is denoted as  $g(y) = \exp \{+j2k_0 \epsilon(y)\}$  and the ideal signal history as  $s(y) = \exp \{-j2k_0 R\}$ . Taking the Fourier transform with respect to the spatial variable  $y$  causes the transforms of these two functions to be related by convolution rather than multiplication. Thus, the PSR

for any  $g(y)$  is given by the following expression:

$$f(y) = \mathcal{F}_{k_y}^{-1} \left\{ \left[ S(k_y) *_{k_y} G(k_y) \right] \overline{S}(k_y) \right\}. \quad (3.4)$$

The spectrum of the ideal signal phase history  $S(k_y)$  is first convolved (denoted by  $*_{k_y}$ ) with the spectrum of the error function  $G(k_y)$ , after which it is multiplied by the complex conjugate of  $S(k_y)$  to effect the correlation operation representing cross range signal compression. An inverse Fourier transform in  $k_y$  yields the corrupted PSR.

### 3.3 Phase Error Analysis

This section presents several canonical examples of synthetic aperture phase errors. The linear, quadratic, and cubic terms may be used to represent more general low-frequency phase errors via truncated Taylor series. The sinusoidal case is presented because it is commonly encountered and can be used to synthesize more complicated phase error functions. Random white noise phase errors arise in practice as a result of hardware problems, medium fluctuations, or errors in the motion estimation/compensation. Another case presented is that of a phase error represented by the sawtooth function. This case is common in SAS imaging, as it arises when a vernier array experiences a static yaw relative to the nominal platform trajectory.

Polynomial errors with terms higher than third order are not considered. According to Carrara *et al.* [24], the quartic and higher-order terms fall into the same category of sinusoids with multiple cycles per synthetic aperture: That is, they are high-frequency errors, and the purpose of the Taylor series analysis is to examine low-frequency phase error functions. Low-frequency phase errors tend to primarily distort the main lobe of the PSR, while high-frequency errors affect the sidelobe structure. The nature of this sidelobe distortion can take several forms, as is shown below.



### 3.3.1 The Principle of Stationary Phase

The principle of stationary phase (PSP) is a technique for finding approximate closed-form solutions to integrals in which the integrand has a large phase variation and a slowly-varying envelope,  $g_0(u)$ . It is frequently discussed in radar and optics texts (for example, [55] and [56]). The expression for the PSP solution is:

$$\int g_0(u) \exp\{j\phi(u)\} du \approx \sqrt{\frac{j2\pi}{\phi''(u^\circ)}} g_0(u^\circ) \exp\{j\phi(u^\circ)\}, \quad (3.5)$$

where  $u^\circ$  is known as the stationary point and is the solution of  $\phi'(u^\circ) = 0$ . The PSP is the primary tool used to carry out the analysis of this chapter.

### 3.3.2 Linear Phase Error

First consider the specific example of a linear motion error:  $\epsilon(y) = \alpha y$  and  $g(y) = \exp\{j2k_0\alpha y\}$ , where  $\alpha$  is a dimensionless scale factor. The Fourier transform of this function is  $G(k_y) = \delta(k_y - 2k_0\alpha)$ . The effect in the  $k_y$  domain of a linear phase error is thus seen to be a shift of the ideal observed signal spectrum. If  $S(k_y)$  is bandlimited, then it is possible to choose  $\alpha$  such that  $S(k_y - 2k_0\alpha)$  has no spectral support in common with  $S(k_y)$ . Roughly speaking, the resolution of the PSR  $f(y)$  will be determined by the spatial bandwidth of this common region of support. The detailed structural appearance of  $f(y)$  will also be determined by the phase of  $F(k_y)$ .

A closed-form expression for the distorted PSR can be obtained when the slope  $\alpha$  of the linear phase error is small. The expression for the Fourier transform of the cross-range signal is:

$$\begin{aligned} S(k_y) &= \sqrt{\frac{\pi x_0}{jk_0}} \text{rect} \left\{ \frac{k_y}{\sin \theta_{3\text{dB}} \sqrt{4k_0^2 + k_y^2}} \right\} \\ &\quad \cdot \exp \left\{ -j \left( \sqrt{4k_0^2 - k_y^2} \right) x_0 \right\} \\ &\approx \sqrt{\frac{\pi x_0}{jk_0}} \text{rect} \left\{ \frac{k_y}{2k_0 \sin \theta_{3\text{dB}}} \right\} \\ &\quad \cdot \exp \left\{ -j2k_0 x_0 + j \frac{k_y^2}{4k_0} x_0 \right\}, \end{aligned} \quad (3.6)$$

where  $\theta_{3\text{dB}}$  is the beamwidth of the sensor. The first expression for  $S(k_y)$  was derived [25] using the principle of stationary phase, which is summarized in Section 3.5. The approximate expression for  $S(k_y)$  was obtained by making a parabolic approximation to  $R$  in (3.1) and using the PSP to take the Fourier transform. The quadratic approximation to  $S(k_y)$  is used below to facilitate the present analysis. It can be substituted into (3.4) to obtain the following:

$$\begin{aligned}
f(y) &= \mathcal{F}_{k_y}^{-1} \left\{ \left[ A(k_y) \exp \left\{ j \frac{k_y^2}{4k_0} x_0 \right\} *_{k_y} \delta(k_y - 2k_0\alpha) \right] \right. \\
&\quad \left. \cdot \bar{A}(k_y) \exp \left\{ -j \frac{k_y^2}{4k_0} x_0 \right\} \right\} \\
&= \mathcal{F}_{k_y}^{-1} \left\{ A(k_y - 2k_0\alpha) \exp \left\{ j \frac{(k_y - 2k_0\alpha)^2}{4k_0} x_0 \right\} \right. \\
&\quad \left. \cdot \bar{A}(k_y) \exp \left\{ -j \frac{k_y^2}{4k_0} x_0 \right\} \right\} \tag{3.7} \\
&= \mathcal{F}_{k_y}^{-1} \left\{ A(k_y - 2k_0\alpha) \bar{A}(k_y) \exp \left\{ j (k_0\alpha^2 - \alpha k_y) x_0 \right\} \right\} \\
&\approx \mathcal{F}_{k_y}^{-1} \left\{ |A(k_y)|^2 \exp \left\{ -j \alpha x_0 k_y \right\} \right\} \\
&= f(y - \alpha x_0).
\end{aligned}$$

The approximation step in (3.7) was made by neglecting  $\alpha$  in the amplitude term and  $\alpha^2$  in the phase term. It can be seen that the effect of a linear phase error is to shift the  $k_y$  spectrum of the received signal. If this shift is small ( $\alpha < 1$ ), then the result of this spectral shifting is simply a shift in the cross-range position of the ideal PSR. This is the central result of this section: spectral shifts  $2k_0\alpha$  translate into spatial shifts  $\alpha x_0$  of the PSR. This fact is used extensively below in order to derive closed-form estimates of PSR deterioration.

### 3.3.3 Quadratic Phase Error

The quadratic phase error is generally considered to be the dominant cause of cross-range defocus in synthetic aperture imagery. Apart from representing the second-order component of the Taylor series expansion of a given motion, the quadratic phase error can be used to model the effect of measurement errors in the wave propagation speed used for image

reconstruction. As such, analysis of the quadratic error is also linked to the notion of depth of focus.

The quadratic phase-error term is given by  $g(y) = \exp\{j2k_0\nu y^2\}$  where  $\nu$  is a scale factor having units of reciprocal length. Its Fourier transform obtained using the PSP is:

$$G(k_y) = \sqrt{\frac{j\pi}{2k_0\nu}} \text{rect}\left(\frac{k_y}{4k_0\nu L_{SA}}\right) \exp\left\{-\frac{jk_y^2}{8k_0\nu}\right\}. \quad (3.8)$$


The error  $g(y)$  has support on  $y \in [-L_{SA}/2, L_{SA}/2]$ . The exact and approximate quadratic phase error spectra  $G(k_y)$  are shown in Figure 3.1(a).

Using the reasoning upon which (3.7) is based, the degradation of the PSR can be estimated if it is assumed that the width of the rectangle function appearing in (3.8) is small. Then, the actual PSR can be thought of as a continuous smear of copies of the ideal PSR. The width of the actual PSR is therefore determined by the width of the rectangle function in (3.8). Consequently, the fundamental type of degradation caused by the quadratic phase error is a broadening of the main lobe of the ideal PSR. In other words, the quadratic phase error causes a loss of cross-range resolution. The expansion of the -3 dB width of the PSR can be roughly estimated by the bandwidth of the rectangle function in (3.8) using the result of the approximation in (3.7). This expansion is found to be equal to  $\Delta_{PSR} = 2|\nu|L_{SA}x_0 = 2|\nu|x_0^2 \sin \theta_{3dB}$ . It is clear that for a given value of  $\nu$ , the PSR width is proportional to the square of range. An example of the PSR degradation caused by the quadratic phase error is shown in Figure 3.1(b).


Now, it can be shown that the effect of an error in the knowledge of the wave propagation speed can be modeled as a quadratic phase error. First, let the received signal be given by:

$$\begin{aligned} \exp\{-j2(k_0 + \Delta k)R\} &= \exp\{-j2k_0R\} \cdot \exp\{-j2\Delta kR\} \\ &= s(y)g(y), \end{aligned} \quad (3.9)$$

where the parabolic approximation to the phase term has replaced its true hyperbolic form,  $R \approx x_0 + \frac{y^2}{2x_0}$ , and the sound speed error has been represented by a change in the temporal



(a)



(b)

**Figure 3.1. (a) Comparison of the actual quadratic phase error (with  $vx_0/D = 0.88$ ) spectrum  $G(k_y)$  and the approximation given by (3.8) which was derived using the principle of stationary phase. (b) Example of an ideal PSR and the PSR corrupted by a quadratic phase error.**

wavenumber,  $\Delta k = -\frac{\omega\Delta c}{c(c+\Delta c)}$ . The phase error function  $g(y)$  has a form that is identical to the cross-range signal  $s(y)$ , and (3.6) may therefore be used to obtain  $G(k_y)$  simply by substituting  $\Delta k$  for  $k_0$ . Consequently, the the region of support for  $G(k_y)$  is  $\{k_y : |k_y| \leq \Delta k \sin \theta_{3\text{dB}}\}$ , and the spread of the PSR corrupted by the sound speed error is:

$$\begin{aligned}\Delta_{\text{PSR}} &= \left| \frac{\Delta k}{k_0} \right| x_0 \sin \theta_{3\text{dB}} \\ &= \left| \frac{\Delta c}{c + \Delta c} \right| L_{\text{SA}}.\end{aligned}\tag{3.10}$$

This corresponds to  $\nu = \frac{\Delta c}{2x_0(c+\Delta c)}$  in the result above. The error in  $c$  causes the PSR degradation to become linearly proportional to the range. It can also be seen that narrow beam sensors are less susceptible to image degradation caused by the sound speed error. The implication is that this error can be reduced by increasing  $D$  at the cost of decreased resolution, or by operating at a higher frequency.

Another useful application for this analysis is quantifying the effects of errors in the estimate of the sensor's forward (namely, cross range) velocity. For this situation, the cross range location of the sensor is (wrongly) believed to be at  $y + \chi y$  when it is actually located at  $y$ . Here, the quantity  $\chi$  represents the fractional error in the velocity estimate. That is,  $\chi = (v_{\text{est}} - v_{\text{act}})/v_{\text{act}}$ . Then, the range becomes  $R = \sqrt{x_0^2 + (y + \chi y)^2}$  which is approximated as  $R \approx x_0 + (1 + 2\chi)\frac{y^2}{2x_0}$  and the phase error function is given by:

$$g(y) = \exp \left\{ -\frac{j2k_0\chi y^2}{x_0} \right\}.\tag{3.11}$$

The corresponding Fourier transform is:

$$G(k_y) = \sqrt{\frac{\pi x_0}{j2k_0\chi}} \text{rect} \left( \frac{x_0 k_y}{4k_0\chi L_{\text{SA}}} \right) \exp \left\{ -j \frac{3x_0 k_y^2}{8k\chi} \right\},\tag{3.12}$$

from which the corrupting PSR spread is found to be

$$\Delta_{\text{PSR}} = 2|\chi|L_{\text{SA}}.\tag{3.13}$$

As with the case of erroneous sound speed estimates, forward velocity errors are less harmful for smaller synthetic aperture lengths and thus favor larger sensors and higher frequencies. According to [57], the results of this and the preceding paragraph can be obtained in

an alternative manner using the principles of geometrical optics.

### 3.3.4 Cubic Phase Error

While the quadratic phase error degrades the PSR symmetrically, the cubic phase error term does so in an asymmetric fashion. Both errors decrease the resolution, which is the -3 dB width of the PSR. The PSP is once again used to facilitate the analysis, as it is necessary to know the Fourier transform of  $g(y) = \exp\{j2k_0\rho y^3\}$ , where  $\rho$  is a constant with units of  $1/\text{length}^2$ :

$$G(k_y) = \int_{-L/2}^{L/2} e^{j2k_0\rho y^3} e^{-jk_y y} dy. \quad (3.14)$$

The stationary point of (3.14) is:

$$y^\circ = \pm \sqrt{\frac{k_y}{6k_0\rho}}, \quad (3.15)$$

and a question arises as to which sign is the correct one to take for the present application.

Substituting  $y^\circ$  into the PSP phase term  $\exp\{j\phi(y^\circ)\}$  gives:

$$\exp\{j\phi(y^\circ)\} = \exp\left\{\mp j \frac{2k_y}{3} \sqrt{\frac{k_y}{6k_0\rho}}\right\}. \quad (3.16)$$

If the signs of  $k_y$  and the constant  $\rho$  differ, then the argument of this exponential becomes purely real. Thus, one finds motivation for choosing the positive sign in (3.15) in order to avoid exponential growth in  $G(k_y)$ . The next term to examine in the PSP analysis is the envelope term. As usual, the envelope of  $g(y)$  is a rectangle function centered at  $y = 0$  and having width  $L_{SA}$ . This leads to the following for the envelope of  $G(k_y)$ :

$$\begin{aligned} g_0(y^\circ) &= \text{rect}\left(\frac{y^\circ}{L_{SA}}\right) \\ &= \text{rect}\left(\sqrt{\frac{k_y}{6k_0\rho L_{SA}^2}}\right). \end{aligned} \quad (3.17)$$

In this equation, it must be recognized that the argument of the rectangle function can become complex. To deal with this situation, the definition of this function will be extended so that  $\text{rect}(x)$  equals one if  $x$  is both real and  $|x| \leq 1/2$ . For all other values of  $x$ ,  $\text{rect}(x) = 0$ .

Therefore, (3.17) indicates that the PSP approximation to the spectrum  $G(k_y)$  is one-sided. It is given by:

$$G(k_y) \approx \frac{\sqrt{j\pi}}{(6\rho k_0 k_y)^{1/4}} \text{rect} \left( \sqrt{\frac{k_y}{6k_0 \rho L_{SA}^2}} \right) \exp \left\{ j \sqrt{\frac{k_y}{6k_0 \rho}} \right\}. \quad (3.18)$$

The region of spectral support of  $G(k_y)$  for the cubic phase error function is  $\{k_y : 0 \leq k_y \leq 3k_0 \rho L_{SA}^2/2\}$  if  $\rho > 0$ . If  $\rho$  is negative, then  $G(k_y)$  has support given by  $\{k_y : 3k_0 \rho L_{SA}^2/2 \leq k_y \leq 0\}$ . Applying the reasoning from (3.7) as before, the mainlobe expansion of the PSR due to the cubic phase error function is:

$$\begin{aligned} \Delta_{\text{PSR}} &= \frac{3}{4} \rho L_{SA}^2 x_0 \\ &= \frac{3}{4} \rho x_0^3 \sin^2 \theta_{3\text{dB}}. \end{aligned} \quad (3.19)$$

The stationary phase approximation to  $G(k_y)$  is compared to the actual spectrum in Figure 3.2(a). The resulting form of PSR defocus is shown in Figure 3.2(b). Unlike the quadratic case above, the spread as given by  $\Delta_{\text{PSR}}$  does not fully describe the appearance of the corrupted PSR. Here, the PSR tapers off from the peak value by several decibels before finally reaching the knee, or cutoff, predicted by  $\Delta_{\text{PSR}}$ .


### 3.3.5 Sinusoidal Phase Error

Another interesting case is that of a sinusoidal phase error,  $g(y) = \exp \{j2k_0 \beta \sin(\gamma y)\}$ . The constant  $\beta$  has units of length, and  $\gamma$  has units of reciprocal length. The analysis of this section is largely based on the work by Fornaro [52]. The function  $g(u)$  can be manipulated into a more useful form by employing the generating function for the Bessel function of the first kind:


$$e^{\frac{x}{2}(t - \frac{1}{t})} = \sum_{n=-\infty}^{\infty} t^n J_n(x), \quad (3.20)$$

and making the substitutions  $t = \exp\{j\gamma y\}$  and  $x = 2k_0 \beta$ . The result is:

$$e^{j2k_0 \beta \sin(\gamma y)} = \sum_{n=-\infty}^{\infty} \left( e^{j\gamma y} \right)^n J_n(2k_0 \beta). \quad (3.21)$$



(a)



(b)

**Figure 3.2.** (a) Comparison of the actual cubic phase error spectrum  $G(k_y)$  and the approximation given by (3.18) which was derived using the principle of stationary phase. (b) Example of an ideal PSR and the PSR corrupted by a cubic phase error with  $\lambda/D = 0.31$  and  $\rho L_{SA}^2 = 0.035$ .



Taking the Fourier transform of this with respect to  $y$  gives  $G(k_y)$ :

$$G(k_y) = \sum_{n=-\infty}^{\infty} J_n(2k_0\beta) \delta(k_y - n\gamma). \quad (3.22)$$


Thus, a sinusoidal phase error will result in the creation of replicas of the ideal observed signal spectrum  $S(k_y)$ . These replicas are weighted according to the value of the coefficient  $J_n(2k\beta)$ . Although there are an infinite number of replicas, the coefficients die off quickly. This rate of decay slows with increasing sinusoidal amplitude  $\beta$ , and the replica spacing is determined by the frequency  $\gamma$  of the phase error function. Figure 3.3(a) shows  $G(k_y)$  for a representative value of  $k_0\beta$  and  $B_{k_y}/\gamma$ . These quantities are dimensionless and represent the relative error magnitude compared to the wavelength and the error frequency relative to the spatial frequency bandwidth observable by the synthetic aperture, respectively. (Recall that  $B_{k_y} \approx 4\pi/D = 4k_0 \sin(\theta_{3\text{dB}}/2)$ .) Assuming  $\beta$  and  $\gamma$  are within certain bounds, then the result from Section (3.3.2) may be used to deduce that a sinusoidal phase error results in multiple scaled copies of the ideal PSR:

$$f_{\text{composite}} = \sum_{n=-\infty}^{\infty} J_n(2k_0\beta) f\left(y - \frac{n\gamma}{2k_0} x_0\right) \quad (3.23)$$


These replicas are located at  $y = \pm \frac{n\gamma}{2k_0} x_0$ . Note that the scale factor  $J_n(2k_0\beta)$  depends on  $k_0$ . Figure 3.3(b) shows an example of these results.

### 3.3.6 Sawtooth Phase Error

The preceding analysis can be used to derive the effect of a sawtooth-shaped phase error  $g(y) = \exp\{j2k_0\xi(y)\}$ , where  $\xi(y)$  is the sawtooth function with period  $P$  and peak-to-peak amplitude  $A$ . This case is important because it is frequently, if not universally, encountered in practice. It is common for towed and autonomous SAS sensors to travel in some degree of cross current. In these situations, the host platform usually experiences a static yaw, known as *crab*, in order to keep the desired course over the bottom; finned bodies are less susceptible to this effect than those without fins. The sawtooth phase error appears if multiple receivers are used, and this is nearly always the case for SAS.



(a)



(b)

**Figure 3.3.** (a) Comparison of the actual sinusoidal phase error spectrum spectrum  $G(k_y)$  and the analytical expression given by (3.22). Note that the actual (or experimental)  $G(k_y)$  was computed discretely, and it therefore does not consist of delta functions as does the continuous result. (b) Example of an ideal PSR and the PSR corrupted by a sinusoidal phase error with  $k_0\beta = 0.63$  and  $\gamma/B_{k_y} = 0.069$ .

The peak-to-peak amplitude of the sawtooth error is given by  $A = ND \sin(\theta_{\text{crab}})/2 \approx ND\theta_{\text{crab}}/2$ , where  $N$  is the number of receivers of length  $D$  and  $\theta_{\text{crab}}$  is the static yaw angle. The period of the sawtooth pattern is, of course,  $P = ND \cos(\theta_{\text{crab}})/2 \approx ND/2$ . Actual crab angles are usually fairly small (less than ten degrees), and the small-angle approximation is sufficiently accurate.

The Fourier series representation of the sawtooth wave is:


$$\begin{aligned}\xi(y) &\approx \frac{A}{\pi} \left( \sin \frac{2\pi y}{P} - \frac{1}{2} \sin \frac{4\pi y}{P} + \frac{1}{3} \sin \frac{6\pi y}{P} - \dots \right) \\ &= \frac{A}{\pi} \sum_{m=1}^{\infty} \frac{(-1)^{m-1}}{m} \sin \frac{2m\pi y}{P} \\ &= \frac{ND\theta_{\text{crab}}}{2\pi} \sum_{m=1}^{\infty} \frac{-1^{m-1}}{m} \sin \frac{4m\pi y}{ND}.\end{aligned}\tag{3.24}$$

This infinite sum appears in the exponential term of  $g(y)$ , and  $g(y)$  can therefore be written as an infinite series of products. Recall from the previous section that the Fourier transform of any single term in this product is an infinite sum in the  $k_y$  domain. It thus may be seen that  $G(k_y)$  for the sawtooth case is an infinite series of convolutions of infinite sums:


$$\begin{aligned}G(k_y) &= \mathcal{F}_y \{g(y)\} \\ &= \mathcal{F}_y \left\{ \prod_{m=1}^{\infty} \exp \left\{ j2k_0 \frac{A(-1)^{m-1}}{m\pi} \sin \frac{2m\pi y}{P} \right\} \right\} \\ &= G_1(k_y) *_{k_y} G_2(k_y) *_{k_y} G_3(k_y) \cdots ,\end{aligned}\tag{3.25}$$

where each of the  $G_m(k_y)$  in the last line is the Fourier transform of a single component of (3.24) and has the form of (3.22). An example is shown in Figure 3.4(a). Unfortunately, there is no convenient closed-form expression for the corrupted PSR due to the sawtooth phase error. It is true that the replicas due to any single sinusoidal component generally decay quickly. The same is not true of the amplitudes of the terms of the Fourier series  $\xi(y)$  because of the discontinuity present in the sawtooth function.

The characteristic feature of the sawtooth phase error is that the infinite series of convolutions behaves such that the replicas of the PSR are not symmetrically weighted. As



(a)



(b)

**Figure 3.4. (a) Comparison of the actual sawtooth phase error spectrum spectrum  $G(k_y)$  and the analytical expression given by (3.25). For this example, the parameters are  $D = 0.4$ ,  $\theta_{\text{crab}} = 2^\circ$ , and  $N = 6$ . (b) Example of an ideal PSR and the PSR corrupted by a sawtooth phase error with the same parameters given in Figure 3.4(a).**

one might expect, the sawtooth error causes a dominant replica to occur either to the left or right of the real object. The distance between the replicas induced by the phase error is:

$$y = \frac{2\pi}{NDk_0}x_0. \quad (3.26)$$


Figure 3.4(b) shows how the replicas of the PSR are evenly spaced but unevenly weighted. It is important to point out that it is possible for one of the replicas of the PSR to have an amplitude greater than the true PSR located at  $y = 0$ . It is also worth noting that the asymmetric behavior predicted by the analysis above is due solely to the phase error. The off-broadside pointing angle of the directivity pattern of the array elements was not considered.

### 3.3.7 Random Normally-Distributed Phase Error

The final case to be considered is a random phase error:  $g(y) = \exp\{j\kappa(y)\}$ , in which  $\kappa(y)$  is a normally-distributed zero-mean white random process with variance  $\sigma_\kappa^2$ . Because the process is stationary, the expected value is independent of  $y$ . Consequently, it is expressed as:

$$\begin{aligned} E\{e^{j\kappa(y)}\} &= E\{e^{j\kappa}\} \\ &= \mu_g \\ &= E\{\cos(\kappa)\} + jE\{\sin(\kappa)\} \\ &= \int_{-\infty}^{\infty} \cos(\kappa) \frac{e^{-\kappa^2/2\sigma_\kappa^2}}{\sqrt{2\pi\sigma_\kappa^2}} d\kappa \\ &\quad + j \int_{-\infty}^{\infty} \sin(\kappa) \frac{e^{-\kappa^2/2\sigma_\kappa^2}}{\sqrt{2\pi\sigma_\kappa^2}} d\kappa. \end{aligned} \quad (3.27)$$

A closed-form expression for this expected value was derived by Richards [58] who obtained  $\mu_g = \exp\{-\sigma_\kappa^2/2\}$ . Figure 3.5 shows the result of carrying out this integration numerically for values of  $\sigma_\kappa^2$  ranging from 0.01 to 10. The plot shows the expected value for the real and imaginary parts separately. The expected value of the complex part of the random process is always zero, agreeing with the analytical result.



**Figure 3.5.** Expected value of the real and imaginary components of (3.27) evaluated for values of  $\sigma_\kappa^2$  ranging from 0.01 to 100.

For small variance in the phase noise, the value of (3.27) is close to one. As the variance becomes large, the expected value approaches zero. This makes sense intuitively because the phase error function  $g(y)$  effectively becomes a vector of unit length whose angle in the complex plane is uniformly distributed. In Figure 3.5, the quotient  $\sigma_\kappa/k_0$  can be used to express the variance in terms of units of length. Similarly, the standard deviation in terms of wavelengths is  $\sigma_\kappa/2\pi$ .

Also of interest is the variance of the function  $g(y)$ :

$$\begin{aligned}
 \text{var}\{g(y)\} &= E\{|g(y) - \mu_g|^2\} \\
 &= E\{|e^{jk(y)} - \mu_g|^2\} \\
 &= E\{|e^{j2k(y)}|^2\} - |\mu_g|^2 \\
 &= 1 - \mu_g^2.
 \end{aligned} \tag{3.28}$$

Interestingly, the variance of  $g(y)$  can be expressed solely in terms of its expected value.

The effect of the random phase error can be summarized by the expected value of the

magnitude of the PSR,  $E(|f(y)|^2)$ . This quantity will be obtained using the following facts regarding linear systems with random processes as inputs [4]:

$$R_{pq}(t_1, t_2) = \int_{-\infty}^{\infty} R_{pp}(t_1, t_2 - \alpha) \bar{h}(\alpha) d\alpha \quad (3.29)$$

$$R_{qq}(t_1, t_2) = \int_{-\infty}^{\infty} R_{pq}(t_1 - \alpha, t_2) h(\alpha) d\alpha, \quad (3.30)$$

where  $h(t)$  represents the system impulse response,  $p(t)$  is its input, and  $q(t)$  is its output. The result above enables one to find the autocorrelation of the output using only the autocorrelation of the input and the system impulse response. The quantity of ultimate interest is the expected value of the squared magnitude of the corrupted PSR,  $E(|f(y)|^2)$ : this is equal to the autocorrelation of  $f(y)$  evaluated for zero lag.

In the following analysis, the corrupted received signal  $s_c(y) = s(y)g(y)$  will be taken as the input  $p(t)$  while  $\bar{s}(y)$  plays the role of the linear system impulse response. The output  $q(t)$  above corresponds to the PSR  $f(y)$ . Recall that the ideal PSR can be thought of as the autocorrelation of  $s(y)$ ,  $R_{ss}(y')$  (that is,  $h = 1$ ). The first step is to find the autocorrelation of  $s_c(y)$  employing the fact that  $g(y)$  can be expressed as the sum of its mean and the centered version of itself,  $g(y) = \mu_g + (g(y) - \mu_g) = \mu_g + v(y)$ :

$$\begin{aligned} R_{s_c s_c}(y_1, y_2) &= E\{s_c(y_1) \bar{s}_c(y_2)\} \\ &= E\{s(y_1)g(y_1) \bar{s}(y_2) \bar{g}(y_2)\} \\ &= E\{[\mu_g s(y_1) + v(y_1)s(y_1)] \\ &\quad \cdot [\mu_g \bar{s}(y_2) + \bar{v}(y_2) \bar{s}(y_2)]\} \\ &= \mu_g^2 s(y_1) \bar{s}(y_2) + E\{v(y_1)s(y_1) \bar{v}(y_2) \bar{s}(y_2)\} \\ &= \mu_g^2 s(y_1) \bar{s}(y_2) + \sigma_g^2 \delta(y_2 - y_1). \end{aligned} \quad (3.31)$$

Note that  $s(y)$  does not appear in the last line because  $s(y) \bar{s}(y)$  evaluated at zero is equal to

one. The next step is to compute  $R_{scf}$ :

$$\begin{aligned}
R_{scf}(y_1, y_2) &= R_{scsc}(y_1, y_2) *_{y_2} s(y_2) \\
&= [\mu_g^2 s(y_1) \bar{s}(y_2) + \sigma_g^2 \delta(y_2 - y_1)] *_{y_2} s(y_2) \\
&= \mu_g^2 s(y_1) \bar{f}(y_2) + \sigma_g^2 s(y_2 - y_1),
\end{aligned} \tag{3.32}$$

followed by  $R_{ff}$ :

$$\begin{aligned}
R_{ff}(y_1, y_2) &= R_{scf}(y_1, y_2) *_{y_1} \bar{s}(y_1) \\
&= [\mu_g^2 s(y_1) \bar{f}(y_2) + \sigma_g^2 s(y_2 - y_1)] *_{y_1} \bar{s}(y_1) \\
&= \mu_g^2 f(y_1) \bar{f}(y_2) + \sigma_g^2 f(y_2 - y_1),
\end{aligned} \tag{3.33}$$

where the symbol  $*$  indicates convolution with respect to the subscripted variable. Finally, if  $y_2 = y_1 = 0$ , the desired result is obtained:

$$R_{ff}(0, 0) = E\{|f(y)|^2\} = \mu_g^2 |f(y)|^2 + \sigma_g^2 f(0). \tag{3.34}$$


Note that  $f(0)$  is real and is equal to  $L_{SA}$ :  $f(0) = \int_{-L_{SA}/2}^{L_{SA}/2} s(y) \bar{s}(y) dy$ . It is now possible to quantify the effect of a normally-distributed white random phase on the ideal PSR. It can be seen that the corrupted PSR consists of a copy of the ideal PSR scaled by  $\mu_g^2$  and added to a background value. As the variance of the phase noise increases, the scaled copy of the ideal PSR diminishes in amplitude relative to the mean background noise. The amplitude of the peak of the corrupted PSR relative to the ideal is given in decibels by:

$$\begin{aligned}
\text{Actual to Ideal Peak Ratio} &= 10 \log_{10} \mu_g^2 \\
&= 10 \log_{10} e^{-\sigma_k^2}.
\end{aligned} \tag{3.35}$$

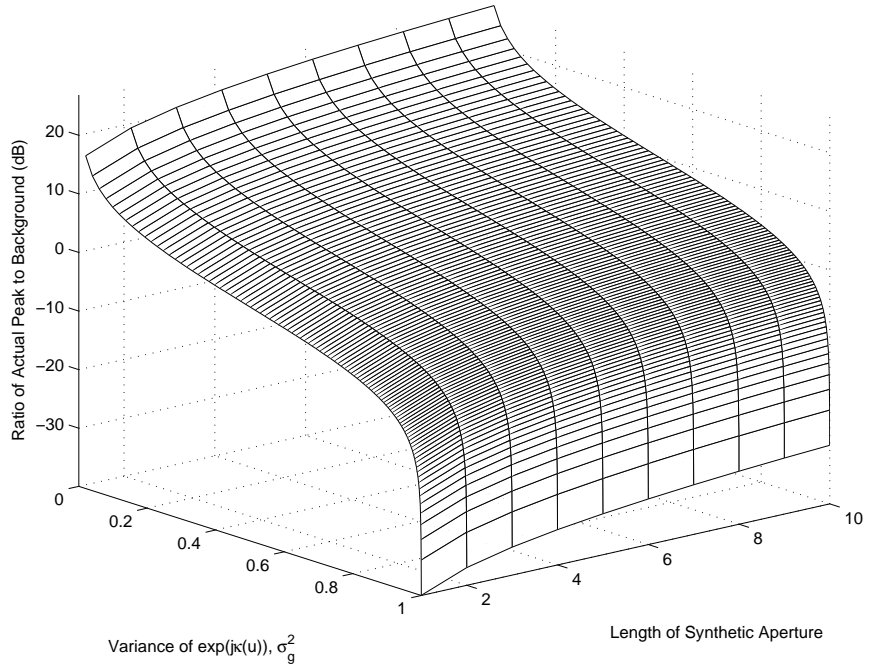
Similarly, the actual peak relative to the background level is given by:

$$\begin{aligned}
\frac{\text{Actual Peak to}}{\text{Mean Background Ratio}} &= 10 \log_{10} \frac{\mu_g^2 L_{SA}}{\sigma_g^2} \\
&= 10 \log_{10} \frac{\mu_g^2 L_{SA}}{1 - \mu_g^2} \\
&= 10 \log_{10} \frac{e^{-\sigma_k^2} L_{SA}}{1 - e^{-\sigma_k^2}}.
\end{aligned} \tag{3.36}$$





(a)



(b)

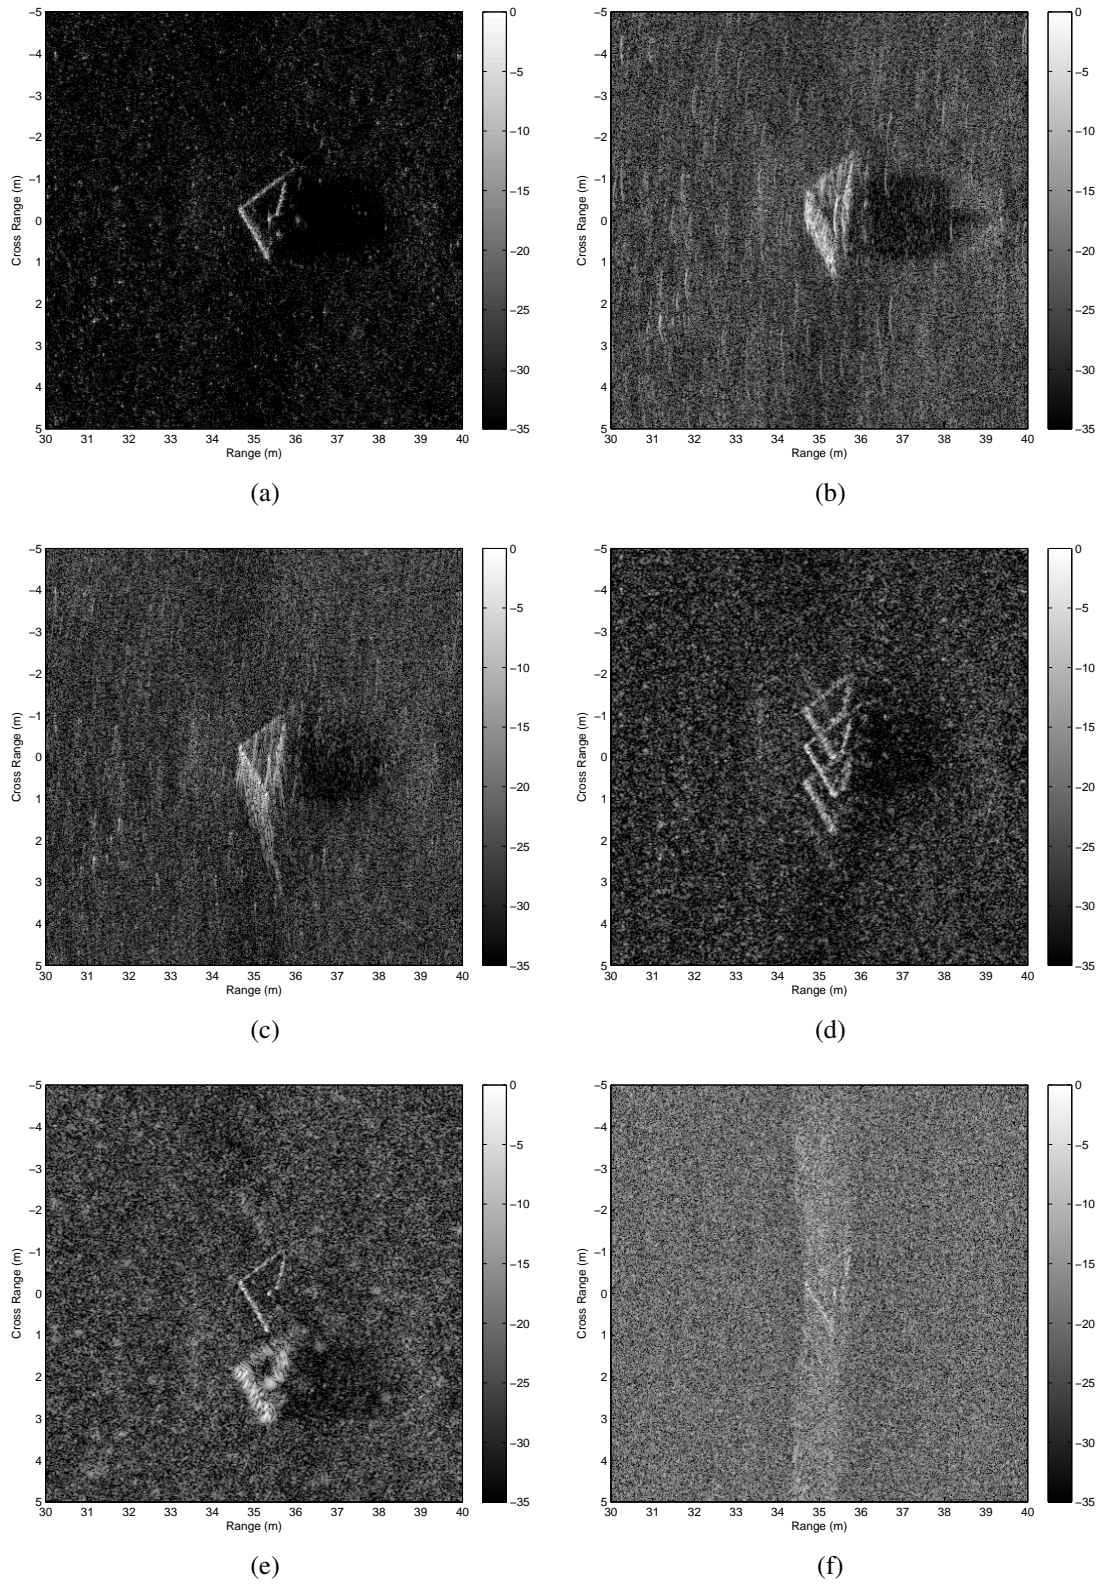
**Figure 3.6. (a) Example of an ideal PSR and the PSR corrupted by a normally-distributed white random phase error whose standard deviation is one quarter of a wavelength. (b) Surface plot depicting the ratio of the actual peak to background ratio (in decibels) as a function of  $\mu_g$  and the synthetic aperture length.**

Equation (3.36) confirms analytically what is well-known experientially: random phase errors result in an overall lack of image contrast but no loss of resolution. For a given noise variance, this effect diminishes with range as the synthetic aperture length grows. Figure 3.6(a) gives an example of an actual PSR and the PSR subject to random phase noise with  $\sigma_\kappa = \lambda/4$ . One thousand realizations were averaged to compute the corrupted PSR. For this example, the theoretical value for the actual peak relative to the ideal is -10.9 dB and the value for the actual peak relative to the background is 28.3 dB. The plot indicates a value slightly less than 28.3 dB for the peak to background ratio. This apparent anomaly is due to the fact that the horizontal extent of the PSR contained in the plot includes the sidelobe structure in addition to the background noise. For cross-range values beyond  $\pm 0.5\text{m}$ , the peak to background ratio approaches the predicted 28.3 dB. Additionally, Figure 3.6(b) depicts the behavior of (3.36) as a function of both  $\mu_g$  and  $L_{SA}$ . The peak to background ratio improves as the synthetic aperture length grows. In practice this means that narrow beam and short range SAS systems are the most susceptible to random phase errors of a given variance. These errors tend to get suppressed due to the averaging that occurs as the aperture length grows.

### 3.4 Results and Discussion

The preceding results are applied to actual SAS field data in Figures 3.7(a)–3.7(f). The test image was first motion compensated in order to obtain proper focus. Then, each of the errors analyzed in Section 3.3 was applied to the data to yield the set of images shown. The parameters used to produce Figure 3.7 are the same as the counterpart examples in Section 3.3. The object shown is a tetrahedron approximately two meters across. The forms of distortion visible in the images are in good agreement with those predicted above.

Recall that the analysis was conducted for a single temporal frequency. The sonar used to collect the field data is considered to be fairly narrowband with a center frequency to bandwidth ratio equal to four. Nevertheless, the finite-bandwidth condition can be modeled




**Figure 3.7. Examples of the effects of the various types of phase errors discussed in Section 3.3: (a) no error; (b) quadratic error; (c) cubic error; (d) sinusoidal error; (e) sawtooth error; (f) random error.**

by computing the degraded PSR for a range of temporal frequencies within a desired band and then summing the results. Figures 3.8(a) and 3.8(b) show two examples of such a computation. The phase error function parameters are the same as those used to produce Figures 3.3(b) and 3.4(b).


### 3.4.1 Stripmap vs. Spotlight Phase Errors

The phase error effects for stripmap imagery are similar to those for the spotlight mode (see Carrara [24]). This similarity is remarkable considering that spotlight phase errors affect the PSR through a completely different mechanism. While stripmap data is collected by accumulating a sequence of signals that are adjacent and parallel, spotlight mode data signals can be thought of as slices through a nominally circular patch on the Earth's surface. Because of the Fourier slice theorem, the constituent time series are Fourier transformed individually to yield a discretely sampled version of the two-dimensional Fourier transform of the scene's reflectivity function [7]. The spotlight image is then constructed by interpolating the polar Fourier data onto a regular rectangular grid so that the inverse Fourier transform can exploit the computational efficiency of the FFT.

The important distinction is that spotlight mode data is collected directly in the spatial Fourier domain, while stripmap data is not. Phase errors across a spotlight mode synthetic aperture are multiplied with the data spectrum, and consequently, the ideal spotlight PSR is convolved with the spectrum of the phase error function to yield the degraded PSR [24]. This stands in contrast to the stripmap case in which the ideal PSR spectrum is convolved with the phase error spectrum, as shown by (3.4). Assuming that the phase error spectrum is known, spotlight phase error analysis is fairly straightforward. For the stripmap case, however, certain restrictions must be placed on  $G(k_y)$  in order to obtain closed-form results that lend themselves to easy interpretation. Specifically,  $G(k_y)$  should be nonzero only within a relatively narrow region centered on  $k_y = 0$ .




(a)




(b)

**Figure 3.8.** The plots above show the result of applying the single-frequency results to a finite-bandwidth signal. Shown are the sinusoidal (a) and sawtooth (b) phase errors for a signal with a center frequency to bandwidth ratio of four. These plots are the finite-bandwidth equivalent of Figures 3.3(b) and 3.4(b) where the ideal PSR is omitted and the horizontal axis is enlarged by a factor of two.



(a)



(b)

**Figure 3.9. Synthetic aperture PSR as influenced by various degrees of quadratic phase error. The upper plot has no weighting applied, while the lower was created using a Hamming window. The horizontal axis is normalized length, where the interval  $[-0.5, 0.5]$  corresponds to the ideal cross range resolution,  $D/2$ . Also, each PSR is displayed relative to its own peak value in order to better illustrate the relative -3 dB widths of the mainlobes.**

### 3.4.2 Spatial Frequency Windowing Effects

The results of Section 3.3 are presented without any sort of spectral weighting, or windowing. This was done in order to facilitate the derivation of closed form analytical expressions for the PSR behavior. Such weighting is, however, generally considered to be standard practice for controlling the sidelobe behavior of the PSR. Figure 3.9 shows the result of Hamming weighting applied to various amounts of quadratic phase error. The figure clearly shows that spectral apodization can go a long way toward blunting the deleterious effects of phase errors on the stripmap PSR.

### 3.4.3 Moving Scatterers

An interesting feature is visible in many of the examples of SAS imagery shown in this thesis, especially those in Chapter 7. Although each image is well-focused with respect to sea floor, there is sporadic localized streaking in the cross-range direction. This streaking is due to scatterers moving in the water, which are likely to be fish. A stationary scatterer would appear in the SAS image as a well-focused pointlike object, and hence not discernible as a feature separate from the sea floor. As the scatterer moves, it introduces phase errors into its cross-range signal history. Therefore, it will appear blurred when the image is reconstructed.

An obvious question to ask is why the smearing is one-dimensional and parallel to the cross-range axis. After all, it seems reasonable to expect fish to be moving in all directions. This question can be answered using the preceding analysis. The  $y$ -axis is defined to be cross-range, while the  $x$ -axis represents range. Also defined are the unit vectors  $\mathbf{e}_x$  and  $\mathbf{e}_y$ . The sensor moves along the  $y$ -axis at a speed (assumed to be positive) of  $v_{\text{SAS}}$ , while the moving scatterer has velocity  $\mathbf{v} = [v_x \ v_y]^T$ . The time required to collect the data for a single synthetic aperture is  $t_{\text{SA}} = L_{\text{SA}}/v_{\text{SAS}}$ , where the synthetic aperture itself occupies the region of the  $y$ -axis satisfying  $-L_{\text{SA}}/2 \leq y \leq L_{\text{SA}}/2$ . The time during which the scatterer will inhabit a given range resolution cell is:  $t_{\text{dwell}} = \min\{\delta x / \mathbf{v} \cdot \mathbf{e}_x, t_{\text{SA}}\}$ .

As the sensor moves along the  $y$ -axis, the range from sensor to scatterer is  $R'(y) =$

$\sqrt{(x_0 - \epsilon_x(y))^2 + (y - \epsilon_y(y))^2}$  in which:

$$\epsilon_x(y) = \mathbf{v} \cdot \mathbf{e}_x \frac{y + L_{SA}/2}{v_{SAS}} \quad (3.37)$$

and

$$\epsilon_y(y) = \mathbf{v} \cdot \mathbf{e}_y \frac{y + L_{SA}/2}{v_{SAS}}. \quad (3.38)$$

Since the scatterer may move in any direction in the  $x$ - $y$  plane, so (3.2) is extended to account for errors in both range and cross-range. The range from sensor to scatterer can be written via truncated Taylor expansion as:

$$\begin{aligned} R' &\approx R + \left. \frac{dR'}{d\epsilon_x} \right|_{\epsilon_x=0} \epsilon_x + \left. \frac{dR'}{d\epsilon_y} \right|_{\epsilon_y=0} \epsilon_y \\ &\approx R - \left( 1 - \frac{y^2}{2x^2} \right) \epsilon_x - \frac{y}{x_0 - \epsilon_x} \epsilon_y \\ &\approx R - \epsilon_x - \frac{y}{x_0} \epsilon_y. \end{aligned} \quad (3.39)$$

The phase error function imposed on the SAS signal history can now be written:

$$\begin{aligned} g(y) &= \exp \left\{ j2k_0 \frac{v_x}{v_{SAS}} (y + L_{SA}/2) \right\} \cdot \exp \left\{ j2k_0 \frac{v_y}{v_{SAS}} \frac{y}{x_0} (y + L_{SA}/2) \right\} \\ &= \exp \{ j2k_0 L_{SA} \} \cdot \exp \left\{ j2k_0 \frac{v_x}{v_{SAS}} y \right\} \cdot \exp \left\{ j2k_0 \frac{v_y}{v_{SAS}} \frac{y^2}{x_0} \right\}. \end{aligned} \quad (3.40)$$

Equation (3.40) is easily analyzed using the results from Sections 3.3.3 and 3.3.2. There are three phase error terms. The first is a constant and is irrelevant in the sense that it doesn't affect the magnitude of the PSR. The second term describes the effect of the  $x$  component of the scatterer velocity and is a linear phase error. A scatterer moving along the  $x$ -axis toward or away from the synthetic aperture will therefore be well-focused, but shifted in cross range by an amount equal to  $x_0 v_x / v_{SAS}$ . The third term in (3.40) is a quadratic phase error and represents the effect of the scatterer's velocity in the  $y$  direction. This motion will cause the PSR to widen, or blur, by a length equal to  $\Delta_{PSR} = 2L_{SA}|v_y|/v_{SAS}$ . Thus, it is shown that the PSR for moving scatterers such as fish will appear primarily as streaks in the cross-range direction. An interesting consequence of this result is that the magnitude of  $v_y$  can be inferred from the width of the observed PSR. For example, the images containing moving scatterers shown in this thesis yield values of  $v_y$  in the neighborhood of 5 cm/s.



### 3.5 Comments

The phase error analysis began with the assumption of small (or residual) motion errors that are less than a resolution cell in extent. This simplification restricts the discussion to phase errors. A similar treatment for the linear and sinusoidal cases is given in the SAR literature by Fornaro *et al.* [52] in which no prior restrictions are placed on the magnitude of the motion errors. The authors then develop an exact expression for the two-dimensional spectrum of the observed synthetic aperture data. This spectral model is then simplified using the small-magnitude assumption. Closed-form expressions for various types of phase errors can be obtained from this simplified model.

It has been shown that the analysis of stripmap phase errors is, in general, much different than for the spotlight mode. Nevertheless, the results are similar for the two imaging modalities. For example, sinusoidal errors generate replicas of the PSR while random high-frequency errors induce a loss of contrast without a loss of resolution. The overall taxonomy for describing the qualitative effects of phase errors on stripmap imaging is essentially the same as for spotlight imaging, as presented by Carrara [24]. These results are useful in a variety of applications such as error budgeting for system design, developing improved strategies for data collection, and diagnosing the sources of image degradation for experimental data.

## CHAPTER 4

### THE REDUNDANT PHASE CENTER TECHNIQUE

#### 4.1 Introduction

Synthetic aperture imaging is ideally done by collecting data on a perfectly straight line. This rarely, if ever, occurs in reality, so the available sensor information must be used to estimate the actual vehicle trajectory in order to obtain properly focused imagery. This chapter describes the redundant phase center (RPC) technique which is the foundation for the motion estimation presented in the next chapter. In the literature, RPC has generally been considered to be a motion estimation technique unto itself (see [59], for example) and aimed purely at improving the quality of the SAS imagery. It is treated here primarily as a method for using the SAS data to accurately measure displacements in the image (or slant) plane. This is a subtle difference, but it opens the door to applying RPC-based motion estimates not only to the SAS imagery but also to improving the navigation of the host vehicle.

The redundant phase center technique is simply a combination of the vernier array concept described in Section 1.4 and the idea of correlating two similar signals to estimate their relative delay. The motivation stated in Section 1.4 for using a vernier array was to increase the area coverage rate of the SAS. An added benefit can be derived from the vernier array if the SAS is moved forward through the water such that the aftmost part of the array at ping  $p$  overlaps the foremost section at ping  $p - 1$  as shown in Figure 1.7. The advance per ping is set to be an integral number of phase center spacings that is less than the total number of receivers in the SAS array. For each overlapping phase center pair, the sensor has two observations of the sea floor which are identical, neglecting noise, except for any delay caused by unwanted platform motion. The delay is estimated by studying the cross-correlation of these two signals, and this is the subject of this chapter. Chapter 5 deals with interpreting the RPC delays to obtain a record of the actual platform trajectory.


**Table 4.1.** Symbols used to describe the delay observed by a pair of overlapping receivers used by the redundant phase center technique.

$\mathbf{d}_n$	Vector from vehicle center to the $n^{\text{th}}$ array element.
$t$	Time of ping transmission.
$\tau^n$	Time from ping transmission to reception by the $n^{\text{th}}$ array element (that is, the time of flight).
$\mathbf{x}(t)$	Position of the vehicle center.
$\Phi(t)$	Rotation matrix indicating the orientation of the vehicle.
$\mathbf{v}(t)$	Linear (translational) velocity of the vehicle.
$\boldsymbol{\omega}(t)$	Angular velocity of the vehicle.
$s$	Scatterer location.
$\mathbf{R}_n(t)$	Vector at time $t$ from the $n^{\text{th}}$ array element to a given scatterer.
$\mathbf{r}_n(t)$	Unit vector in the direction of $\mathbf{R}_n(t)$ .

## 4.2 Modeling the RPC Time Delays

This section gives a continuous-time model for the time delays observed by a redundant phase center pair. Table 4.1 contains the definition of the symbols used. It is assumed that the signal transmission and reception each occur instantaneously: the temporal Doppler shifting caused by motion of the array is ignored.

The diagram shown in Figure 4.1 depicts the array locations over the course of two consecutive pings used to estimate the time delay observed by a pair of redundant phase centers. For the sake of illustration, the array shown has eleven elements, and the transmitter is positioned such that it is coincident with the fifth element. It can be seen from the previous discussion regarding vernier arrays that the required overlap actually occurs in the phase-center space. In order to estimate the time delays, two overlapping channels are required. In the following, this overlapping pair is referred to an RPC pair. SAS arrays are often characterized by the number of RPC pairs they employ. This number can be adjusted by changing the forward advance between pings of the array. Using more RPC pairs affords an opportunity to average the motion estimates. However, this comes at the expense of having a shorter, non-overlapping, section of the array that is actually used for imaging. The consequence is a reduction of the area coverage rate. Conversely, fewer RPC pairs



**Figure 4.1. Schematic of two consecutive SAS pings with the array in motion. The physical arrays are shown. A similar diagram of the phase center arrays would show the expected overlap of elements 10 and 11. The transmitter is colocated with the fifth element of the receiver array.**

means that the imaging portion of the array is effectively longer resulting in an extension of the range and area coverage rate of the SAS.

There are four times of interest depicted in Figure 4.1. The transmissions occur at times  $t = t_{p-1}$  and  $t = t_p$ , respectively. Since the array is in continuous motion, the receptions occur at different locations at times  $t = t_{p-1} + \tau_{p-1}^n$  and  $t = t_p + \tau_p^n$ . Thus, the symbol  $\tau$  is used to denote the time between signal transmission and reception. The subscripts on both  $t$  and  $\tau$  indicate the time of transmission. Furthermore,  $\tau$  has a superscript indicating the position of the array element under consideration. The superscripts in Figure 4.1 are

numbers. In the diagram,  $\tau_{p-1}^1$  indicates the time of flight between transmission at time  $t_{p-1}$  and reception by the first element. Clearly,  $\tau$  can vary along the array. In the following equations the superscript “fore” denotes the element in the forward part of the array at ping  $p - 1$  that overlaps with the aft part of the array, hence the designation “aft”, at ping  $p$ . The quantity observed by an RPC pair comprising a forward channel from ping  $p - 1$  and the appropriate aft channel from ping  $p$  is  $\Delta\tau = \tau_p^{\text{aft}} - \tau_{p-1}^{\text{fore}}$ , where the respective times of flight can be expressed in terms of the kinematics:

$$\begin{aligned}\tau_{p-1}^{\text{fore}} &= \frac{1}{c} \left\{ \left| \mathbf{R}_{\text{tx}}(t_{p-1}) \right| + \left| \mathbf{R}_{\text{fore}}(t_{p-1} + \tau_{p-1}^{\text{fore}}) \right| \right\} \\ &= \frac{1}{c} \left\{ \left| \mathbf{s} - \mathbf{x}(t_{p-1}) - \Phi(t_{p-1}) \mathbf{d}_{\text{tx}} \right| + \left| \mathbf{s} - \mathbf{x}(t_{p-1} + \tau_{p-1}^{\text{fore}}) - \Phi(t_{p-1} + \tau_{p-1}^{\text{fore}}) \mathbf{d}_{\text{fore}} \right| \right\} \\ \tau_p^{\text{aft}} &= \frac{1}{c} \left\{ \left| \mathbf{R}_{\text{tx}}(t_p) \right| + \left| \mathbf{R}_{\text{aft}}(t_p + \tau_p^{\text{aft}}) \right| \right\} \\ &= \frac{1}{c} \left\{ \left| \mathbf{s} - \mathbf{x}(t_p) - \Phi(t_p) \mathbf{d}_{\text{tx}} \right| + \left| \mathbf{s} - \mathbf{x}(t_p + \tau_p^{\text{aft}}) - \Phi(t_p + \tau_p^{\text{aft}}) \mathbf{d}_{\text{aft}} \right| \right\}\end{aligned}\tag{4.1}$$

To simplify matters, let the position and orientation at the time of reception be approximated by way of a truncated first-order Taylor series expansion. The position of the origin of the vehicle’s body coordinates expressed in terms of a fixed external coordinate system is:

$$\mathbf{x}(t_{p-1} + \tau_{p-1}^{\text{fore}}) \approx \mathbf{x}(t_{p-1}) + \tau_{p-1}^{\text{fore}} \mathbf{v}(t_{p-1})\tag{4.2}$$

Prior to expanding the orientation of the vehicle at reception, let the rotation matrix  $\Phi(t_{p-1} + \tau_{p-1}^{\text{fore}})$  be factored into two sequential rotations:  $\Phi(t_{p-1} + \tau_{p-1}^{\text{fore}}) = \Psi(\tau_{p-1}^{\text{fore}}) \Phi(t_{p-1})$ . The latter term is the orientation of the vehicle at the time of transmission, while the former represents the additional rotation that occurs during the time of flight. The Taylor expansion is then:

$$\Phi(t_{p-1} + \tau_{p-1}^{\text{fore}}) \approx \Phi(t_{p-1}) + \tau_{p-1}^{\text{fore}} \dot{\Psi}(t_{p-1}) \Phi(t_{p-1}).\tag{4.3}$$

However,  $\dot{\Psi}$  evaluated at  $t_{p-1}$  is simply  $\dot{\Phi}(t_{p-1})$ , which is assumed to be known from the vehicle instrumentation. Thus,

$$\Phi(t_{p-1} + \tau_{p-1}^{\text{fore}}) \approx \Phi(t_{p-1}) + \tau_{p-1}^{\text{fore}} \dot{\Phi}(t_{p-1}) \Phi(t_{p-1}).\tag{4.4}$$

Further simplifications can be made if the rotation  $\Psi$  is small. In (4.4), the rotation rate  $\dot{\Phi}$  is a matrix. However, multiplying  $\dot{\Phi}$  times a vector can be shown to be equivalent to the cross product between the instantaneous angular velocity  $\boldsymbol{\omega} = [\dot{\psi} \ \dot{\phi} \ \dot{\theta}]^T$  and that same vector. This form is used henceforth. Substituting the Taylor approximations into the time-of-flight equations yields:

$$\begin{aligned}
c\tau_{p-1}^{\text{fore}} &= \left| \mathbf{s} - \mathbf{x}(t_{p-1}) - \Phi(t_{p-1})\mathbf{d}_{\text{tx}} \right| \\
&\quad + \left| \mathbf{s} - \left( \mathbf{x}(t_{p-1}) + \tau_{p-1}^{\text{fore}}\mathbf{v}(t_{p-1}) \right) - \left( \Phi(t_{p-1})\mathbf{d}_{\text{fore}} + \tau_{p-1}^{\text{fore}}\boldsymbol{\omega}(t_{p-1}) \times \Phi(t_{p-1})\mathbf{d}_{\text{fore}} \right) \right| \\
c\tau_p^{\text{aft}} &= \left| \mathbf{s} - \mathbf{x}(t_p) - \Phi(t_p)\mathbf{d}_{\text{tx}} \right| \\
&\quad + \left| \mathbf{s} - \left( \mathbf{x}(t_p) + \tau_p^{\text{aft}}\mathbf{v}(t_p) \right) - \left( \Phi(t_p)\mathbf{d}_{\text{aft}} + \tau_p^{\text{aft}}\boldsymbol{\omega}(t_p) \times \Phi(t_p)\mathbf{d}_{\text{aft}} \right) \right|
\end{aligned} \tag{4.5}$$

In an effort to facilitate matters later on, it is also assumed that the linear and angular velocities change negligibly over the duration of two pings. The linear translations can then be written as:

$$\mathbf{x}(t_p) \approx \mathbf{x}(t_{p-1}) + \Delta t \mathbf{v}(t_{p-1}) \tag{4.6}$$

$$\mathbf{x}(t_p + \tau_p^{\text{aft}}) \approx \mathbf{x}(t_{p-1}) + (\Delta t + \tau_p^{\text{aft}})\mathbf{v}(t_{p-1}) \tag{4.7}$$

and the rotations are:

$$\Phi(t_p) \approx \Phi(t_{p-1}) + \Delta t \boldsymbol{\omega}(t_{p-1}) \times \Phi(t_{p-1}) \tag{4.8}$$

$$\Phi(t_p + \tau_p^{\text{aft}}) \approx \Phi(t_{p-1}) + (\Delta t + \tau_p^{\text{aft}})\boldsymbol{\omega}(t_{p-1}) \times \Phi(t_{p-1}). \tag{4.9}$$

Making this approximation is convenient because it allows the motion to be expressed in terms of the position, orientation, and velocity of the first ping of the pair. The time delay

observed by the RPC pair shown in Figure 4.1 becomes:

$$\begin{aligned}
c\Delta\tau &= c(\tau_p^{\text{aft}} - \tau_{p-1}^{\text{fore}}) \\
&= \left| s - \left\{ \mathbf{x}(t_{p-1}) + \Delta t \mathbf{v}(t_{p-1}) \right\} - \left\{ \Phi(t_{p-1}) \mathbf{d}_{\text{tx}} + \Delta t \boldsymbol{\omega}(t_{p-1}) \times \Phi(t_{p-1}) \mathbf{d}_{\text{tx}} \right\} \right| \\
&\quad - \left| s - \mathbf{x}(t_{p-1}) - \Phi(t_{p-1}) \mathbf{d}_{\text{tx}} \right| \\
&\quad + \left| s - \left\{ \mathbf{x}(t_{p-1}) + (\Delta t + \tau_p^{\text{aft}}) \mathbf{v}(t_{p-1}) \right\} - \left\{ \Phi(t_{p-1}) \mathbf{d}_{\text{aft}} + (\Delta t + \tau_p^{\text{aft}}) \boldsymbol{\omega}(t_{p-1}) \times \Phi(t_{p-1}) \mathbf{d}_{\text{aft}} \right\} \right| \\
&\quad - \left| s - \left( \mathbf{x}(t_{p-1}) + \tau_{p-1}^{\text{fore}} \mathbf{v}(t_{p-1}) \right) - \left( \Phi(t_{p-1}) \mathbf{d}_{\text{fore}} + \tau_{p-1}^{\text{fore}} \boldsymbol{\omega}(t_{p-1}) \times \Phi(t_{p-1}) \mathbf{d}_{\text{fore}} \right) \right| \\
&= \mathbf{R}_{\text{tx}}(t_p) - \mathbf{R}_{\text{tx}}(t_{p-1}) + \mathbf{R}_{\text{aft}}(t_p + \tau_p^{\text{aft}}) - \mathbf{R}_{\text{fore}}(t_{p-1} + \tau_{p-1}^{\text{fore}}) \\
&= \mathbf{R}_a - \mathbf{R}_b + \mathbf{R}_c - \mathbf{R}_d,
\end{aligned} \tag{4.10}$$

where the notation in the last line has been introduced for the sake of brevity in the equations used subsequently.

The details of the mathematics presented in this section come into play in the next chapter. For the present discussion of measuring RPC time delays, the most important concept from this section is the fact that these delays vary with range, or acquisition time, unless (1) the sensor lies in the plane of the sea floor, and (2) its velocities  $\mathbf{v}$  and  $\boldsymbol{\omega}$  are both zero. This is, of course, never the case in practice, and it is customary to attempt to estimate the delay locally by using a short-time sliding window. Of these two effects, the range variation of the delay caused by nonzero altitude is the most significant. The component of the array velocity in the range direction is generally small enough so that the array movement within the duration of the correlation window may be neglected.

### 4.3 Short-Time Correlation-Based Time Delay Estimation

A short-time cross-correlation scheme is used in order to capture the range-dependence of the RPC time delays as well as obtain estimates that are robust to local decorrelation. A flowchart is provided in Figure 4.2 for reference in this description. The process begins with two time series- one from each channel of the RPC pair. A local short-time window

is then applied at the same position for both signals. The temporal length of the window typically corresponds to 0.5–2 m of range. For common SAS systems, this range of window lengths produces low-variance estimates of the delay while also being short enough to justify the assumption of constant delay over the window duration. It is tacitly assumed that signals 1 and 2 are not grossly time-shifted relative to one another. If they were, the windowed signals might not be highly correlated. This is not a problem in practice for motion estimation, although it may be for interferometry. Gross delays could be accounted for by lengthening the window, estimating the gross delay, shifting the shorter windows appropriately, and then proceeding as described below.

The next step in the time delay estimation process is to compute the correlation coefficient between the two signals:

$$\rho_{12}(t) = \frac{\text{cov}(s_1, s_2)}{\sqrt{\text{var}(s_1)\text{var}(s_2)}}. \quad (4.11)$$

Note that  $\rho_{12}$  is complex because of the use of the complex baseband representation of the received SAS signals. Examination of the cross-covariance function in more detail serves to illuminate the method used for time-delay estimation. Recall the form of a (noise-free) baseband complex echo return from a point scatterer at range  $R$ :

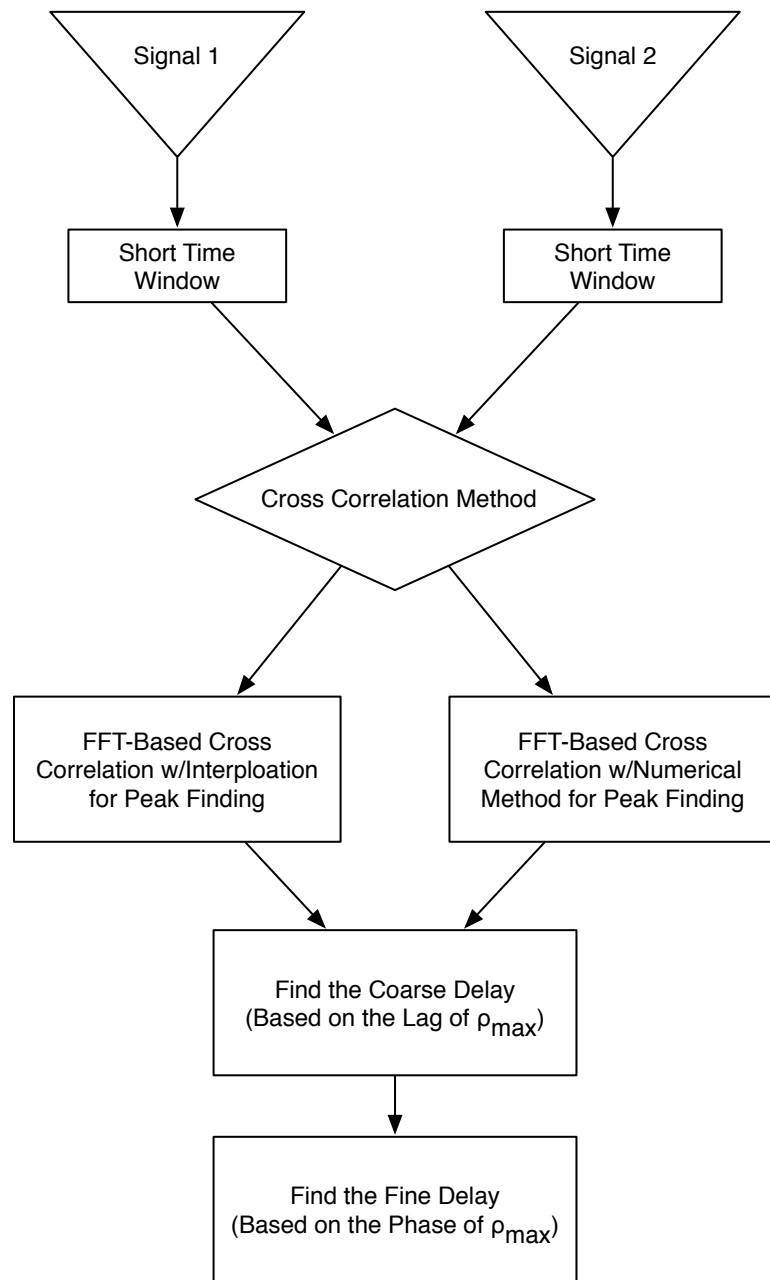
$$\begin{aligned} s_1(t) &= p_m \left( t - \frac{2R_1}{c} \right) \cdot \exp \{ -j\omega_0 t \} \\ &= p_b \left( t - \frac{2R_1}{c} \right) \cdot \exp \left\{ j\omega_0 \left( t - \frac{2R_1}{c} \right) \right\} \cdot \exp \{ -j\omega_0 t \} \\ &= p_b \left( t - \frac{2R_1}{c} \right) \cdot \exp \left\{ j\omega_0 \frac{2R_1}{c} \right\}. \end{aligned} \quad (4.12)$$

Likewise for  $s_2$ :

$$s_2(t) = p_b \left( t - \frac{2R_2}{c} \right) \cdot \exp \left\{ j\omega_0 \frac{2R_2}{c} \right\}. \quad (4.13)$$

Assuming  $s_1$  and  $s_2$  are zero-mean, wide-sense stationary, and ergodic, the cross-covariance





**Figure 4.2. Flowchart of the time delay estimation process. Signals 1 and 2 are typically much longer than the window length. The process shown in the diagram is repeated many times, and the output is a vector of time delays, one for each position of the local short-time window.**

function is:

$$\begin{aligned}
\text{cov}(s_1, s_2; t) &= \int s_1(\lambda) \bar{s}_2(\lambda - t) d\lambda \\
&= \int p_b\left(\lambda - \frac{2R_1}{c}\right) \cdot \exp\left\{j\omega_0 \frac{2R_2}{c}\right\} \cdot \bar{p}_b\left(\lambda - t - \frac{2R_2}{c}\right) \cdot \exp\left\{-j\omega_0 \frac{2R_2}{c}\right\} d\lambda \\
&= \int p_b\left(\lambda - \frac{2R_1}{c}\right) \cdot \bar{p}_b\left(\lambda - t - \frac{2R_2}{c}\right) \cdot \exp\left\{-j\omega_0 \frac{2(R_2 - R_1)}{c}\right\} d\lambda
\end{aligned} \tag{4.14}$$

where the overbar indicates complex conjugation and  $\omega_0$  is the center frequency of the signal. If the cross-covariance is evaluated at the lag that is equal to the relative time delay between  $s_1$  and  $s_2$ , then the following is obtained:

$$\text{cov}(s_1, s_2; t = 2(R_1 - R_2)/c) = \exp\left\{-j\omega_0 \frac{2(R_2 - R_1)}{c}\right\} \int \left|p_b\left(\lambda - \frac{2R_1}{c}\right)\right|^2 d\lambda. \tag{4.15}$$

This equation represents the (complex) value of the cross-covariance corresponding to the maximum value of  $|\text{cov}(s_1, s_2; t)|$ , and it has two important properties. First, the peak magnitude occurs when the lag is equal to the relative time delay between  $s_1$  and  $s_2$ . Second, this value is in general complex and its phase is proportional to the time delay being estimated, modulo  $2\pi$ . This suggests a two step approach for computing the time delay [60, 61]. First, compute the *coarse delay*,  $\Delta\tau_{\text{coarse}}$ , and then follow on with the *fine delay*,  $\Delta\tau_{\text{fine}}$ , such that:

$$\Delta\tau = \Delta\tau_{\text{coarse}} + \Delta\tau_{\text{fine}}. \tag{4.16}$$

The coarse delay is equal to the location of the peak of the magnitude of  $\text{cov}(s_1, s_2; t)$ . It serves as an indicator of the integral number of wavelengths at the center frequency spanned by the relative time shift. The fine delay represents the contribution to the time delay that is a fractional part of a wavelength. It is found from the phase of the value of  $\text{cov}(s_1, s_2; t)$  that has the largest magnitude, which is the value used for the coarse delay. The fine delay is computed according to:

$$\Delta\tau_{\text{fine}} = -\frac{\alpha}{2\pi f_0} + \frac{m}{f_0} \tag{4.17}$$

in which  $\alpha$  is the phase of (4.15) and  $m \in \mathbb{Z}$  is the *ambiguity number* satisfying the condition that  $|\Delta\tau_{\text{coarse}} - \Delta\tau_{\text{fine}}|$  is minimized.

Uncertainty in the estimate arises if noise is introduced to the signals being compared. For this case, Bellettini and Pinto [62, 61] have reported the Cramér-Rao lower bounds on the coarse and fine time delay estimates. They are:

$$\sigma_{\text{coarse}} = \frac{\sqrt{3}}{\pi B_w \sqrt{B_w T}} \sqrt{\frac{1}{\nu} + \frac{1}{2\nu^2}} \quad (4.18)$$

and

$$\sigma_{\text{fine}} \approx \frac{1}{2\pi f_0 \sqrt{B_w T}} \sqrt{\frac{1}{\nu} + \frac{1}{2\nu^2}}. \quad (4.19)$$


In these calculations  $B_w$  is the signal bandwidth,  $T$  is the temporal duration of the correlation window, and  $\nu$  is the signal-to-noise ratio expressed in terms of the peak magnitude of the correlation coefficient:  $\nu = |\rho|/(1 - |\rho|)$ . The product  $B_w T$  is defined as the number of independent samples in the correlation. The signals  $s_1$  and  $s_2$  are bandwidth sampled. It is sometimes useful to rewrite  $T$  in terms of the spatial window length (in meters) used for the correlation:  $T = 2L_{\text{win}}/c$ .

The equation in the paragraph above relating SNR and the correlation coefficient is easily derived as follows (see [63]). Consider a pair of signals  $s_1 = s + n_1$  and  $s_2 = s + n_2$ , in which  $s$  is the part that is common to both signals, and  $n_1$  and  $n_2$  are zero-mean iid noise processes that are not correlated with  $s$ . The magnitude of the correlation coefficient is:


$$\begin{aligned} |\rho| &= \frac{|\text{cov}(s_1, s_2)|}{\sqrt{\text{var}(s_1)\text{var}(s_2)}} \\ &= \frac{\sigma_s^2}{\sigma_s^2 + \sigma_n^2}. \end{aligned} \quad (4.20)$$

Factoring  $\sigma_s^2$  out of the numerator and denominator and defining the SNR as  $\nu = \sigma_s^2/\sigma_n^2$  gives the relationship used above:  $|\rho| = 1/(1 + \nu^{-1})$ .

Figure 4.3 provides a real-world example of the potential accuracy of the RPC time delay estimation. Figure 4.3(a) shows a set of delays obtained from SAS field data using RPC in addition to a plot of the curve describing the delays whose parameters were computed using the nonlinear least squares technique given in Section 5.6. The standard deviation cannot be computed directly since the delays vary with range and do not have a



(a)



(b)

**Figure 4.3. (a) Time delays estimated using RPC (diamonds) plotted along with the closed-form description of the RPC delay (solid line) found using the nonlinear least squares technique of Section 5.6. The RPC delays shown correspond to  $\rho \geq 0.85$ . Delays with  $\rho < 0.85$  are not shown and were rejected for the least-squares computation. (b) The lower plot shows the result of subtracting the least-squares delay from the RPC observations (black diamonds, left-hand scale). Also shown are the correlation coefficients for the RPC delays (circles, right-hand scale). Note that all delays are expressed as distances.**

meaningful expected value that can be removed for the computation of  $\sigma_{\text{fine}}$ . To circumvent this problem, the range-varying behavior is accounted for by subtracting the least-squares estimate of the RPC delay from the observed values. The result of this operation is shown in Figure 4.3(b), as are the correlation coefficients for the delays. In this example the average correlation coefficient is  $\rho = 0.95$ , and the standard deviation expressed in terms of wavelengths is  $\sigma_{\text{fine}} = \lambda/185$ . By comparison, the CRLB computed from (4.19) is  $\sigma_{\text{fine}} = \lambda/216$ . The performance in this instance is exceptionally good, and in practice the results often come reasonably close to the CRLB. This extraordinary accuracy is the key to the successful implementation of the motion estimation and compensation techniques presented in Chapters 5 and 6.

The experimental  $\sigma_{\text{fine}}$  is occasionally observed to be less than the predicted CRLB. This might seem to be problematic at first, as it can be proven that the CRLB represents the best possible performance of any estimator. There are several possible sources of the discrepancy, however. First, one must consider that the assumptions used to derive the CRLB may not apply to the experimentally-observed data. For example, closed-form expressions for the CRLB are often more easily obtained by assuming that the estimator has zero bias. In this case, the result cannot preclude the existence of a biased estimator with a smaller variance; such CRLBs only guarantee that no unbiased estimator exists with better performance. CRLB derivations are also frequently simplified by modeling the noise as an additive white Gaussian process, when the true noise may be multiplicative, correlated, and/or colored. Another potential source of inconsistency between the theoretical CRLB and the experimental result is that the relationship between signal-to-noise and correlation coefficient (4.20) used in (4.19) may differ from the actual relationship for the observed SAS data. It should also be remembered that the *ad hoc* procedure used above to account for the range variation of the delay may affect the result. Regardless of these concerns, the CRLB given by (4.19) is accurate and is an invaluable tool for the design and analysis of SAS systems and processing algorithms. The reader wishing to consider in depth the

development of this result is referred to [64, 65, 66, 67, 62]. Additionally, Johnson and Dudgeon [68] give a brief discussion of bounds other than the CRLB that may be more applicable to the problem of time delay estimation in low SNR conditions.

#### 4.3.1 Algorithms for Coarse Time Delay Estimation

To close the discussion of the time delay estimation algorithm, two methods are presented that can be used to accurately locate the peak of  $\rho_{12}$  without resorting to highly oversampling  $s_1$  and  $s_2$ . The first and easiest to implement is simply to fit a parabola using the discrete peak of  $|\rho_{12}|$  and its two neighboring samples. Because only three samples are employed, a perfect parabolic fit can be guaranteed. Furthermore, an explicit analytical expression exists for locating the peak value of the parabola. This expression is derived in Section III.2.2.2 of [17] and is:

$$\Delta n = \frac{-0.5[f(n+1) - f(n-1)]}{f(n+1) - 2f(n) + f(n-1)}. \quad (4.21)$$

Given  $f = |\rho_{12}|$  evaluated at the peak value and the two adjacent samples, this equation yields the fractional sample offset of the parabolic peak. Thus, the coarse time delay equals the integer sample lag of the peak of  $|\rho_{12}|$  plus the offset computed from (4.21). This interpolative approach is robust and easy to compute. However, its accuracy is ultimately limited by the fact that the peak of the correlation coefficient is approximated by a parabola. Oversampling the correlation coefficient tends to improve the approximation and thus the accuracy of the interpolated peak location. This is a consequence of the Weierstrass approximation theorem, a result from the field of real analysis [69]. In essence, it states that a continuous function can be piecewise approximated by polynomials to any degree of accuracy.

The second technique for finding the value of  $\rho_{12}$  corresponding to the maximum of  $|\rho_{12}|$  is iterative in nature. It can be used to locate the peak to an arbitrary degree of accuracy without oversampling. The procedure is to shift  $s_2$  relative to  $s_1$  until the maximum of  $|\rho_{12}|$  occurs precisely on a discrete sample. The first step is to compute  $|\rho_{12}|$  as well as

its derivative with respect to the lag  $t$ ,  $\rho'_{12}$ . The value of  $\rho'_{12}$  evaluated at the peak sample of  $|\rho_{12}|$  serves as an indication of how much  $s_2$  should be shifted relative to  $s_1$  in order to result in the highest possible value of  $|\rho_{12}|$ . In other words, the probability of the peak of  $|\rho_{12}|$  falling exactly on a discrete sample is zero. However, this condition can be forced to occur by iteratively shifting  $s_2$ . Newton's method is used to accomplish the goal. Recall that Newton's method for finding a zero crossing is:  $x_{n+1} = x_n - f(x_n)/f'(x_n)$ . In the present context, the application is to solve iteratively for the coarse time delay,  $\Delta\tau_{\text{coarse}}$ , such that  $|\rho_{12}(\Delta\tau_{\text{coarse}})|^2 = \max\{|\rho_{12}(t)|^2\}$ . Newton's method applied to the problem of finding the peak magnitude of the cross-correlation is:

$$\Delta\tau_{\text{coarse},n+1} = \Delta\tau_{\text{coarse},n} + \frac{\frac{d}{dt}|\rho_{12}(\Delta\tau_{\text{coarse},n})|^2}{\frac{d^2}{dt^2}|\rho_{12}(\Delta\tau_{\text{coarse},n})|^2}. \quad (4.22)$$

The first and second derivatives of the squared envelope of  $\rho_{12}$  are required in order to execute this computation. Fortunately, analytical expressions exist that make these derivatives easy to compute without resorting to finite-differencing schemes. They are:

$$\begin{aligned} \frac{d}{dt}|\rho_{12}(t)|^2 &= \frac{d}{dt} \bar{\rho}_{12}(t) \cdot \rho_{12}(t) \\ &= \bar{\rho}_{12}(t) \cdot \mathcal{F}_\omega^{-1}\{j\omega P_{12}(\omega)\} + \rho_{12}(t) \cdot \mathcal{F}_\omega^{-1}\{j\omega P_{12}(-\omega)\} \end{aligned} \quad (4.23)$$

and

$$\begin{aligned} \frac{d^2}{dt^2}|\rho_{12}(t)|^2 &= \frac{d^2}{dt^2} \bar{\rho}_{12}(t) \cdot \rho_{12}(t) \\ &= \bar{\rho}_{12}(t) \cdot \mathcal{F}_\omega^{-1}\{-\omega^2 P_{12}(\omega)\} + \rho_{12}(t) \cdot \mathcal{F}_\omega^{-1}\{-\omega^2 P_{12}(-\omega)\} \\ &\quad + 2\mathcal{F}_\omega^{-1}\{j\omega P_{12}(\omega)\} \cdot \mathcal{F}_\omega^{-1}\{j\omega P_{12}(-\omega)\} \end{aligned} \quad (4.24)$$

where  $P_{12}(\omega)$  is the Fourier transform of  $\rho_{12}(t)$ . Note also that the overbar indicates complex conjugation and that the terms  $P_{12}(-\omega)$  can be readily computed using the time-reversal property of the Fourier transform:  $P_{12}(-\omega) = \mathcal{F}_t\{\bar{\rho}_{12}(t)\}$ . The differentiation property of the Fourier transform has also been used:  $d\rho_{12}(t)/dt = j\omega\mathcal{F}_t\{\rho_{12}(t)\}$ . With the required derivatives in hand, the computation of  $\Delta\tau_{\text{coarse},n+1}$  is repeated until the norm of  $\frac{d}{dt}|\rho_{12}(t)|$  is below a specified threshold.


## 4.4 Comments

The question of optimal window length naturally arises when performing RPC time delay estimation. This issue is complicated in general if one demands a truly optimal solution, and the methods of harmonic analysis may prove to be of use in this case. This subject is not explored further here. However, there are a few general observations that lead to window lengths that perform well and are almost always suitable, if not optimal.


It is important to recognize from (4.19) that, for a fixed bandwidth, the delay estimates improve as the window length increases because the number of independent samples increases. However, the window length should not be increased arbitrarily. One reason for this restriction is the underlying assumption that the delays are stationary over the duration of the window, while the geometry of the problem causes the delays to vary with range (*cf.* (4.10) which is a function of scatterer position  $\mathbf{s}$ , Figure 5.6, and Section 6.2). Another argument for shorter window lengths is that the correlation on the RPC channels degrades with range due to the SNR drop associated with effects such as spreading loss, absorption, medium instability, and multipath (see Figure 4.4(a)). Thus, adding a large number of samples to the window can actually have a negative impact on the quality of the delay estimates. Furthermore, the correlation may degrade locally. The most frequent causes of this are biologics and shadowing. (Fish and Carr [1, 2] and McCarthy [70] discuss the presence of biologics, such as fish and vegetation, in RAS imagery.) Figure 4.4 shows a SAS image and its corresponding RPC correlation map. The figures show several small fish schools and the resulting localized decorrelation they cause. In this case the motivation for shorter correlation windows is clear, as this provides a degree of robustness to the motion estimation by enabling the locally unreliable delays to be rejected.

A general rule for choosing the minimum length of the correlation window can be obtained from (4.19). Figure 4.5 shows a family of curves representing the Cramér-Rao lower bound on the time delay estimate for an example system with a bandwidth of 30 kHz and center frequency of 120 kHz. These curves exhibit a knee, beyond which there is little






(a)



(b)

**Figure 4.4. RPC correlation map (a) and the corresponding SAS image (b). The dark holes in the correlation map are caused by schools of fish that are visible as blurred patches in the image. Also note the bright wishbone-shaped object in the image. The strong scattering from this object causes an increase in the local RPC correlation between 45 m and 50 m range.**



**Figure 4.5. Cramér-Rao lower bound on the time delay estimates (4.19) for a SAS with a 30 kHz bandwidth and a center frequency of 120 kHz. The curves represent correlation coefficients ranging from 0.5 (upper curve) to 0.9 (lower curve) in increments of 0.1.**

advantage to increasing the correlation window length. Thus, the minimum window length is suggested by the location of the bend in the CRLB curve for the typical correlation value expected. The maximum length, on the other hand, is generally driven by the amount of delay variation expected within the duration of the window. This variation diminishes with range (see Figure 4.3(a)), so that longer windows may then be used to counter the drop in correlation that also occurs at range.

## CHAPTER 5

### MOTION ESTIMATION TECHNIQUES

#### 5.1 Introduction


This chapter outlines a number of ways to estimate the motion of the SAS array and the vehicle carrying it. Most of these are based on the redundant phase center technique described in the previous chapter. While RPC is the thread connecting the motion estimation schemes presented, it is shown that there are numerous ways to usefully interpret and exploit the RPC time-delay measurements. Hence, it is difficult to identify a particular algorithm as being the best. The choice is largely driven by the number of RPC channels in the array, the amount and quality of auxiliary navigation data, and the computational resources available.

#### 5.2 Defining the Nominal Sensor Trajectory

The ultimate goal is to derive the information needed to properly compensate the SAS data for the unwanted motion of the array. The overall approach to motion estimation is to first establish a baseline, or *nominal trajectory*, that will be used as a reference for all motion. The motion is then estimated and expressed relative to the coordinate system defined by the nominal trajectory (NT). The  $y$ -axis is taken to be along the NT, the  $z$ -axis is up, toward the surface of the water, and the  $x$ -axis is range. In the discussion to follow, the NT is understood to be computed using a method like the following least-squares regression. First, the latitude and longitude of the platform at ping  $p$  are converted into meters traveled north and east with respect to the first ping. The conversion used is

$$x_{\text{east}}[p] = (\theta_{\text{long}}[p] - \theta_{\text{long}}[1]) \cdot R_{\text{earth}} \cos \theta_{\text{lat,avg}} \quad (5.1)$$

$$x_{\text{north}}[p] = (\theta_{\text{lat}}[p] - \theta_{\text{lat}}[1]) \cdot R_{\text{earth}} \quad (5.2)$$



**Figure 5.1.** The nominal trajectory (NT) is defined on a scene-by-scene basis. The  $y$ -axis and the yaw relative to it are defined by the mean track over the sea floor.

In computing  $x_{\text{east}}[p]$ , the mean scene latitude has been used, and the radius of the earth is 6,378,140 m. This conversion is done on a scene-by-scene basis. Since the typical scene dimension is on the order of tens of meters, this locally-flat approximation is perfectly acceptable. The next step is to actually compute the best fit line to the vehicle positions. This is done via:

$$A = [x_{\text{east}} \mathbf{1}] \quad (5.3)$$

$$\begin{bmatrix} \alpha_{\text{slope}} & \alpha_{\text{intercept}} \end{bmatrix} = (A^T A)^{-1} A^T x_{\text{north}} \quad (5.4)$$

The quantity  $\alpha_{\text{slope}}$  gives the slope of the best straight line fit to the actual vehicle positions over the ping interval contained in the scene. Note that if the vehicle course is close to being due north or due south,  $|\alpha_{\text{slope}}|$  would be large, possibly giving rise to errors in the NT. In this case, the roles of  $x_{\text{east}}$  and  $x_{\text{north}}$  could be swapped. What matters is that the  $x$ -axis in this regression should be the cardinal direction closest to the actual vehicle track.

It is important to know  $\alpha_{\text{slope}}$  because navigation sensors record compass heading, but what is required for motion estimation is the vehicle yaw relative to the track used for image formation. Once  $\alpha_{\text{slope}}$  has been found, the yaw can be computed from the following equation.

$$\theta_{\text{yaw}} = -(\theta_{\text{heading}} - \pi/2 + \tan^{-1} \alpha_{\text{slope}}). \quad (5.5)$$

Recall that the compass measures angles in the clockwise direction, while the coordinate system employed in this thesis uses counterclockwise angles. The NT can be chosen somewhat arbitrarily. However, some choices are better than others in practice. For example, see Section 6.5. The approach taken here is probably the most straightforward in that it is simple and intuitive. Another useful way to define the NT would be to find the track that minimizes the mean square yaw. In any case, the NT is a line in three dimensional space, and it is considered to be perfectly horizontal (that is, parallel to the flat earth) at a height equal to the mean vehicle altitude over the ping interval forming the scene.

If one did not wish to make the locally-flat earth assumption, the least squares computation of the NT heading would have to be carried out in spherical coordinates. In this setting, the problem goes from finding a best-fit line to points on a plane to finding the best-fit great circle to points on the surface of a sphere. This approach is described by de Witte [71]. This added complexity is unnecessary for the distances relevant to making SAS imagery.

### 5.3 IMU-Based Motion Estimation

For host vehicles equipped with an inertial measurement unit (IMU), the most straightforward approach to synthetic aperture motion estimation is to simply rely on the vehicle's navigation records [72]. Indeed, this technique is frequently used for airborne SAR imaging. The IMU-based motion estimate of the position of the  $m^{\text{th}}$  receiver on the SAS array is given by:

$$\mathbf{R}_m(t) = \mathbf{R}_{\text{IMU}}(t) + \Phi(t)\mathbf{d}_m, \quad (5.6)$$

where  $\mathbf{R}_m(t)$  is equal to the IMU's estimate of its position,  $\mathbf{R}_{\text{IMU}}(t)$ , plus the distance from the IMU to the  $m^{\text{th}}$  receiver,  $\Phi(t)\mathbf{d}_m$ . The vector  $\mathbf{d}_m$  pointing from the IMU to the  $m^{\text{th}}$  receiver is expressed in the vehicle body coordinates; hence, it is fixed with respect to time. The rotation matrix  $\Phi(t)$  is provided by the IMU and transforms from the body coordinates into the external (earth) reference frame.

The drawback to this simple technique is that its effectiveness depends on the accuracy of the IMU which, in turn, is generally tied to its cost, with good accuracy being expensive. In nearly all SAS literature dealing with the problem of motion estimation, it is stated that IMU accuracy is generally insufficient for SAS motion estimation and compensation, regardless of the cost. This statement is justified by requiring that the motion be known to within  $\lambda/8$  over the entire synthetic aperture. While this is a sufficient requirement, the analysis of Chapter 3 can be used to derive more realistic accuracy requirements for the IMU. It should also be pointed out that inferior motion compensation is often better than none. For this reason, the IMU motion estimate makes a good backup solution for situations in which the RPC-based techniques fail; for example, one or more RPC channels malfunction or the correlations drop due to environmental effects or the presence of biologics.

## 5.4 Slant Plane Redundant Phase Centers


This section discusses the most widely used and most fundamental of SAS motion estimation schemes. In the literature, it is often referred to as the displaced phase center antenna (DPCA) technique. It is referred to here as the slant plane redundant phase center (SPRPC) technique. There are two reasons for the departure from the common terminology: The first is to avoid confusion with another, related, radar concept known as DPCA used for moving target identification (MTI) [73, 74]; hence the use of RPC in place of DPCA. The “slant plane” prefix serves to emphasize the fact that the motion is estimated without regard to any external frame of reference. At any given range, the motion estimate is computed relative only to the scatterers at that range. Therefore, the motion estimate does not require knowledge of the bathymetry, as it is computed in the instantaneous slant plane.

The slant plane RPC technique is described in the literature and in patents by several authors. However, the most thorough description is arguably that given by Pinto and Belletini in 2002 [62]. The authors describe the DPCA motion estimation scheme and provide

an analysis of the theoretical limits on its accuracy; some of these results were used earlier in Section 4.3. The content of the present section is drawn primarily from this source.

SPRPC begins with the same raw material as most other SAS motion estimation schemes; the time delays measured using the redundant phase center channels as described in Chapter 4. The RPC delays, in general, vary with range because of the sensor altitude and the bathymetry. It is assumed that RPC time-delay measurements have been made at all ranges in the scene. If this is not the case, then the solution must be interpolated/extrapolated to those ranges for which estimates do not exist. At any given range, the observed RPC time delay comprises contributions from three sources. The first is the translation that occurred between pings  $p - 1$  and  $p$  (that is, the translational velocity). The second source of delay is the static attitude of the array relative to the sensor baseline, or nominal trajectory. The third contribution to the delay is the change in attitude from ping  $p - 1$  to  $p$ , or the angular velocity.

In [62], the first source of time delay is called DPCA sway and the latter two are termed DPCA yaw (see Figure 5.2). The rationale for calling them sway and yaw is as follows: Consider the point on the sea floor corresponding to a given RPC time delay measurement. A plane is defined by this point and the along-track positions of the sensor at pings  $p - 1$  and  $p$ . This plane can be thought of as the slant plane for that range; it is the instantaneous slant plane. The RPC time delays measure the projection into this plane of the actual sensor locations. In other words, the RPC portion of the SAS array effectively senses only sway and yaw. This information is sufficient to perform compensation at the range corresponding to the RPC delay, and no external reference frame is required. Only sway and yaw as projected into the slant plane are sensed because the RPC portion of the SAS array has been assumed to be one-dimensional. If the overlapping section of a vernier array were two-dimensional, then the motion could be resolved more completely. The paper by Doisy [75] gives an analysis of two- and three-dimensional sonar arrays used for correlation-based motion estimation.



**Figure 5.2.** The diagram above, after Pinto and Bellettini [62], shows the DPCA sway and yaw for two consecutive pings. The shaded box indicates the overlapping portion of the SAS phase center array used for motion estimation. The DPCA sway,  $\gamma$ , is referenced to the middle of this region.

According to [62], the Cramer-Rao lower bounds on the DPCA sway ( $\gamma$ ) and yaw ( $\psi$ ) estimates are:

$$\begin{aligned}\sigma_\gamma &= \frac{1}{2\pi} \frac{\lambda_0}{2} \frac{1}{\sqrt{B_w T}} \frac{1}{\sqrt{K}} \sqrt{\frac{1}{\nu} + \frac{1}{2\nu^2}} \\ \sigma_\psi &= \frac{\sqrt{3}}{\pi} \sqrt{\frac{K-1}{K+1}} \frac{\lambda_0}{(K-1)D} \frac{1}{\sqrt{B_w T}} \frac{1}{\sqrt{K}} \sqrt{\frac{1}{\nu} + \frac{1}{2\nu^2}},\end{aligned}\tag{5.7}$$


where  $\nu$  is the signal-to-noise ratio,  $B_w T$  is the number of independent samples in the short-time window used for RPC time delay estimation (see Section 4.3),  $D$  is the horizontal length of a single array element, and  $K$  is the number of independent elements in the RPC portion of the SAS array. The signal-to-noise ratio is related to the magnitude of the correlation coefficient via  $\nu = |\rho|/(1 - |\rho|)$ .

The sway and yaw in the image plane are measured by the RPC portion of the SAS array and then applied to the imaging (nonredundant) portion of the array to effect the motion compensation. The apparent sonar trajectory at the time corresponding to ping  $p$  is given by (see Section V of [62]):

$$\begin{aligned}y_p &= y_{p-1} + D_{\text{ping}} \\ x_p &= x_{p-1} + \gamma_{p-1} + \frac{D}{2}\theta_{p-1} + \frac{D}{2}\theta_p \\ \theta_p &= \theta_{p-1} + \psi_{p-1},\end{aligned}\tag{5.8}$$

where  $D_{\text{ping}}$  is the ideal advance per ping,  $\gamma_p$  and  $\psi_p$  are the sway and yaw from ping  $p-1$  to  $p$  as projected into the instantaneous image plane, and the angles  $\theta_p$  are assumed to be small so that  $\sin \theta_p \approx \theta_p$ . Figure 5.3 illustrates these quantities. The inter-ping sway of the





**Figure 5.3.** The diagram above, after Pinto and Bellettini [62], shows the trajectory of the SAS phase center array in the slant (or image) plane. The first array is shown as being parallel to the  $x$ -axis to indicate that the integrated trajectory estimate given by (5.10) is referenced to the first ping.

geometric center of the SAS array is given by  $(x_p - x_{p-1})$ , and it is composed of the sway seen by the RPC portion of the array as well as lever arm effects due to the in-plane yaw at pings  $p - 1$  and  $p$ . Equation (5.8) can be manipulated into the following form:

$$\begin{aligned}
 y_p &= (p - 1)D_{\text{ping}} \\
 x_p &= \sum_{l=1}^{p-1} \gamma_l + D \sum_{l=1}^{p-1} \left( p - l - \frac{1}{2} \right) \psi_l \\
 \theta_p &= \sum_{l=1}^{p-1} \psi_l,
 \end{aligned} \tag{5.9}$$

where  $l$  is the ping index. The accumulated position error can be written from this form of the equations:

$$\begin{aligned}
 \delta y_p &= 0 \\
 \delta x_p &= \sum_{l=1}^{p-1} \delta \gamma_l + D \sum_{l=1}^{p-1} \left( p - l - \frac{1}{2} \right) \delta \psi_l \\
 \delta \theta_p &= \sum_{l=1}^{p-1} \delta \psi_l.
 \end{aligned} \tag{5.10}$$

The along track error is zero by assumption. Pinto and Bellettini point out that the cross-track error accumulates like a random walk in the presence of pure sway error ( $\delta \psi_p = 0$

for all  $p$ ) and like an integrated random walk in the presence of pure DPCA yaw error ( $\delta\gamma_p = 0$  for all  $p$ ). Error accumulates much more rapidly in the latter case resulting in a highly correlated pattern of phase error along the synthetic aperture. It was shown in Chapter 3 that uncorrelated phase errors are less destructive to SAS focus than correlated phase errors.

The SPRPC (or DPCA) technique is the quintessential method of SAS motion estimation. One advantage is that it requires no input from other motion sensors and requires no external frame of reference. It is possible, at least in theory, to perform motion compensation using only the SAS data itself— an advantage that can be used to reduce the cost and complexity of the overall SAS/vehicle system. It is for this reason that motion estimation and compensation are sometimes called autofocus in the literature. In practice, it is often necessary to extrapolate the SPRPC solution to certain locations in the SAS image. In this situation, it often becomes necessary to make assumptions about the bathymetry and sensor location. Another weakness of SPRPC is that the in-plane yaw estimates achievable using redundant phase centers are generally inferior compared to the attitude measurements that are likely to be provided by the on-board navigation system of an AUV. For any given set parameters, the accuracy of the DPCA sway and yaw can be improved by increasing the number of elements used for RPC delay estimation. This option carries a steep cost, however, as it implies operating with a lower area coverage rate or using sonars with extra elements and the added electronics and power consumption associated with them.

## 5.5 Nonlinear Least Squares

This section outlines the use of nonlinear least squares (NLLS) to find the vector  $\mathbf{v}$  in (4.10) under the assumption that the array's attitude and angular velocity are known from an additional sensor such as an IMU. The solution is Newton's method and is described in Chapter 10 of [76]. The NLLS solution is similar in form to the one-dimensional Newton's

method solution for finding local maxima or minima:

$$x_{n+1} = x_n - f'(x_n)/f''(x_n). \quad (5.11)$$

Many estimates of the time delay are typically available for a given RPC pair because of the sliding short-time window described in the previous chapter. The problem is therefore overdetermined, and the RPC delays are used to find the best estimate of the three unknown parameters,  $v_x$ ,  $v_y$ , and  $v_z$ , expressed as the vector  $\mathbf{v} = [v_x \ v_y \ v_z]^T$ . The goal is to find the  $\mathbf{v}$  that minimizes the quantity  $f(\mathbf{v}) = \frac{1}{2}D(\mathbf{v})^T D(\mathbf{v})$ , where  $D(\mathbf{v})$  is known as the *residual function*. The  $m^{\text{th}}$  entry of  $D(\mathbf{v})$  is written, using (4.10), as:

$$D_m(\mathbf{v}) = c\Delta\hat{\tau}_m - c\Delta\tau_m, \quad (5.12)$$

where there are  $M$  measured RPC time delays and  $D(\mathbf{v}) \in \mathbb{R}^{M \times 1}$ . The residual function is simply the product of the sound speed and the vector of the time delays computed using an estimate of  $\mathbf{v}$  (that is,  $c\Delta\hat{\tau}_m$ ) minus the observed RPC time delays,  $c\Delta\tau_m$ .

The first and second derivatives of  $f(\mathbf{v})$  are required in order to form the solution using Newton's method. The first derivative of  $f(\mathbf{v})$  is

$$\begin{aligned} \nabla f(\mathbf{v}) &= \sum_{m=1}^M D_m(\mathbf{v}) \cdot \nabla D_m(\mathbf{v}) \\ &= J(\mathbf{v})^T D(\mathbf{v}), \end{aligned} \quad (5.13)$$

where  $J(\mathbf{v}) \in \mathbb{R}^{M \times 3}$  is the Jacobian matrix with  $J(\mathbf{v})_{mn} = \partial R_m(\mathbf{v})/\partial v_n$ . Here, the  $m$  index denotes the  $m^{\text{th}}$  observed time delay and  $n$  indexes over the components of  $\mathbf{v}$ . The  $m^{\text{th}}$  row of  $J(\mathbf{v})$  is

$$J(\mathbf{v})_{m*} = \left[ \Delta t \mathbf{r}_a + (\Delta t + \tau_f^{\text{aft}}) \mathbf{r}_c - \tau_i^{\text{fore}} \mathbf{r}_d \right]^T. \quad (5.14)$$

The rows of the matrix  $J(\mathbf{v})$  are evaluated using known quantities and the current estimate of  $\mathbf{v}$ . The lowercase  $\mathbf{r}$  is used to indicate the unit vector pointing in the direction of the corresponding  $\mathbf{R}$  in (4.10).

The second-derivative of  $f(\mathbf{v})$  is:

$$\begin{aligned}\nabla^2 f(\mathbf{v}) &= \sum_{m=1}^M \left( \nabla D_m(\mathbf{v}) \cdot \nabla D_m(\mathbf{v})^T + D_m(\mathbf{v}) \cdot \nabla^2 D_m(\mathbf{v}) \right) \\ &= J(\mathbf{v})^T J(\mathbf{v}) + S(\mathbf{v}),\end{aligned}\tag{5.15}$$

where

$$S(\mathbf{v}) \equiv \sum_{m=1}^M D_m(\mathbf{v}) \cdot \nabla^2 D_m(\mathbf{v}).\tag{5.16}$$

The second derivative, or Hessian, of  $D_m(\mathbf{v})$  is  $\nabla^2 D_m(\mathbf{v}) \in \mathbb{R}^{3 \times 3}$  whose entries are given by:

$$\nabla^2 D_m(\mathbf{v})_{ij} = \frac{\partial^2 D_m(\mathbf{v})}{\partial v_i \partial v_j}.\tag{5.17}$$

It is evaluated using the same information used to evaluate the Jacobian matrix. Based on (4.10), the entries of the Hessian matrix are:

$$\begin{aligned}\nabla^2 D_m(\mathbf{v})_{ij} &= (\Delta t)^2 \left\{ \frac{\mathbf{R}_{a,i} \mathbf{R}_{a,j}}{|\mathbf{R}_a|^3} - \frac{\delta_{ij}}{|\mathbf{R}_a|} \right\} \\ &+ (\Delta t + \tau_f^{\text{aft}})^2 \left\{ \frac{\mathbf{R}_{c,i} \mathbf{R}_{c,j}}{|\mathbf{R}_c|^3} - \frac{\delta_{ij}}{|\mathbf{R}_c|} \right\} - (\tau_i^{\text{fore}})^2 \left\{ \frac{\delta_{ij}}{|\mathbf{R}_d|} - \frac{\mathbf{R}_{d,i} \mathbf{R}_{d,j}}{|\mathbf{R}_d|^3} \right\}.\end{aligned}\tag{5.18}$$

The symbol  $\delta_{ij}$  in (5.18) represents the Kronecker delta, and it is used to delete the term in which it appears for entries off the main diagonal of  $\nabla^2 D_m(\mathbf{v})$ . In practice, one does not know the value of  $\tau_f^{\text{aft}}$ . The definition of the RPC time delay,  $\Delta\tau = \tau_f^{\text{aft}} - \tau_i^{\text{fore}}$ , is used to replace  $\tau_f^{\text{aft}}$  in (5.18) with  $\Delta\tau + \tau_i^{\text{fore}}$ , where  $\tau_i^{\text{fore}}$  is assumed to be known from the geometry and  $\Delta\tau$  is the observed time delay. Using the above expressions for the first and second derivatives of the residual function, and by analogy with (5.11), Newton's method applied to the nonlinear least squares problem is written as:

$$\mathbf{v}^c = \mathbf{v}^p - \left( J(\mathbf{v}^p)^T J(\mathbf{v}^p) + S(\mathbf{v}^p) \right)^{-1} J(\mathbf{v}^p)^T D(\mathbf{v}^p),\tag{5.19}$$

where  $\mathbf{v}^c$  is the current estimate of the vector  $\mathbf{v}$  computed from the previous estimate  $\mathbf{v}^p$ , or from the initial guess. Equation (5.19) is iterated until the norm of the residual is below a specified threshold:

$$\|D(\mathbf{v})\| = \left[ D(\mathbf{v})^T D(\mathbf{v}) \right]^{1/2} < \epsilon_{\text{tol}}.$$

Once this condition is met, the iteration stops and estimates for the  $x$ ,  $y$ , and  $z$  components of the array velocity are obtained. Note that RPC time delays are computed using two consecutive pings. According to the model described earlier, it was assumed that the velocity remains constant over the duration of these two pings. Thus, the estimate represents the velocity at ping  $p - 1$ .

A refinement can be introduced to the solution (5.19) by recalling that the RPC time delays are derived from cross-correlations. These computations yield the lag as well as the correlation coefficient. Thus, it is useful to weight each component of  $D(\mathbf{v})$  by the corresponding cross-correlation coefficient. Let these coefficients be represented by the vector  $\mathbf{w}$  and by the matrix  $W$  which is all zeros except for  $\mathbf{w}$  appearing along the main diagonal. The weighted nonlinear least-squares solution then becomes:


$$\mathbf{v}^c = \mathbf{v}^p - \left( J(\mathbf{v}^p)^T W J(\mathbf{v}^p) + S(\mathbf{v}^p) \right)^{-1} J(\mathbf{v}^p)^T W D(\mathbf{v}^p), \quad (5.20)$$

where  $S(\mathbf{v})$  is modified thus:


$$S(\mathbf{v}) = \sum_{m=1}^M \mathbf{w}_m D_m(\mathbf{v}) \cdot \nabla^2 D_m(\mathbf{v}). \quad (5.21)$$

A useful strategy is to place a threshold on  $\mathbf{w}$  such that components whose value is below the threshold are set to zero. Thus, the unreliable time delay estimates are effectively discarded and the solution is based only on the best measurements available. A sample image that has been motion compensated using this technique is shown in Figure 5.4.

It should be pointed out that the quality of the estimate of the the velocity vector  $\mathbf{v}$  is not the same for all components. In general, poor estimates are obtained for components of the velocity that are orthogonal to the range vectors implicit in the RPC delay estimates. Conversely, more accurate velocity estimates are obtained for components that are nearly parallel to the range vectors. In other words, the NLLS estimate of  $v_x$  is generally good, that of  $v_z$  is less accurate because it is determined mostly by short-range delay estimates, and that of  $v_y$  is generally poor. In fact, the  $v_y$  estimate is typically useless, and the preceding result can be simplified by omitting the  $y$  component of the solution.



(a)



(b)

**Figure 5.4.** The example above demonstrates the effectiveness of the nonlinear least squares (NLLS) motion estimation. The upper image (a) was created without motion compensation, and the lower image (b) was compensated using the NLLS estimate. The scene depicts scattered rocks and a pair of lobster traps on the sea floor. A rope is visible connected to the trap in the lower right quadrant of the image.

## 5.6 Simplified Nonlinear Least Squares

This section outlines the use of nonlinear least squares to find the ping-to-ping displacement vector  $\mathbf{x} = [\Delta x \ \Delta z]^T$  associated with a given pair of RPC channels. This technique is documented in the literature by Cook *et al.* [77]. The solution is identical in structure to that found in the previous section, but in the present case the stop-and-hop assumption is made. This eliminates two of the four range vectors described in (4.10). Furthermore, only the relative horizontal and vertical displacements of the phase centers are estimated using least squares. Once these have been computed, the contribution due to rotation is removed leaving behind the sway and heave estimates.

Referencing Figure 5.5(b), which is a cross-section of Figure 5.5(a), the required time delay is given by:


$$\begin{aligned}\Delta\tau &= 2\Delta R/c \\ &= \frac{2}{c}(R_p - R_{p-1}) \\ &= \frac{2}{c} \left( \sqrt{(x + \Delta x)^2 + (z + \Delta z)^2} - \sqrt{x^2 + z^2} \right),\end{aligned}\tag{5.22}$$

where the subscript  $p$  is the ping number index. As before, it is assumed that the sea floor is a horizontal plane. If bathymetric information is available, for example, from an interferometric array, it can be incorporated in a straightforward manner into (5.22). It is also assumed that the RPC pair channels overlap perfectly in the  $y$  direction, or equivalently, there is no surge error.

In the interest of brevity, the necessary residual, Jacobian, and Hessian are given without added derivations:

$$D_m = c\Delta\hat{\tau}_m - c\Delta\tau_m,\tag{5.23}$$

$$J_{m*} = \left[ \frac{2(x + \Delta x)}{\sqrt{(x + \Delta x)^2 + (z + \Delta z)^2}} \quad \frac{2(z + \Delta z)}{\sqrt{(x + \Delta x)^2 + (z + \Delta z)^2}} \right],\tag{5.24}$$



**Figure 5.5.** The left diagram (a) shows two consecutive pings of a SAS array with two RPC channel. The figure on the right (b) shows a cross-section in the  $x$ - $z$  plane of a pair of redundant phase centers. When the stop-and-hop assumption is made,  $\Delta x$  and  $\Delta z$  do not vary during the ping reception time.

$$\nabla^2 D_{m,ij} = \begin{bmatrix} \frac{2}{\hat{R}_p} - \frac{2(x+\Delta x)^2}{\hat{R}_p^3} & -\frac{2(x+\Delta x)(z+\Delta z)}{\hat{R}_p^3} \\ -\frac{2(x+\Delta x)(z+\Delta z)}{\hat{R}_p^3} & \frac{2}{\hat{R}_p} - \frac{2(z+\Delta z)^2}{\hat{R}_p^3} \end{bmatrix}, \quad (5.25)$$

where the hat indicates quantities evaluated using the current estimate of  $\mathbf{x}$ . As before, the (weighted) numerical solution is written as:


$$\mathbf{x}^c = \mathbf{x}^p - \left( J^T W J + S \right)^{-1} J^T W D, \quad (5.26)$$

with

$$S = \sum_{m=1}^M \mathbf{w}_m D_m \cdot \nabla^2 D_m. \quad (5.27)$$

Figure 5.6 shows an example of this nonlinear least squares technique applied to field data from the SAS12 system developed by the US Office of Naval Research [78]. The center frequency of this research sonar is 180 kHz, making the wavelength equal to 8.33 mm. Note that the total variation of the observed RPC time delays is less than a single wavelength over the 70 m swath shown in the plot. It is clear that the NLLS scheme does an excellent job of estimating the model parameters. This solution is a simplified version of that given in the previous section. However, it is important to note that the curve describing the RPC time delay tends to flatten out with range. Thus, the near range time delays are the






**Figure 5.6.** The plot above demonstrates the typical performance of the nonlinear least squares technique. The broken line shows the measured RPC time delays as a function of range, and diamonds indicate the points used to estimate  $[\Delta x \ \Delta z]^T$  via nonlinear least squares. These points correspond to delays for which the correlation coefficient was equal to, or greater than, 0.9. The solid line shows the time delay as a function of range as computed using the  $[\Delta x \ \Delta z]^T$  estimated by nonlinear least squares.

most important for computing  $\mathbf{x}$ . Fortunately, the near range is precisely where the stop-and-hop approximation holds and where the correlation (alternatively, the signal-to-noise ratio) is the highest.

At this point, only the inter-ping displacement of one element in an RPC pair has been estimated. This displacement can be caused by translations (sway, surge, heave) and rotations (pitch, roll, yaw). AUVs are often outfitted with good-quality angular sensors, so it is assumed that the rotations are known from these. It is then a simple matter to subtract the rotational contribution from  $[\Delta x \ \Delta z]^T$  yielding the desired sway and heave estimates according to:

$$\begin{bmatrix} \Delta x_{\text{sway}} \\ \Delta z_{\text{heave}} \end{bmatrix} = \begin{bmatrix} \Delta x \\ \Delta z \end{bmatrix} - \Phi_p \mathbf{d}_{\text{aft}} + \Phi_{p-1} \mathbf{d}_{\text{fore}}, \quad (5.28)$$


where  $\Phi_p$  is the rotation matrix describing the orientation of the array at ping  $p$  while  $\mathbf{d}_{\text{fore}}$  and  $\mathbf{d}_{\text{aft}}$  represent the positions, relative to the rotation sensor, of the phase centers making




**Figure 5.7.** The inter-ping heave ( $\Delta z$ ) estimate is integrated to yield an estimate of the altitude history for a segment of SAS12 data (solid line). The broken line shows the altitude measured by the on-board navigation system.

up the RPC pair. As an example of the quality of the motion estimates achievable using this technique, Figure 5.7 shows the integrated NLLS heave estimate compared to the altitude as recorded by the AUV's navigation computer. The constant of integration was chosen such that the mean of the NLLS curve coincides roughly with the mean of the measured altitude. Figure 5.8 shows a SAS image produced (a) without motion compensation and (b) using the NLLS technique. This image shows some large-scale sand ripples adjacent to fine-scale ripples. The visibility of the fine-scale ripples demonstrates the quality of focus achieved by the motion estimation and compensation.

As an alternative to using IMU angular measurements to resolve the translational components of the NLLS solution, the array rotations can be estimated and removed from the NLLS estimate using RPC delays just as they are for the SPRPC technique. Briefly, the procedure is as follows. At each ping, one obtains a separate estimate of  $\Delta x$  and  $\Delta z$  for each RPC pair. Simple linear regression can be used to estimate the pitch and yaw by




(a)



(b)

**Figure 5.8.** SAS image before (a) and after (b) motion compensation via simplified nonlinear least squares. The scene depicts the edge of a region of sand ripples on the sea floor. (Figure reproduced from [77].)




**Figure 5.9.** Linear regression used to estimate the sway, heave, yaw, and pitch using the  $[\Delta x \ \Delta z]^T$  estimates for ten pairs of redundant phase centers. The data shown in the plot contains an outlier which was excluded in the computation. The circles represent the points  $[\Delta x \ \Delta z]^T$  in 3D space, while the crosses indicate the projections into the  $x$ - $y$  and  $y$ - $z$  planes.


projecting the  $[\Delta x \ \Delta z]^T$  measurements onto the  $y$ - $z$  and  $x$ - $y$  planes, respectively. This operation is shown in Figure 5.9. The success of this approach depends on the number of RPC channels, the SNR, the precision of the array construction, and the array calibration. Figure 5.10 shows a sample result obtained from this technique. Figure 5.11 shows an image that was compensated using the RPC estimate of the angular motion in lieu of the IMU measurements.

## 5.7 Dual-Sided Closed-Form Solution

The previous two sections detailed the use of nonlinear least squares to estimate the sonar array velocity. The merits of this approach are twofold: first, only a single-sided sonar is required, port or starboard, but not both. Secondly, it can account for the motion of the





(a)




(b)

**Figure 5.10. Vehicle yaw (a) and pitch (b) estimated using the simplified nonlinear least squares RPC technique. The attitude estimates from the AUV's inertial measurement unit (IMU) are shown for reference. Both estimates are unsmoothed.**




**Figure 5.11. SAS image before (a) and after (b) motion compensation showing a trap lying on the bottom of Boston harbor in Massachusetts.**



**Figure 5.12.** The diagram above illustrates the process of computing the time-delay correction needed to map a physical array of receivers into an array of phase centers located along the centerline of the AUV.

array during the time of reception, so the stop-start assumption is not needed. In contrast, this section describes a technique, based on approximations to the model presented above, that can be used in conjunction only with a dual-sided sonar.

A closed-form solution for the sway and heave velocities,  $v_x$  and  $v_z$ , is desired. To derive this solution, several simplifications are made. First, the stop-and-hop assumption is made so that the vehicle is motionless between transmission and reception for any given ping. Second, the along-track velocity is assumed to be ideal (no surge error). Third, the port and starboard arrays are mapped into phase center arrays located along the longitudinal axis, or centerline, of the vehicle. This is shown for a single side in Figure 5.12. To transform the received signal from a bistatic transmitter/receiver pair into an equivalent phase center, all that is required is the appropriate time delay,  $\Delta t_{pc}^m$ . The superscript  $m$  refers to the  $m^{\text{th}}$  receiver. This time delay accounts for the difference between the out-and-back distance associated with the physical transmitter/receiver and the two-way distance associated with



**Figure 5.13.** A dual-sided RPC configuration can be used to derive a closed-form solution to the problem of estimating the ping-to-ping sway and heave.

the corresponding phase center (see Section 1.4). It can be written as:

$$\begin{aligned}\Delta t_{\text{pc}}^m &= 2|\mathbf{R}_{\text{pc}}| - (|\mathbf{R}_{\text{tx}}| + |\mathbf{R}_{\text{rx}}^m|) \\ &= 2 \left| \mathbf{s} - \mathbf{x}(t_i) - \mathbf{p}_{\text{pc}}^m \right| - \left( \left| \mathbf{s} - \mathbf{x}(t_i) - \mathbf{d}_{\text{tx}}^m \right| + \left| \mathbf{s} - \mathbf{x}(t_i) - \mathbf{d}_{\text{rx}} \right| \right)\end{aligned}\quad (5.29)$$

where  $\mathbf{d}_{\text{pc}}^m$  is the vector pointing to the  $m^{\text{th}}$  phase center location. This is defined to be the projection of the usual phase center onto the vehicle centerline:

$$\mathbf{p}_{\text{pc}}^m = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \left( \frac{\mathbf{d}_{\text{rx}}^m - \mathbf{d}_{\text{tx}}}{2} \right). \quad (5.30)$$

The notation of (5.29) closely follows that used to develop (4.10). In fact, the point of mapping the port and starboard arrays to the centerline is to simplify (4.10). The  $m^{\text{th}}$  phase center is now in the same location regardless of whether it is on the port or starboard side, and (4.10) can be simplified:

$$\begin{aligned}c\Delta\tau_s &= 2 \left( R_p^{\text{stbd}} - R_{p-1}^{\text{stbd}} \right) \\ &= 2 \left| \mathbf{s} - \left\{ \mathbf{x}(t_{p-1}) + \Delta t \mathbf{v}(t_{p-1}) \right\} - \left\{ \Phi(t_{p-1}) \mathbf{p}_{\text{aft}} + \Delta t \boldsymbol{\omega}(t_{p-1}) \times \Phi(t_{p-1}) \mathbf{p}_{\text{aft}} \right\} \right| \\ &\quad - 2 \left| \mathbf{s} - \mathbf{x}(t_{p-1}) - \Phi(t_{p-1}) \mathbf{p}_{\text{fore}} \right|\end{aligned}\quad (5.31)$$

In this equation, a subscript has been introduced to distinguish between the RPC time delay of the port and starboard sides. Recall that  $\mathbf{s}$  indicates the location of a scatterer on the



bottom and is taken to be  $\mathbf{s} = [0 \ s_x \ 0]^T$  for a flat bottom. Thus, to write (5.31) for the port side, it is only necessary to change  $\mathbf{s}$  to  $-\mathbf{s}$ :

$$\begin{aligned} c\Delta\tau_p &= 2 \left( R_p^{\text{port}} - R_{p-1}^{\text{port}} \right) \\ &= 2 \left| -\mathbf{s} - \left\{ \mathbf{x}(t_{p-1}) + \Delta t \mathbf{v}(t_{p-1}) \right\} - \left\{ \Phi(t_{p-1}) \mathbf{p}_{\text{aft}} + \Delta t \boldsymbol{\omega}(t_{p-1}) \times \Phi(t_{p-1}) \mathbf{p}_{\text{aft}} \right\} \right| \\ &\quad - 2 \left| -\mathbf{s} - \mathbf{x}(t_{p-1}) - \Phi(t_{p-1}) \mathbf{p}_{\text{fore}} \right| \end{aligned} \quad (5.32)$$

For both of these equations, all the terms not containing the velocity vector are moved to the left-hand side and the result denoted as  $\alpha$ . Therefore, (5.31) and (5.32) become:

$$\begin{aligned} \alpha_s &= 2 \left| \mathbf{s} - \left\{ \mathbf{x}(t_{p-1}) + \Delta t \mathbf{v}(t_{p-1}) \right\} - \left\{ \Phi(t_{p-1}) \mathbf{p}_{\text{aft}} + \Delta t \boldsymbol{\omega}(t_{p-1}) \times \Phi(t_{p-1}) \mathbf{p}_{\text{aft}} \right\} \right| \\ \alpha_p &= 2 \left| -\mathbf{s} - \left\{ \mathbf{x}(t_{p-1}) + \Delta t \mathbf{v}(t_{p-1}) \right\} - \left\{ \Phi(t_{p-1}) \mathbf{p}_{\text{aft}} + \Delta t \boldsymbol{\omega}(t_{p-1}) \times \Phi(t_{p-1}) \mathbf{p}_{\text{aft}} \right\} \right|. \end{aligned} \quad (5.33)$$


Each side can be squared and the two equations solved explicitly for  $v_x$  and  $v_z$ ; remember that  $v_y$  is ideal. The right-hand sides of (5.33) are vector magnitudes. To square these each component is treated separately where  $\mathbf{s}$  is as defined above,  $\mathbf{v} = [v_x \ v_y \ v_z]^T$ , and  $\mathbf{x} = [0 \ -D_p/2 \ z]^T$ , where  $D_p$  is the distance along-track between pings and  $z$  is the altitude of the vehicle at the first ping of the RPC pair. Lastly, define a new vector,  $\hat{\mathbf{p}} = \left\{ \Phi(t_{p-1}) \mathbf{p} + \Delta t \boldsymbol{\omega}(t_{p-1}) \times \Phi(t_{p-1}) \mathbf{p} \right\}$ , which simplifies the task of explicitly writing the components of (5.33).

$$\begin{aligned} \alpha_s^2/4 &= s_x^2 + (\Delta t v_x)^2 + \hat{p}_{\text{aft},x}^2 - 2(s_x \Delta t v_x) - 2s_x \hat{p}_{\text{aft},x} + 2\Delta t v_x \hat{p}_{\text{aft},x} \\ &\quad + (\Delta t v_y - D_p/2)^2 + \hat{p}_{\text{aft},y}^2 + 2(\Delta t v_y - D_p/2) \hat{p}_{\text{aft},y} \\ &\quad + (z + \Delta t v_z)^2 + \hat{p}_{\text{aft},z}^2 + 2(z + \Delta t v_z) \hat{p}_{\text{aft},z} \end{aligned} \quad (5.34)$$

$$\begin{aligned} \alpha_p^2/4 &= s_x^2 + (\Delta t v_x)^2 + \hat{p}_{\text{aft},x}^2 + 2(s_x \Delta t v_x) + 2s_x \hat{p}_{\text{aft},x} + 2\Delta t v_x \hat{p}_{\text{aft},x} \\ &\quad + (\Delta t v_y - D_p/2)^2 + \hat{p}_{\text{aft},y}^2 + 2(\Delta t v_y - D_p/2) \hat{p}_{\text{aft},y} \\ &\quad + (z + \Delta t v_z)^2 + \hat{p}_{\text{aft},z}^2 + 2(z + \Delta t v_z) \hat{p}_{\text{aft},z} \end{aligned} \quad (5.35)$$

Although somewhat menacing at first glance, these two equations are virtually identical. Subtracting the latter from the former eliminates almost all the terms and it becomes easy to solve for  $v_x$ :

$$v_x = \frac{-1}{s_x \Delta t} \left( \frac{\alpha_s^2 - \alpha_p^2}{16} + s_x \hat{p}_{\text{aft},x} \right). \quad (5.36)$$



**Figure 5.14.** The dual-sided closed-form solution was used to compensate the image above. The data was collected by the SAS21 system sponsored by the Office of Naval Research [79, 80]. The uncompensated image (a) is on the left, and the compensated image (b) is on the right. The image shows a small depression in a sandy sea floor that is otherwise covered with sand ripples. Such holes have a number of causes, and a notable example is the crater feeding behavior of dolphins [81]. (Figure reproduced from [80].)

Solving for  $v_z$  is also straightforward, but more tedious. To find  $v_z$  a quadratic equation must be solved,  $v_z = (-b \pm \sqrt{b^2 - 4ac})/(2a)$ , where the positive sign on the discriminant is used, and the  $a$ ,  $b$ , and  $c$  terms are:

$$\begin{aligned}
 a &= 2 \Delta t^2 \\
 b &= 4\Delta t(z + \hat{p}_{\text{aft},z}) \\
 c &= 2(z + \Delta t v_z)^2 + 4(z + \Delta t v_z)\hat{p}_{\text{aft},z} - 2z^2 - 4z\hat{p}_{\text{aft},z}.
 \end{aligned} \tag{5.37}$$

An example showing a SAS image compensated using the dual-sided closed-form solution appears in Figure 5.14. One important aspect of this approach is a consequence of the stop-and-hop assumption. It was earlier shown that the RPC time delays vary with range. However, if the array is stationary during transmit/receive, one is really measuring

the ping-to-ping  $x$  and  $z$  displacements of the vehicle, which do not vary with range (see Figure 5.13). RPC time delay estimates are computed using a sliding window, so for any given pair of pings this solution yields many estimates of  $v_x$  and  $v_z$ . These can then be subjected to a weighted average, as suggested above, where the weight is related to the local cross-correlation coefficient. If the cross-correlation coefficient is below a specified threshold, a weight of zero can be used.

## 5.8 Estimating Platform Surge Using RPC

Another problem of interest in the field of SAS is that of estimating errors in the along-track, or cross-range, direction. The RPC technique relies on the correlation between overlapping channels between two consecutive pings. This implies that the SAS system is capable of measuring its speed with sufficient accuracy to ensure that the RPC channels overlap properly. This distance covered between pings is known as the advance per ping or APP. Errors in APP have two deleterious effects on the SAS operation. First, the array may become spatially aperiodically sampled. As long as the APP error is not severe, this poses little difficulty if the true APP is known. The other effect of APP errors is to reduce the correlation of the RPC pairs. Section 3.3.3 discusses how a constant bias in the APP degrades the focus of the synthetic aperture if it is uncompensated. This section describes three related methods for measuring the APP.

### 5.8.1 Array Correlation Matrix Analysis

The first of the surge estimation techniques addressed here is described in the literature by Groen [82] (also see [83] and [84]). The surge occurring between pings  $p - 1$  and  $p$  is computed from the correlation matrix  $\mathbf{C}^p$  which is formed from the received signals of the

$N$  channels of the SAS array:

$$\mathbf{C}^p = \begin{bmatrix} c_{11} & \cdots & \cdots & c_{1N} \\ \vdots & \ddots & & \vdots \\ \vdots & c_{mn} & \ddots & \vdots \\ c_{N1} & \cdots & \cdots & c_{NN} \end{bmatrix}, \quad (5.38)$$

where

$$c_{mn} = \frac{|\langle \mathbf{s}_m^{p-1}, \mathbf{s}_n^p \rangle|}{|\mathbf{s}_m^{p-1}| |\mathbf{s}_n^{p-1}|}, \quad (5.39)$$

and the angled brackets indicate the inner product operation. Each entry of  $\mathbf{C}^p$  is computed using signals from two consecutive pings. The superscript indicates the ping index and the subscript denotes the array channel index. The signals used in (5.39) are generally subjected to short-time windowing as is done for RPC delay estimation.

Assuming a favorable SNR, the matrix  $\mathbf{C}^p$  exhibits a diagonal with high correlation coefficients, and the position of this diagonal indicates the forward advance of the SAS array between pings  $p - 1$  and  $p$ . If the advance is zero, then the strong correlation would occur along the main diagonal since  $\mathbf{s}_n^{p-1}$  would be nearly the same as  $\mathbf{s}_n^p$ . If the advance is equal to  $\tilde{N} \leq N - 1$  times the distance between phase centers,  $D/2$ , then  $\mathbf{s}_1^{p-1}$  correlates strongly with  $\mathbf{s}_{\tilde{N}+1}^p$ , and the strongly correlated diagonal would begin with the entry  $c_{\tilde{N}1}$  rather than  $c_{11}$ . The other entries of this diagonal are given in Table 5.1. In general, the advance per ping may not be exactly equal to  $\tilde{N}D/2$ , and interpolation must be carried out within  $\mathbf{C}^p$  in order to estimate the true advance per ping, or surge.

This approach to surge estimation has the advantage of being simple and computationally efficient. It does have one serious drawback, however. If the array moves laterally by an appreciable amount, then the received signals for ping  $p - 1$  are no longer time-aligned with the signals for ping  $p$ . In this case, the signals decorrelate and the surge estimate will degrade. The way to circumvent this problem is to abandon the inner product given in

**Table 5.1. Intended RPC channel pairing for an  $N$ -channel array with  $M = N - \tilde{N}$  overlaps.**

Ping $p - 1$	Ping $p$
1	$N - M$
2	$N - M + 1$
3	$N - M + 2$
$\vdots$	$\vdots$
$M$	$N$

(5.39) and replace the entries of  $C^p$  with the peak of the cross-correlation coefficient:

$$\rho_{mn} = \frac{\text{cov}(s_m^{p-1}, s_n^p)}{|s_m^{p-1}| |s_n^p|}. \quad (5.40)$$

This solution is more robust, but it incurs the expense of computing  $N^2$  cross-correlations and finding their peaks as opposed to computing  $N^2$  inner products.

### 5.8.2 Exhaustive Surge Error Estimation

If there is no error in the advance per ping (APP), then the overlapped RPC channels in Table 5.1 will be highly correlated. As the APP error grows, the correlation between the intended RPC pairs diminishes and the correlation between other combinations of channels will rise. As above, this effect is exploited in order to measure APP. Using the RPC channels to estimate the surge requires multiple-channel overlap between pings, and thus there is a certain price to pay in terms of area coverage rate for the ability to measure APP errors. (The relationship between range coverage and array length is discussed in Section 1.3.4.) This section presents a surge estimation technique similar to the one above and was first set forth in the SAS literature by Oeschger [85].

As before, it is assumed that there are  $N$  channels in the array and  $M = N - \tilde{N}$  of them overlap for RPC use. The intended channel pairs are given in Table 5.1. For example,  $M = 2$  in Figure 1.7, and  $M = 3$  in Figures 5.2 and 5.3. If the APP error is zero, then this table accurately describes the phase center positions relative to one another. Conceptually, APP error causes the right column of the table to shift up or down by some unknown amount.

A positive APP error means that the vehicle speed is estimated to be less than it actually is. Consequently, the SAS ping will be triggered late and the distance between consecutive pings too large. A negative APP error indicates that the vehicle speed is estimated to be faster than its true value resulting in the SAS being triggered too soon. The distance between pings will then be too small.

Since the true APP may be grossly in error, the pairwise signal correlations given in Table 5.1 may be too low to use for motion estimation. This problem can be solved by using the available overlapping channels to search for the combination that yields the highest correlation. In other words, a given RPC channel (say,  $N - M + 1$ ) is not only correlated with its intended counterpart (2), but also with the neighbors of that counterpart (1 and 3, for example). This notion suggests the following algorithmic procedure for estimating the actual APP.

1. Define a reference channel,  $s_{\text{ref}}$ , from the overlapping portion of the array at ping  $p$ . This is the right-hand column of Table 5.1.
2. Correlate  $s_{\text{ref}}$  with each channel of the previous ping,  $p - 1$ , using the short-time procedure described in Chapter 4. This requires  $N$  correlation operations.
3. Find the channel,  $n$ , that is most highly correlated with  $s_{\text{ref}}$ . This gives the APP error to within an integral number of phase centers.
4. Using the three correlations of  $s_{\text{ref}}$  with channels  $n - 1$ ,  $n$ , and  $n + 1$ , perform a quadratic interpolation (4.21) to estimate the fractional part of the APP error.

This procedure is capable of measuring any amount of negative APP error, since ping  $p$  cannot occur before ping  $p - 1$ . However, the observable positive APP error is limited by the number of overlapping phase centers. This is seen by considering Table 5.1. Positive APP error is sensed by correlating the signal  $s_{\text{ref}}$  with channels forward of that channel that matches  $s_{\text{ref}}$  in Table 5.1. In the table, these channels are found in the left-hand column.

The largest observable positive APP error corresponds to  $M - 1$  phase center spacings, or  $(M - 1)D/2$ , where  $D$  is the element width. If  $s_{\text{ref}} = s_N$ , then correlation with  $s_1$  from ping  $p - 1$  represents the upper bound on the observable positive APP error.

This technique for measuring the APP error is computationally expensive. The cost can be mitigated to a certain extent, however. As a matter of practice, it is not always necessary to correlate  $s_{\text{ref}}$  with all  $N$  elements of the array. For example, if one knows that the vehicle's on-board navigation can reliably estimate the APP to within  $\pm 1$  phase center spacings,  $\pm D/2$ , then fewer correlations need to be performed. If the APP is expected to be in gross error, then it must be estimated before performing the cross-track motion estimation: the APP estimate will be used to first find the RPC channel combination yielding the highest correlation. Otherwise, using the intended RPC correlation pairs may result in bad cross-track motion estimates.

### 5.8.3 Surge Estimation via Eigendecomposition

A simple and elegant solution exists to the problem of surge estimation if it can be assumed that the delay observed by the RPC channels is less than a range resolution cell. The first step is to choose the  $n^{\text{th}}$  channel from ping  $p$  as a reference signal,  $s_{\text{ref}} = s_n^p$ . The reference channel is generally one of the aft channels used for RPC delay estimation. Next, form a matrix whose columns are the time series from the previous ping and whose locations along the array bracket the cross-range position of the receiver,  $m$ , which would ordinarily form the RPC match with the reference channel in the absence of any surge error:

$$\mathbf{S} = \begin{bmatrix} s_{m-u}^{p-1} & \cdots & s_m^{p-1} & \cdots & s_{m+v}^{p-1} \end{bmatrix}. \quad (5.41)$$

The indices  $u$  and  $v$  determine how many channels are searched in the surge estimation. For example, an array of seven channels might have three of them dedicated for RPC use (see Figure 5.3). One choice for  $s_{\text{ref}}$  would be channel six, assuming the array is indexed from one to seven beginning at the foremost element. Then, setting  $u = v = 1$  gives  $\mathbf{S} = \begin{bmatrix} s_3^{p-1} & s_2^{p-1} & s_1^{p-1} \end{bmatrix}$ . Using this particular configuration allows the surge to be estimated as

long as it is within  $\pm D/2$ , where  $D/2$  is the horizontal separation of the phase centers of the vernier array. Using more channels in  $S$  allows for the estimation of greater amounts of surge.

In the absence of APP error, the signal  $s_{\text{ref}} = s_n^p$  will be highly correlated with  $s_m^{p-1}$ . The APP error is estimated by finding the weight vector  $\mathbf{w}$  that maximizes the quantity

$$\begin{aligned} \|(S\mathbf{w})^* s_{\text{ref}}\|^2 &= (S\mathbf{w})^* s_{\text{ref}} s_{\text{ref}}^* S\mathbf{w} \\ &= \mathbf{w}^* S^* s_{\text{ref}} s_{\text{ref}}^* S\mathbf{w} \\ &= \mathbf{w}^* R\mathbf{w}, \end{aligned} \tag{5.42}$$

where  $R$  is a correlation matrix. In other words, the weight vector sought combines the columns of  $S$  so as to produce the best estimate of the vector  $s_{\text{ref}}$ . The condition is also imposed that the norm of  $\mathbf{w}$  be equal to one. From these requirements, the following Lagrangian is formed:


$$L(\mathbf{w}, \lambda) = -\mathbf{w}^* R\mathbf{w} + \lambda (\mathbf{w}^* \mathbf{w} - 1). \tag{5.43}$$

Taking the derivative of  $L(\mathbf{w}, \lambda)$  with respect to  $\mathbf{w}$  and setting the result equal to zero gives  $-R\mathbf{w} + \lambda\mathbf{w} = 0$ , or  $(R - \lambda I)\mathbf{w} = 0$ , whose solutions are the eigenvalues and eigenvectors of the matrix  $R$ . The matrix  $R$  has certain properties because it is a correlation matrix. Among these are that  $R$  is Hermitian symmetric ( $R$  equals its conjugate transpose,  $R^*$ ) and that it is positive definite (that is,  $\mathbf{w}^* R\mathbf{w} > 0$ ,  $\forall \mathbf{w} \neq \mathbf{0}$ ). These properties imply that the eigenvalues of  $R$  are real and strictly positive.

The solutions of (5.42) and (5.43) represent the minimum, maximum, and saddle points, where the maximum (minimum) value corresponds to the largest (smallest) eigenvalue. This result, known as a maximum (or minimum) principle, is found in standard linear algebra texts in connection with the discussion of quadratic forms and the Rayleigh quotient [86, 87]. It is also a consequence of the Courant-Fischer min-max theorem.

Since the maximum value of (5.43) is sought for the present application, the desired solution is the eigenvector corresponding to the largest eigenvalue. The columns of  $S$  correspond to physical locations along the receiver array. Thus, the APP can be thought of as





**Figure 5.15. Comparison of the speed over ground (SOG) estimated computed using the exhaustive correlation and eigendecomposition methods. The speed estimate computed by the AUV's on-board navigation system is also shown for reference.**

the location of the peak found by interpolating within the entries of the weight vector  $\mathbf{w}$ . A comparison of the results obtained using the exhaustive search and eigenanalysis methods is shown in Figure 5.15.

As with the other operations using RPC, this computation is performed locally using a short-time sliding window. The results are then averaged using the correlation coefficients as weights. In fact, this technique fits nicely with the time delay estimation, as the necessary correlation coefficients will have already been computed. The matrix  $R$  will usually be of fairly low rank, so the required eigendecomposition is inexpensive compared to the cost of computing the extra cross-correlations used for the general APP estimation. In fact, a complete eigendecomposition is unnecessary since only the largest eigenvalue is required. Thus, this approach is amenable to numerical techniques, such as the power method [88], which can compute the eigenvalues in descending order in terms of magnitude. Lastly, it is worth observing that this technique for estimating the APP error of a SAS array bears a

resemblance to certain methods of adaptive array processing [68] used for signal detection and estimation.

## 5.9 Non-RPC Data-Driven Motion Estimation Techniques

The majority of this chapter is dedicated to motion estimation techniques based on redundant phase centers. These are robust and commonly used in practice. However, there are other data-driven methods for estimating the motion of the SAS array which do not use redundant phase centers. One significant factor motivating interest in these approaches is the fact that, for a given number of array elements, dedicating receivers to RPC usage lowers the area coverage rate of the SAS.

A technique for estimating sway is given by Johnson *et al.* [89]. In this case, the sway is estimated by examining the statistics of the ping-to-ping phase variation of the received echoes. This method is based on several assumptions, not the least of which is the assumed absence of isolated bright scatterers. Fortunately, it is relatively simple to detect and reject the returns from these scatterers, and the backscatter from the sea floor is otherwise frequently observed to be reasonably homogeneous.

Another example of non-RPC-based motion estimation is given by Gough and Miller [90] and is capable of estimating both the ping-to-ping sway and yaw of a SAS vernier array. The method is based on the idea of producing low-resolution images using only the receivers associated with a single ping (namely, all the receivers in the array). The resulting images from adjacent pings are compared using cross-correlation, from which the differential sway and yaw are inferred.

## 5.10 Comments

With the exception of the slant plane RPC motion estimation, the cross-track motion estimation techniques presented above contain explicit reference to an average location of a scatterer on the sea floor. The typical workaround is to assume that the entire sea floor is

a horizontal plane. This is a reasonable assumption for many environments, but there are circumstances in which it may cause unacceptable error in the motion estimate. In general, assuming a planar bottom is acceptable. Prior knowledge of the bathymetry or an on-board interferometric capability could be used to dispose of any assumptions regarding the contours of the sea floor.

Another subtlety that bears mentioning regards the short-time correlation coefficient obtained in the time-delay computations. It is common practice to use these values as both a threshold and weight in subsequent motion estimation schemes. The obvious approach is to simply use the correlation coefficients directly as weights after thresholding at a particular value, keeping  $\rho \geq 0.85$  for example. However, it may be preferable to give substantially greater weighting to high values of  $\rho$ . In such cases, one might weight by  $1/(1 - \rho)$ . Regardless, it should be noted that the performance of SAS motion estimation is dependent on the choice for threshold value and weighting method, and that it is important to discard the RPC time delays corresponding to low SNR conditions.

## CHAPTER 6


### MOTION COMPENSATION

#### 6.1 Introduction


This chapter addresses the actual process of compensating for unwanted platform motion. There is a variety of tools for this purpose, and the specific choice of technique depends strongly on the desired image quality, system parameters, and the permissible computational expense. Motion compensation amounts to merely reversing the time delays induced by the platform motion, but this can be a complicated matter in general. The subsequent discussion assumes that the motion compensation is a separate step in the overall processing flow. That is to say, the compensation serves to condition the data prior to being reconstructed with an algorithm that only treats the ideal straight line collection geometry. It is this approach that necessitates the creation of the various compensation schemes. No separate compensation step is required when the general form of spatiotemporal backprojection is used for image reconstruction. In this case the price for perfect motion compensation is steep, and excellent imagery can usually be obtained with much less computational effort.

#### 6.2 The Range-Varying Nature of Motion Compensation

The amount of the shift required for compensation depends on the particulars of the motion encountered, the range, and the bathymetry. (See [91] and [92] for a discussion in the context of SAR.) This is illustrated in Figure 6.1 in which  $\delta$  is the vector pointing from the nominal trajectory (NT) to the actual location of the sensor in question, that is, a particular channel of a given ping. It is important to point out that  $\delta$  is not a ping-to-ping motion vector. Instead, it is the result of the motion estimation in which the platform velocity has been integrated to give a history of the array position relative to the NT. Recall from Chapter 5 that the nominal trajectory is the line in 3D space to which the SAS data is being corrected. In the present discussion, the NT points out perpendicularly from the



(a)



(b)

Figure 6.1. The upper figure (a) shows the geometry used to describe the range varying time delay correction. The drawing is not to scale in the sense that a far-field assumption is made in which  $|\delta| \ll |R|$  implying  $R$  is nearly parallel to  $R'$  and thus  $R \approx R' + \delta \cdot e$ . This simplification leads to the idea that the compensation delay for a particular range is given by the projection of  $\delta$  onto a given  $R$ . Using this idea, the total variation over the range swath of the time delay correction is found to be proportional to  $\delta \cdot (e_{\text{far}} - e_{\text{near}})$  as shown in the lower figure (b) (in which  $R'$  is omitted). The maximum variation occurs when  $\delta \parallel (e_{\text{far}} - e_{\text{near}})$ .

page in Figure 6.1. Also shown are the vectors  $\mathbf{R}$  and  $\mathbf{R}'$ , along with their corresponding unit vectors, which point to a given location on the sea floor from the NT and the sensor, respectively. The time delay necessary for motion compensation is proportional to the difference between  $\mathbf{R}$  and  $\mathbf{R}'$ ,

$$t_{\text{correct}} = \frac{2}{c} (|\mathbf{R}| - |\mathbf{R}'|). \quad (6.1)$$

At this point a far-field approximation is made in which it is assumed that  $|\delta| \ll |\mathbf{R}|$ . This implies that  $\mathbf{R}$  is effectively parallel to  $\mathbf{R}'$  which in turn implies:

$$\begin{aligned}\mathbf{R} &\approx \mathbf{R}' + \delta \cdot \mathbf{e} \\ &= \mathbf{R}' + |\delta| \cos \theta_{\text{err}},\end{aligned}\tag{6.2}$$

where  $\theta_{\text{err}}$  is the angle between  $\delta$  and  $\mathbf{e}$ . It is here that the bathymetry comes into play and the common assumption of a horizontal planar sea floor occurs. It is impossible to determine  $\theta_{\text{err}}$  without knowledge of both  $\delta$  and  $\mathbf{e}$ . The previous chapter explained how to estimate  $\delta$ , but  $\mathbf{e}$  is unlikely to be known unless the SAS in question happens to possess an interferometric capability that can provide an estimate of the sea floor topography. If the RPC motion estimation delays are known for all ranges in the swath, then knowledge of the bathymetry is not required. However, it becomes necessary to interpolate or extrapolate delays for those ranges for which RPC delays are unavailable. The planar assumption regarding the sea floor is simply one possible way of determining  $\mathbf{e}$  so that the motion compensation delay can be computed for all ranges.

### 6.3 Range Compensation Using a Single Delay

Referring to Figure 6.1, the required time shift is clearly a function of range. If this variation can be ignored, then motion compensation can be done by applying the same time shift at all ranges, for example, the shift associated with the scene center,  $R_0$ . This is the best-case scenario as it is easy to implement using the time-shift property of the Fourier transform. As such, it is well-matched to the wavenumber-based image reconstruction algorithms since it can be performed after the temporal Fourier transform and prior to the spatial Fourier transform. The single-delay approach can also be used to create a piecewise approximation to a continuous delay. For example, a scene might be broken into two or more subswaths with each receiving a single compensation delay. Then, each of the subswaths could be imaged independently using wavenumber-based reconstruction.

Another single-delay compensation scheme applicable for narrowband systems is to

divide the delay by the temporal sampling period giving the delay in terms of samples. The integral part of the delay is then effected by reindexing the signal, and the fractional part (expressed in seconds) is introduced via phase rotation:  $\exp \{-j2\pi f_0 t_{\text{correct,frac}}\}$ .


## 6.4 Range Compensation Using Variable Delays

There is often the need to apply a time (or range) varying delay. For small displacements, relative to the size of a resolution cell, a simple phase rotation suffices for narrowband sonars. For a given cross-range position this is implemented by multiplying each sample in the received time series by  $\exp \{j2k_0 \Delta R(t)\}$ , where  $\Delta R(t)$  is the time-varying displacement to be compensated. If the motion error is severe enough to shift the envelope of the received signal through resolution cells, then the phase shift alone cannot achieve the desired result. One option for dealing with this is to perform a bulk shift of the entire time series as described in Section 6.3. The amount of this bulk shift is chosen by the user and might correspond to the range-averaged delay or perhaps the value of the delay at the swath center. Any residual delay (that is, the delay after the bulk shift is removed) can be removed using a phase rotation provided it does not also shift the envelope of the received signal through resolution cells.

Sometimes the compensation delay varies by more than a resolution cell over the range swath, and/or it is not possible or desirable to subdivide the swath. In these circumstances, there is little choice but to resample the time series by interpolating it from the actual acquisition times onto the set of times dictated by the continuously-varying compensation delays.

## 6.5 Crab Compensation

Crab occurs when an AUV or towfish points into a cross-current in order to maintain the desired track over the sea floor (see Figure 6.2). The result of crab is that the boresight of the physical SAS array is no longer perpendicular to the nominal trajectory. In other words,



**Figure 6.2.** Crab occurs when an AUV is forced to place its angle of attack into a cross current in order to maintain a desired track over the sea floor. Crab compensation can be performed in two ways. The first is to beam steer the physical arrays by the angle  $\theta_{\text{crab}}$  in order to align them with the nominal trajectory (NT) defined by the vehicle's track over the sea floor. The second approach is to define another compensation baseline  $NT_{\text{sway}}$  making an angle  $\theta_{\text{crab}}$  with the vehicle track and then treat the crab as sway.

crab is a form of array squint.

Crab can be compensated in two different ways. The first and most obvious is to beam steer the signals for each ping, effectively moving the squint angle of the physical array back to zero. This approach is easy to implement, but not all SAS array designs are amenable to it. At the time of acquisition, the projector and receivers are always pointed in the same direction. The notion of crab only arises when the synthetic aperture is formed according to the chosen NT and the constituent physical arrays are found to be looking off axis. When these arrays are steered in order to correct for the perceived crab, the array elements must be spaced closely enough to ensure that the beam steering doesn't cause spatial aliasing. Unfortunately, this requirement demands that the array element spacing be finer than  $D/2$ . When designing the SAS array, it is wise to consider the amount of crab that might be encountered by the system.

The second approach for compensating crab is to reassign the nominal trajectory so that it is aligned with the crabbed longitudinal axis of the vehicle, as shown in Figure 6.2. By doing so, the crab ceases to look like a squint and is transformed into a sway. Recall that the choice of nominal trajectory is completely arbitrary, although its placement is generally



dictated by both common sense and processing constraints. This technique has the distinct advantage of liberating arrays sampled at  $D/2$  from the problem described in the previous paragraph. The drawback is that the near-range and far-range boundaries of the scene appear to creep, or skew, in range. Although the contents of the scene are not geometrically distorted by this operation, a large crab angle can cause some image content to be lost and is visually unappealing. A practical compromise is to correct for the crab by beam steering as much as is allowed by the physical array. The remainder of the crab is then treated as a sway, thus minimizing the loss of image content.

Crab is a common form of unwanted motion in AUV-borne SAS, and it should be considered when planning data collection missions. For example, if an AUV is scheduled to image an area of the sea floor, then it is best to operate at slack tides when currents are at a minimum. Barring that option, then it is best to operate the vehicle parallel to the current, either with or against it. Commanding the AUV to drive perpendicularly to a swift current can result in SAS data that is difficult, if not impossible, to properly compensate.

## **6.6 Compensation for Angular Motion**

Compensation for more general angular motion requires procedures similar to those described for crab. The difference, though, is that crab is a static angular displacement of the array, while the typical movements of an AUV are oscillatory in nature. There is therefore little choice for compensation except to perform beam steering at each ping to shift the squint angle to zero. Compensating for large amplitude oscillations can result in localized spatial aliasing of the physical arrays and, consequently, localized image degradation that would appear as a blur running down range. For this reason, the specification of the SAS array phase center spacing should take into account the statistics of the motions of the vehicle intended to carry the SAS.

Unwanted angular motion can change the performance of the SAS even if the geometric errors are perfectly compensated using appropriate time delays. The synthetic aperture


length is idealized as being equal to the -3 dB width of the array element horizontal beam-pattern. The synthetic aperture length, in turn, determines the cross-range resolution. It is possible for angular motions to change the pointing angle of the array so as to extend or diminish the synthetic aperture, resulting in potential spatial aliasing and/or resolution loss.

## 6.7 Cross-Range Compensation

There is the potential need to compensate in the cross-range, or along-track, direction. Here, there is little choice except to interpolate in the cross-range direction. This operation has its limits. As above, the sampling of the SAS array influences the effectiveness of the compensation. Sampling a synthetic aperture at  $D/2$  phase center spacing is analogous to sampling a time series at the Nyquist rate. This analogy is only approximate. See Chapter 5 of [8] and [13] for an in-depth treatment of SAS sampling. If the  $D/2$  sampling requirement is met at all points on the synthetic aperture, then it is possible to interpolate the data onto any desired cross-range grid. If the synthetic aperture is undersampled at any point, then the cross-range interpolation may give results that degrade the imagery rather than improving it. Since vernier SAS arrays are usually designed with phase center spacing satisfying  $du \leq D/2$ , undersampling can only occur when the advance per ping is so large that there is no overlap between consecutive pings. This fact is yet another motivation for designing SAS arrays with some degree of overlap for RPC use.



## 6.8 Examples of Range-Varying Motion Compensation

Equation (6.2) along with the assumption of a planar sea floor indicate that the compensation delay always approaches some asymptotic value as the range is increased. This is true to the extent that the stop-start assumption is valid. If the intra-ping array motion must be considered, then a limiting delay value does not necessarily exist. In this case, the  $\delta$  above would need to be rewritten as a function of time, or range, rather than being assumed constant for the entire time of reception.



**Figure 6.3. Helical motion with a diameter of 0.25 m at a mean altitude of 7 m.**

Figures 6.3 and 6.4 show an example of how the compensation error can vary if the SAS array experiences a helical motion. Such motion is not uncommon in practice, as it can occur when the direction of the AUV's travel makes a nonzero angle with the direction of the ocean waves. A helical trajectory often results if the vehicle is close enough to the sea surface to be influenced by the wave action. The vehicle motion is shown in Figure 6.3 along with the nominal trajectory. Figure 6.4(a) shows the compensation delay that remains after compensating the entire range swath using the delay computed for a range of 45 m. Cross-range slices of this residual delay are shown in Figure 6.4(b) for ranges of 15 m and 80 m. Curves such as these can be used in conjunction with the results of Chapter 3 to predict how the PSR degrades as a result of using a single delay to compensate all ranges. For example, suppose the synthetic aperture length were 5 m at the closest range of 15 m. Then, the residual compensation delay behaves like a phase error that is locally quadratic




**Figure 6.4.** The upper figure (a) shows the residual motion compensation error visualized as a surface. Note that the residual error at a range of 45 m is zero, because this is the range at which the bulk compensation was performed. The lower plot (b) shows the residual error at ranges of 15 m and 80 m. To analyze the effect of the residual compensation error on the PSR, the error must be considered only over the spatial extent of the synthetic aperture.

or linear in nature, depending on cross-range position. Thus, alternating smeared and well-focused regions in the cross-range direction would be observed in the image, provided that the wavelength is not significantly larger than the residual error. The period of the example motion is 15 m. At a range of 80 m where the hypothetical synthetic aperture length is 25 m, the phase error is sinusoidal in nature. At this range, the sea floor would be well focused, but would exhibit replicas (or ghosts) in the cross-range direction.


For well-tuned AUVs operating in calm conditions, the range varying nature of the compensation can often be ignored without having a dramatic impact on image quality. This is not always the case, however. The angle  $\theta_{\text{err}}$  changes rapidly at short ranges and rather slowly at long ranges. Thus, it is reasonable to expect range-varying compensation to be a requirement for proper short range imaging. An example of this effect is shown in Figure 6.5. For this scene, the range of the object shown and the altitude of the SAS are such that the depression angle of the object relative to the array boresight falls close to the first null of the vertical beam pattern of the array. In fact, the null is visible in the imagery as a dark vertical band at a range of 13 m. An adaptive gain has been applied to the imagery to equalize the overall brightness in the scene. It is important to remember that the visual appearance due to this gain can be deceptive: the signal-to-noise ratio outside the main lobe is often poor.

## 6.9 Comments

Figure 6.6 shows a diagram of the motion compensation taxonomy that is suggested by the preceding discussion. It shows the algorithmic options that are available to address the various types of motion compensation encountered in SAS imaging. Note that spatiotemporal backprojection is a world unto itself, so to speak, in that it is the exact solution for any array trajectory. All other motion compensation is approximate to some extent. Nevertheless, the combination of a relatively stable vehicle and good estimates of the motion can often be used with fast compensation schemes and Fourier image reconstruction to

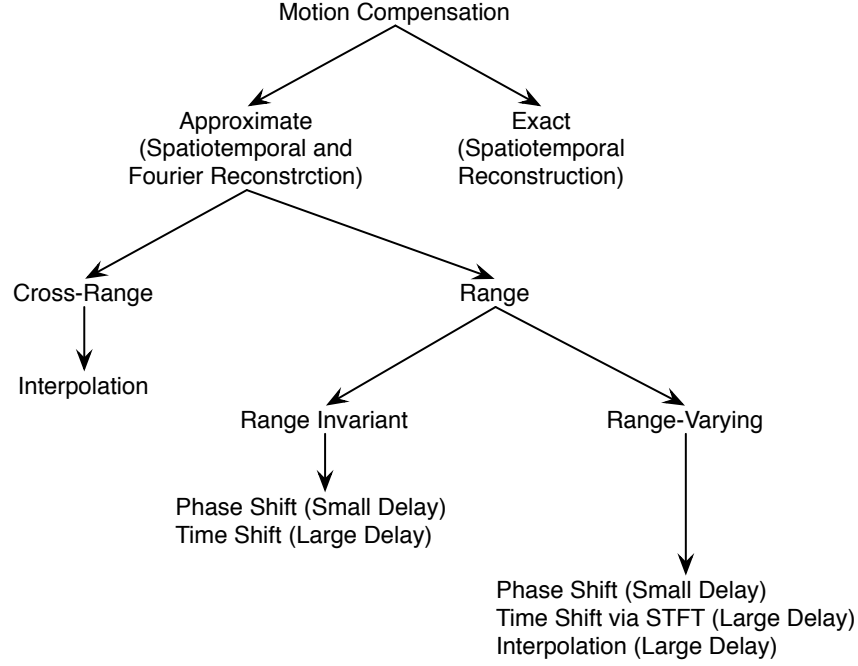


(a)



(b)

**Figure 6.5.** Short ranges are especially susceptible to the defocus caused by compensating for the motion using a single delay for all ranges. The top image (a) shows an example of this effect, while the bottom image (b) resulted from applying a range-varying compensation scheme. The dark vertical band appearing at a range of 13 m is the first null of the vertical beam pattern.



**Figure 6.6. A taxonomy of motion compensation schemes.**

efficiently yield imagery that is indistinguishable from its backprojected counterpart.

Another complication to the problem of motion compensation arises for widebeam SAS. The discussion throughout this document assumes that, at a given range, the compensation delay is the same for all scatterers in the sensor's beam. This is an idealization since all real beamwidths are finite, and the error for scatterers located at an angle of  $\theta$  relative to boresight is  $\epsilon_R = (1 - \cos \theta)$ . Techniques for widebeam compensation have been proposed by several researchers [93, 94, 95, 96]. A solid overview is given by Callow [37]. As widebeam SAS becomes more prevalent, widebeam compensation will become increasingly important.

## **CHAPTER 7**


### **CONCLUSION**

The potential usefulness of synthetic aperture sonar is far-reaching. SAS can be used in any application that is suitable for high-resolution real aperture sidescan sonar. Some examples are given below.

- Maritime archaeology [97, 98, 99, 100]
- Ecological and fisheries management [101, 102, 103]
- Environmental remediation [104]
- Naval applications
- Oceanography
- Salvage
- Search and recovery [105, 106]
- Surveying
- Crime scene investigation and law enforcement [107]

Apart from having range-invariant resolution, SAS has the added benefit that its imagery is geometrically accurate when it is well-focused. This feature makes it relatively easy to georeference SAS imagery, resulting in detailed maps of the sea floor in which distances can be accurately measured. A well-focused SAS image may exhibit geometric image distortion such as layover, but it will not show the warping due to platform motion that is commonly seen in real aperture sidescan sonar. This property is a consequence of the fact that the sensor motion relative to a point on the sea floor must be known in order to properly reconstruct the image in the first place.






**Figure 7.1.** The photos above show a pair of World War II vintage aircraft. The left photo (a) depicts a Fairey Firefly, a British carrier-based fighter. The right photo (b) shows a Curtiss SB2C Helldiver, an American dive bomber. The Helldiver was responsible for sinking more tonnage of shipping than any other aircraft during World War II. SAS imagery of these aircraft is shown in Figure 7.2.


## 7.1 Example SAS Imagery

This section presents a collection of sample imagery that highlights the ability of SAS to image relatively large regions of the sea floor at high resolution. All of these images were created using the motion estimation and compensation techniques described in this thesis, and no autofocus schemes were used for enhancing the reconstructed imagery. The first of these is Figure 7.2 showing SAS images of the aircraft pictured in Figure 7.1. Both are World War II era carrier-based airplanes: the first a Fairey Firefly fighter and the second a Curtiss SB2C Helldiver dive bomber. It can be clearly seen in the imagery that the port wing of the Firefly is buried in the sediment, and the upper two blades of its four-blade propeller are unbent. The implication of the latter observation is that the propeller was not turning when the aircraft landed in the water. The Firefly rests in the waters of Jervis Bay, Australia. The image of the Helldiver plainly shows that the horizontal stabilizer was detached and now sits just aft of the port wing. The image also shows the internal structure of the wings. This aircraft was ditched by its pilot in Buzzards Bay, Massachusetts.

A shipwreck is shown in Figure 7.3. The vessel is the tugboat *Vittoria*. It was fitted with armament during World War II and sank as a result of a torpedo attack on the vessel it was towing. It rests upright on the floor of the Ligurian Sea off the coast of Italy. Many details are clearly visible on the deck of the boat, such as the outlines of the railing and the




(a)



(b)

**Figure 7.2.** The top image (a) shows a SAS image of a Fairey Firefly lying on the sea floor in Jarvis Bay, Australia. It crashed as a result of a midair collision on November 27, 1956. Today, the wreckage is popular with scuba divers. The bottom image (b) is of a Curtiss SB2C Helldiver in Buzzards Bay, Massachusetts. It was ditched by Naval reservist Ensign John L. Hagerman after the aircraft's engine lost oil pressure and stalled during a training mission on July 7, 1947.




**Figure 7.3.** The image above shows the World War II era tugboat *Vittoria* sitting upright on the sea floor off the coast of Italy.

superstructure. The shadow cast by the wreck clearly reveals that it is a tug, and the rudder is even visible beneath the stern.

Figure 7.4 contains another image from the Ligurian Sea. In this case, the scene depicts segments of a pipeline that has fallen into disuse. The fact that most of the pipe segments are not connected is easily discerned. The pipe flanges are also visible as bright highlights. Nearby, a narrower modern pipeline is visible. The newer pipeline is secured by regularly spaced anchor blocks.

## **7.2 Approaches to Synthetic Aperture Imaging**

In reading the literature pertaining to synthetic aperture imaging, one notices that the writers come from a variety of backgrounds and that they bring their own perspective to bear



**Figure 7.4.** The image above shows a pair of underwater pipelines off the coast of Italy. The larger of the two is clearly not functional because many of its segments have become disconnected.

on the problem. Before concluding, some of the major approaches to the study of synthetic aperture are highlighted. There are at least four ways to view the topic:

1. **Mathematics:** The mathematical viewpoint is attractive because of its power and clarity. However, it can be difficult to convey the details of implementation in this context. The mathematical approach was used in Chapter 2 to show how synthetic aperture imaging is developed from linear acoustics and the wave equation. Mathematical results also provide the certainty of convergence for techniques such as the nonlinear least squares used in Chapter 5. Another advantage to the mathematical standpoint is that the creation of a synthetic aperture image is an inverse problem, and there is a vast body of literature on this subject.
2. **Signal Processing:** The signal processing point of view is an important one because

most of the familiar results from one-dimensional processing of time series have direct and easily-understood analogs in the cross-range spatial dimension. These provide many of the tools needed to properly design a SAS and predict its performance. The paper by Munson and Visentin [108] gives an excellent signal processing interpretation of synthetic aperture imaging. The signal processing approach is important also because it is the language for actually implementing synthetic aperture algorithms efficiently (for example, sampling at the bandwidth and processing at baseband as opposed to processing data at twice the Nyquist rate).

3. Optics: The fields of geometric optics and statistical optics serve as the bases for explaining many aspects of synthetic aperture imaging. For example, Equations (3.9) and (3.13) in Section 3.3.3 were derived mathematically. It was mentioned that the same pair of results can be obtained using simple analogies from geometrical optics. In addition, the important results presented by Pinto and Bellettini [62] regarding the Cramer-Rao lower bounds on redundant phase center time delay estimation rely partially on a result from statistical optics known as the van Cittert-Zernike theorem [109, 110].
4. Probability: The probabilistic view of synthetic aperture imaging discards the idea that the operation of image reconstruction is the coherent integration of echoes along a deterministic path (see (2.21)). Rather, this integration is viewed as being over all the possible paths, the distribution of which is weighted by the appropriate probability density function. In this respect, the image reconstruction problem is not unlike other problems in navigation and optimal signal estimation. Indeed, it is often stated that the image reconstruction process is a sort of spatial matched filter. This is true, but the statement is usually made primarily as an analogy to one-dimensional operations like pulse compression, and isn't meant to describe an actual methodology for computing imagery. To date, the probabilistic approach is the least explored.

Clearly, one should not be surprised to find that there are often several ways to explain any given aspect of synthetic aperture imaging. It should also be recognized that there is great potential for making advances by bringing seemingly unrelated ideas from other fields to bear on the problem. For example, the wavenumber domain image reconstruction was well-known in the geophysics community for years before it came to revolutionize the creation of synthetic aperture imagery. It is currently the dominant technique for image reconstruction. Although synthetic aperture sonar is a largely mature technology [111], there are likely to be many more breakthroughs that will enhance the performance and expand the capability of SAS systems.

### **7.3 Summary of Contributions**

Properly accounting for the platform motion is an essential part of a robust SAS processing suite, and a thorough treatment of the subject has heretofore been absent from the literature. The overall contribution of this work is to provide a comprehensive overview of the problem of SAS motion estimation and compensation. In addition, several specific contributions to the field of synthetic aperture imaging were realized in the course of this research:

- Chapter 3 gives an analysis of the effects of phase errors for stripmap mode synthetic aperture processing [112]. The analysis for the companion problem of spotlight mode operation is well-documented [24], but prior coverage of the stripmap case is less complete [52, 54].
- Section 4.3.1 contains a novel method for finding the peak of the magnitude of the cross-correlation function. The traditional approach is to oversample the correlation function and then locate the peak using a parabolic fit [17]. The proposed technique solves iteratively for the peak without the need for using oversampling as a tool to achieve improved accuracy.

- Three techniques for estimating platform motion were developed: the two-sided closed-form solution [80], the nonlinear least squares approach [77], and the surge estimation using eigendecomposition.

## 7.4 Final Comments

It goes without saying that the quality of SAS imagery is directly related to the quality of the motion estimation and compensation. As a result, the purpose of this document is to offer a survey of relevant techniques and to provide insight into how they are used in practice. The preceding chapters give an overview of the problem of SAS motion estimation and compensation. A deliberate attempt is made to highlight those details which are often less than obvious, but absolutely necessary for success. Such knowledge is usually hard-won through trial and error because it is largely absent in the literature.

In many SAS applications, motion estimation and compensation are just part of the overall processing flow. The ultimate image quality is also determined by any autofocus applied as well as any processing for image enhancement. The purpose of autofocus is to correct any residual defocus left after motion compensation. However, SAS autofocus schemes are generally computationally expensive compared to motion estimation and compensation. They may also have the unintended effect of worsening image quality, as autofocus schemes usually assume a point scatterer for the purpose of estimating the corrupting phase function. It is therefore possible to see a well-focused extended target wrongly compressed to a point. For this reason, it is wise to emphasize the motion compensation/estimation in order to reduce dependence on autofocus.

Nevertheless, autofocus is sometimes a necessary aspect of SAS processing. When imaging at long range, the synthetic aperture length grows making it increasingly difficult to keep the integrated error in the motion estimate within tolerable levels. Also, long range imaging carries the risk of having the medium properties, most notably the sound speed, change as a function of position and time. Such variation can cause image blurring and

is difficult to combat except through the use of autofocus algorithms. There is a large body of literature on the subject of SAR and SAS autofocus. While there are a number of techniques used for SAS autofocus, the most successful appear to trace their lineage to the phase gradient algorithm originally developed for spotlight-mode SAR [113, 114] and subsequently adapted to stripmap mode [115].

The best motion compensation solution for any given SAS depends on factors such as the resolution, available computing power, system cost, vehicle size, and expected operating conditions. The analysis and techniques presented here can be used to cover many possible combinations of these parameters. However, there remain a number of outstanding problems. An important example is widebeam motion compensation. This is a non-trivial problem, and its relevance is increasing as the applications for SAS move beyond high-resolution imaging and into areas like buried object detection and sediment characterization. Another avenue for research is in the area of using the SAS redundant phase center motion estimates to aid the navigation of the host vehicle. Some initial work has been done [116], but the potential for improvement is as yet unclear.



## REFERENCES

- [1] J. P. Fish and H. A. Carr, *Sound Underwater Images*. Lower Cape Publishing, 1990.
- [2] J. P. Fish and H. A. Carr, *Sound Reflections: Advanced Applications of Side Scan Sonar*. Lower Cape Publishing, 2001.
- [3] C. S. Clay and H. Medwin, *Acoustical Oceanography: Principles and Applications*. John Wiley and Sons, Inc., 1977.
- [4] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, 1984.
- [5] M. I. Skolnik, *Radar Handbook*. McGraw-Hill, Inc., 1990.
- [6] B. D. Steinberg and H. M. Subbaram, *Microwave Imaging Techniques*. John Wiley and Sons, Inc., 1991.
- [7] C. V. Jakowatz, 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, 1996.
- [8] D. W. Hawkins, *Synthetic Aperture Imaging Algorithms: with application to wide bandwidth sonar*. PhD thesis, Department of Electrical and Electronic Engineering, University of Canterbury, Christchurch, New Zealand, 1996.
- [9] K. D. Le Page and H. Schmidt, "Bistatic synthetic aperture imaging of proud and buried targets from an AUV," *IEEE Journal of Oceanic Engineering*, July 2002.
- [10] S. K. Mitchell, K. N. Scarborough, S. P. Pitt, and T. S. Kooij, "Bistatic SAS imaging studies," in *Proceedings of the Institute of Acoustics*, vol. 28, Pt. 5, pp. 57–64, 2006.
- [11] S. K. Mitchell, K. N. Scarborough, S. P. Pitt, and T. S. Kooij, "High resolution circular SAS with controlled focus," in *Proceedings of the Institute of Acoustics*, vol. 28, Pt. 5, pp. 65–71, 2006.
- [12] M. Soumekh, *Synthetic Aperture Radar Signal Processing with MATLAB Algorithms*. John Wiley and Sons, Inc., 1999.
- [13] K. D. Rolt and H. Schmidt, "Azimuthal ambiguities in synthetic aperture sonar and synthetic aperture radar imagery," *IEEE Journal of Oceanic Engineering*, vol. 17, pp. 73–79, January 1992.
- [14] J. C. Curlander and R. N. McDonough, *Synthetic Aperture Radar Systems and Signal Processing*. John Wiley and Sons, Inc., 1991.

- [15] P. J. Barclay, *Interferometric Synthetic Aperture Sonar Design and Performance*. PhD thesis, University of Canterbury, Christchurch, New Zealand, 2006.
- [16] G. A. Gilmour, "Synthetic aperture side-looking sonar system, review of U.S. Patent 4,088,978," *Journal of the Acoustical Society of America*, vol. 65, no. 2, 1978.
- [17] W. W. Bonifant, "Interferometric synthetic aperture sonar processing," Master's thesis, Georgia Institute of Technology, 1999.
- [18] D. C. Brown, D. A. Cook, and J. E. Fernandez, "Results from a small synthetic aperture sonar," in *MTS/IEEE Oceans 2006 Boston*, 2006.
- [19] R. Fattal, D. Lischinski, and M. Werman, "Gradient domain high dynamic range compression," *ACM Transactions on Graphics*, vol. 21, pp. 249–256, July 2002.
- [20] S. Kuttikkad and R. Chellappa, "Statistical modeling and analysis of high-resolution synthetic aperture radar images," *Statistics and Computing*, vol. 10, pp. 133–145, 2000.
- [21] D. E. Kreithen, W. W. Irving, and S. M. Crooks, "Generating correlated gamma random fields with application to synthesis of simulated SAR imagery," in *International Conference on Acoustics, Speech, and Signal Processing, 1991 (ICASSP-91)*, vol. 4, pp. 2601–2604, 1991.
- [22] H. J. Muller, "Characterization of radar clutter by gamma induced distributions," in *IEEE International Geoscience and Remote Sensing Symposium Proceedings, 1998 (IGARSS '98)*, vol. 3, pp. 1216–1218, 1998.
- [23] M. Sekine, S. Ohtani, T. Musha, T. Irabu, E. Kiuchi, T. Hagiwara, and Y. Tomita, "Weibull distributed ground clutter," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-17, pp. 596–598, July 1981.
- [24] W. S. Carrara, *Spotlight Synthetic Aperture Radar Signal Processing Algorithms*. Artech House, 1995.
- [25] M. Soumekh, *Fourier Array Imaging*. Prentice-Hall, Inc., 1994.
- [26] M. Cheney, "A mathematical tutorial on synthetic aperture radar," *SIAM Review*, vol. 43, no. 2, pp. 301–312, 2001.
- [27] J. D. Jackson, *Classical Electrodynamics*. John Wiley and Sons, Inc., second edition ed., 1975.
- [28] A. D. Pierce, *Acoustics: An Introduction to Its Physical Principles and Applications*. Acoustical Society of America, 1994.
- [29] A. W. Naylor and G. R. Sell, *Linear Operator Theory in Engineering and Science*. Springer-Verlag, 1982.

- [30] P. M. Morse and K. U. Ingard, *Theoretical Acoustics*. Princeton University Press, 1968.
- [31] F. Rocca, "Synthetic aperture radar: A new application for wave equation techniques (Report SEP-56)," tech. rep., Stanford University, 1987.
- [32] G. Franceschetti and G. Schirinzi, "A SAR processor based on two-dimensional FFT codes," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 26, pp. 356–366, 1990.
- [33] C. Cafforio, C. Prati, and F. Rocca, "SAR data focussing using seismic migration techniques," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 27, pp. 199–207, 1991.
- [34] R. Bamler, "A comparison of range-doppler and wavenumber domain SAR focusing algorithms," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 30, pp. 706–713, July 1992.
- [35] R. Stolt, "Migration by fourier transform," *Geophysics*, vol. 43, pp. 23–48, 1978.
- [36] J. F. Claerbout, "Imaging the earth's interior (Report SEP-40)," tech. rep., Stanford University, 1984.
- [37] H. J. Callow, *Signal Processing for Synthetic Aperture Sonar Image Enhancement*. PhD thesis, Department of Electrical and Electronic Engineering, University of Canterbury, Christchurch, New Zealand, 2003.
- [38] L. Ulander, H. Hellsten, and G. Stenstrom, "Synthetic aperture radar processing using fast factorised back-projection," in *EUSAR 2000, 3<sup>rd</sup> European Conference on Synthetic Aperture Radar*, (Munich, Germany), pp. 753–756, 2000.
- [39] S. Xiao, D. Munson, Jr., 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*, vol. 1, pp. 3–7, 2000.
- [40] 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*, pp. 529–534, June 2002.
- [41] L. M. H. Ulander, H. Hellsten, and G. Stenstrom, "Synthetic aperture radar processing using fast factorized back-projection," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, pp. 760–776, July 2003.
- [42] P. J. Westervelt, "Parametric acoustic array," *Journal of the Acoustical Society of America*, vol. 35, pp. 535–537, 1963.
- [43] L. Kopp, "Arrays trilogy: adaptive-parametric-synthetic," in *Sensor Array and Multichannel Signal Processing Workshop Proceedings*, pp. 59–67, 2004.

- [44] M. A. Pinto, A. Bellettini, R. Hollett, and A. Tesei, "Real- and synthetic-array signal processing of buried objects," *IEEE Journal of Oceanic Engineering*, vol. 27, pp. 484–494, July 2002.
- [45] C. L. Nesbitt and J. L. Lopes, "Subcritical detection of an elongated target buried under a rippled interface," in *MTS/IEEE Oceans 2004*, pp. 1945–1952, 2004.
- [46] M. Jonsson, J. Pihl, and M. Aklint, "Imaging of buried objects by low frequency SAS," in *Oceans 2005 Europe*, pp. 669–673, 2005.
- [47] M. E. Zakharia, C. Pollet, and E. Rigaud, "Combined parametric synthetic and interferometric sonar for the detection of buried objects," in *Oceans 2005 Europe*, pp. 522–526, 2005.
- [48] W. J. Zehner and R. L. Thompson, "Methods of estimating allowable motion perturbations in side-scan sonar systems," *IEEE Journal of Oceanic Engineering*, vol. 21, pp. 245–255, July 1996.
- [49] R. O. Harger, *Synthetic Aperture Radar Systems*. Academic Press, 1970.
- [50] W. M. Brown, "SAR resolution in the presence of phase errors," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 24, pp. 808–814, November 1988.
- [51] D. Blacknell, A. Freeman, S. Quegan, I. A. Ward, I. P. Finley, C. J. Oliver, R. G. White, and J. W. Wood, "Geometric accuracy in airborne SAR images," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-25, pp. 241–258, March 1989.
- [52] G. Fornaro, "Trajectory deviations in airborne SAR: Analysis and compensation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 35, pp. 997–1009, July 1999.
- [53] G. Fornaro, E. Sansosti, G. Franceschetti, and S. Perna, "Phase accuracy of motion compensated airborne SAR images," in *Proceedings of the IEEE Geoscience and Remote Sensing Symposium (IGARSS '03)*, 2003.
- [54] G. Fornaro, G. Franceschetti, and S. Perna, "Motion compensation errors: Effects on the accuracy of airborne SAR images," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, pp. 1338–1352, October 2005.
- [55] C. E. Cook and M. Bernfeld, *Radar Signals: An Introduction to Theory and Application*. Academic Press, 1967.
- [56] M. Born and E. Wolf, *Principles of Optics: Electromagnetic Theory of Propagation, Interference, and Diffraction of Light*. Pergamon Press, 1975.
- [57] J. D. Lathrop and I. C. Paustian, "The defocusing of long-range synthetic aperture sonar images due to errors in either the assumed speed of sound or the assumed forward speed." Unpublished manuscript, 2006.

- [58] M. A. Richards, "Coherent integration loss due to white gaussian phase noise," *IEEE Signal Processing Letters*, vol. 10, pp. 208–210, July 2003.
- [59] 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*, pp. 236–245, IEEE, 1992.
- [60] M. A. Pinto, F. Fohanno, O. Trémois, and S. Guyonic, "Autofocusing a synthetic aperture sonar using the temporal and spatial coherence of sea-floor reverberation," in *High Frequency Acoustics in Shallow Water* (N. G. P. et al., ed.), pp. 417–424, La Spezia, Italy: SACLANTCEN, 1997.
- [61] A. Bellettini. Personal communication.
- [62] A. Bellettini and M. A. Pinto, "Theoretical accuracy of synthetic aperture sonar micronavigation using a displaced phase-center antenna," *IEEE Journal of Oceanic Engineering*, vol. 27, pp. 780–789, October 2002.
- [63] H. A. Zebker and J. Villasenor, "Decorrelation in interferometric radar echoes," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 30, pp. 950–959, September 1992.
- [64] H. L. Van Trees, *Detection, Estimation, and Modulation Theory, Part I*. Wiley, 1968.
- [65] C. H. Knapp and G. C. Carter, "The generalized correlation method for estimation of time delay," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 24, pp. 320–327, August 1976.
- [66] G. C. Carter, "Time delay estimation for passive sonar signal processing," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 29, pp. 463–470, June 1981.
- [67] A. H. Quazi, "An overview on the time delay estimate in active and passive systems for target localization," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 29, pp. 527–533, June 1981.
- [68] D. H. Johnson and D. E. Dudgeon, *Array Signal Processing*. Prentice-Hall, Inc., 1993.
- [69] R. G. Bartle, *The Elements of Real Analysis*. John Wiley and Sons, Inc., 1964.
- [70] E. McCarthy, "Acoustic characterization of submerged aquatic vegetation," in *High Frequency Acoustics in Shallow Water* (N. G. P. et al., ed.), pp. 417–424, La Spezia, Italy: SACLANTCEN, 1997.
- [71] L. de Witte, "Least squares fitting of a great circle through points on a sphere," *Communications of the ACM*, vol. 3, no. 11, pp. 611–613, 1960.
- [72] J. L. Farrell, *Integrated Aircraft Navigation*. Academic Press, Inc., 1976.

- [73] C. E. Muehe and M. Labitt, "Displaced-phase-center antenna technique," *Lincoln Laboratory Journal*, vol. 12, no. 2, 2000.
- [74] M. A. Richards, *Fundamentals of Radar Signal Processing*. McGraw-Hill, 2005.
- [75] Y. Doisy, "General motion estimation from correlation sonar," *IEEE Journal of Oceanic Engineering*, vol. 23, pp. 127–140, April 1998.
- [76] J. E. Dennis and R. B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. Prentice-Hall, Inc., 1983.
- [77] D. A. Cook, D. C. Brown, and J. E. Fernandez, "Synthetic aperture sonar motion estimation using nonlinear least squares," in *Proceedings of the Institute of Acoustics*, vol. 28, Pt. 5, pp. 176–182, 2006.
- [78] A. D. Matthews, T. C. Montgomery, D. A. Cook, J. W. Oeschger, and J. S. Stroud, "12.75-inch synthetic aperture sonar (SAS), high resolution and automatic target recognition," in *MTS/IEEE Oceans 2006 Boston*, 2006.
- [79] G. S. Sammelmann, J. E. Fernandez, J. T. Christoff, L. Vaizer, J. D. Lathrop, R. W. Sheriff, and T. C. Montgomery, "High frequency/low frequency synthetic aperture sonar," *Naval Research News*, vol. 49, no. 3, pp. 3–8, 1997.
- [80] D. A. Cook, J. S. Stroud, J. E. Fernandez, and J. D. Lathrop, "Results from a hybrid synthetic aperture sonar motion compensation scheme," in *MTS/IEEE Oceans 2005 Europe*, pp. 1376–1381, 2005.
- [81] D. L. Herzing, "Vocalizations and associated underwater behavior of free-ranging atlantic spotted dolphins, *stenella fontalis* and bottlenose, *tursiops truncatus*," *Aquatic Mammals*, vol. 22, no. 2, pp. 61–69, 1996.
- [82] J. Groen, *Adaptive Motion Compensation in Sonar Array Processing*. PhD thesis, Technische Universiteit Delft, 2006.
- [83] A. Bellettini and M. A. Pinto, "Experimental investigation of synthetic aperture sonar micronavigation," in *Proceedings of the Fifth ECUA*, pp. 445–450, 2000.
- [84] S. Banks, *Studies in High Resolution Synthetic Aperture Sonar*. PhD thesis, University College London, 2002.
- [85] J. W. Oeschger, "Estimating along-track displacement using redundant phase centers," in *Proceedings of the Institute of Acoustics*, vol. 28, 2006.
- [86] G. Strang, *Linear Algebra and Its Applications*. Thompson Learning, Inc., 1988.
- [87] C. D. Meyer, *Matrix Analysis and Applied Linear Algebra*. Society for Industrial and Applied Mathematics, 2000.
- [88] S. D. Conte and C. de Boor, *Elementary Numerical Analysis*. McGraw-Hill, Inc., 1980.

- [89] K. A. Johnson, M. P. Hayes, and P. T. Gough, "A method for estimating the sub-wavelength sway of a sonar towfish," *IEEE Journal of Oceanic Engineering*, vol. 20, pp. 258–267, October 1995.
- [90] P. T. Gough and M. A. Miller, "Displaced ping imaging autofocus," *IEE Proceedings on Radar, Sonar, and Navigation*, vol. 151, pp. 163–170, June 2004.
- [91] J. C. Kirk, Jr., "Motion compensation for synthetic aperture radar," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-11, pp. 338–348, May 1975.
- [92] K. A. C. de Macedo and R. Scheiber, "Precise topography- and aperture-dependent motion compensation for airborne SAR," *IEEE Geoscience and Remote Sensing Letters*, vol. 2, pp. 172–176, April 2005.
- [93] W. Carrara, S. Tummala, and R. Goodman, "Motion compensation for widebeam stripmap SAR," *Proceedings of the SPIE*, vol. 2487, pp. 13–23, 1995.
- [94] R. Goodman, S. Tummala, and W. Carrara, "Issues in ultra-wideband, widebeam SAR image formation," *IEEE International Radar Conference*, pp. 479–485, 1995.
- [95] 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*, vol. 37, no. 23, pp. 1405–1407, 2001.
- [96] M. Soumekh, D. A. Nobles, M. C. Wicks, and G. J. Genello, "Signal processing of wide bandwidth and wide beamwidth P-3 SAR data," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 37, pp. 1122–1141, October 2001.
- [97] H. Singh, J. Adams, D. Mindell, and B. Foley, "Imaging underwater for archaeology," *Journal of Field Archaeology*, vol. 27, no. 3, pp. 319–328, 2000.
- [98] R. D. Ballard, F. T. Hiebert, D. F. Coleman, C. Ward, J. S. Smith, K. Willis, B. Foley, K. Croff, C. Major, and F. Torre, "Iron age shipwrecks in deep water off Ashkelon, Israel," *American Journal of Archaeology*, vol. 105, April 2002.
- [99] R. D. Ballard, F. T. Hiebert, D. F. Coleman, C. Ward, J. S. Smith, K. Willis, B. Foley, K. Croff, C. Major, and F. Torre, "Deepwater archaeology of the Black Sea: The 2000 season at Sinop, Turkey," *American Journal of Archaeology*, vol. 105, April 2002.
- [100] R. Quinn, W. Forsythe, C. Breen, M. Dean, M. Lawrence, and S. Liscoe, "Comparison of the maritime sites and monuments record with side-scan sonar and diver surveys: A case study from Rathlin Island, Ireland," *Geoarchaeology*, vol. 17, no. 5, pp. 441–451, 2002.
- [101] J. P. Fish and H. A. Carr, "Integrated remote sensing of dive sites," in *OCEANS '88: A Partnership of Marine Interests*, pp. 1309–1314, 1988.

- [102] A. M. Friedlander, G. W. Boehlert, M. E. Field, J. E. Mason, J. V. Gardner, and P. Dartnell, "Sidescan-sonar mapping of benthic trawl marks on the shelf and slope off Eureka, California," *Fishery Bulletin*, vol. 97, no. 4, pp. 786–801, 1999.
- [103] V. E. Kostylev, B. J. Todd, G. B. J. Fader, R. Courtney, G. Cameron, and R. Pickrill, "Benthic habitat mapping on the scotian shelf based on multibeam bathymetry, surficial geology, and sea floor photographs," *Marine Ecology Progress Series*, vol. 219, pp. 121–137, 2001.
- [104] D. J. Keith, V. Capone, G. S. Cook, D. A. Carey, D. N. Wiley, and J. P. Fish, "Target detection and mapping of aquatic hazardous waste sites in Massachusetts Bay utilizing sidescan sonar," in *Proceedings of OCEANS '92: Mastering the Oceans Through Technology*, pp. 497–502, 1992.
- [105] H. Momma, M. Watanabe, K. Mitsuzawa, K. Danno, and M. Ida, "Search for the Japanese H-II Rocket Flight No. 8," in *MTS/IEEE Oceans 2000*, pp. 1115–1120, 2000.
- [106] M. K. Kutzleb, "The search for South African Airways Flight 295," in *OCEANS '88: A Partnership of Marine Interests*, pp. 1444–1447, 1988.
- [107] J. W. Will, "SONAR: Underwater search and recovery," *FBI Law Enforcement Bulletin*, pp. 1–5, May 1987.
- [108] D. C. Munson, Jr. and R. L. Visentin, "A signal processing view of strip-map synthetic aperture radar," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 37, no. 12, pp. 2131–2147, 1989.
- [109] J. W. Goodman, *Statistical Optics*. John Wiley and Sons, Inc., 1985.
- [110] R. Mallart and M. Fink, "The van Cittert-Zernike theorem in pulse echo measurements," *Journal of the Acoustical Society of America*, vol. 90, pp. 2718–2727, November 1991.
- [111] M. P. Hayes and P. T. Gough, "Synthetic aperture sonar: A maturing discipline," in *Proceedings of the Seventh European Conference on Underwater Acoustics (ECUA)*, 2004.
- [112] D. A. Cook and D. C. Brown, "Analysis of phase error effects on stripmap SAS." Submitted for review to the *IEEE Journal of Oceanic Engineering*.
- [113] D. E. Wahl, P. H. Eichel, D. C. Ghiglia, and C. V. Jakowatz, Jr., "Phase gradient autofocus— a robust tool for high resolution SAR autofocus," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 30, pp. 827–835, July 1994.
- [114] P. Eichel, D. Ghiglia, and C. Jakowatz, Jr., "Speckle processing method for synthetic-aperture-radar phase correction," *Optics Letters*, vol. 14, October 1989.



- [115] D. E. Wahl, C. V. Jakowatz, Jr., P. A. Thompson, and D. C. Ghiglia, “New approach to strip-map SAR autofocus,” in *Sixth IEEE Digital Signal Processing Workshop*, pp. 53–56, 1994.
- [116] J. D. Campbell and J. R. Pearson, “A synthetic aperture sonar micronavigation Kalman filter,” in *Proceedings of the Institute of Acoustics*, vol. 28, Pt. 5, pp. 183–190, 2006.