


Seafloor Depth Estimation by means of Interferometric Synthetic Aperture Sonar



Torstein Olsmo Sæbø

A dissertation for the degree of
Philosophiae Doctor

September 2010



«If you knew what you were doing
it wouldn't be called research»

Albert Einstein

«Hvem sa at tiden leger alle sår
det er løgn, tiden bare går og går»

Dumdum Boys

Abstract

The topic of this thesis is relative depth estimation using interferometric sidelooking sonar. We give a thorough description of the geometry of interferometric sonar and of time delay estimation techniques. We present a novel solution for the depth estimate using sidelooking sonar, and review the cross-correlation function, the cross-uncertainty function and the phase-differencing technique. We find an elegant solution to co-registration and unwrapping by interpolating the sonar data in ground-range. Two depth estimation techniques are developed: Cross-correlation based sidescan bathymetry and synthetic aperture sonar (SAS) interferometry. We define flank length as a measure of the horizontal resolution in bathymetric maps and find that both sidescan bathymetry and SAS interferometry achieve theoretical resolutions. The vertical precision of our two methods are close to the performance predicted from the measured coherence. We study absolute phase-difference estimation using bandwidth and find a very simple split-bandwidth approach which outperforms a standard 2D phase unwrapper on complicated objects. We also examine advanced filtering of depth maps. Finally, we present pipeline surveying as an example application of interferometric SAS.

Acknowledgments

I would like to thank everyone who has helped and inspired me during my doctoral study.

I especially want to thank my two advisors. Professor Alfred Hanssen at the University of Tromsø have contributed with detailed and precise remarks on fields where I originally lacked the knowledge. Without his help I may never have achieved the appropriate quality of this thesis. Without the optimism of my co-advisor, Doctor Roy Edgar Hansen at the Norwegian Defence Research Establishment (FFI), I believe I would not have been able to finished this thesis. His suggestions for topics and his fast and useful feedback during the writing have been a necessity. Thanks for excellent work.

The University of Tromsø I thank for support, and for patience when I was late with the yearly progress reports. I would like to thank FFI for giving me the opportunity to write this thesis within my work hours.

I also thank Kongsberg Maritime for a very fruitful collaboration. Special thanks goes to Per Espen Hagen, Bjørnar Langli and Terje Gunnar Fossum.

All my lab colleagues at FFI has made it a convivial place to work. In particular, I would like to thank Hayden John Callow for excellent advices and for his friendship and help in the past seven years. I would also like to mention Stig Asle Vaksvik Synnes, Øivind Midtgaard and Herman Midelfart, and all the members of the HUGIN group.

Without my friends, and that includes the above persons, I would have struggled to keep my spirit up through the years I have worked on this thesis. They have patiently listened to all my complaints and distracted me with delightful social events. I also believe that without the recreation in road biking, my thesis work would have been much harder. To the people I shared this interest with: Thank you for the companionship.

Nothing is more important than family. To my wife Cesilie: Special thanks for endless support, for making me smile, and for your love. And to my daughter, Klara. Every day, when I come home tired, you greet me with a contagious smile. You make everything worth it. And to my unborn son: You make our family complete.

Torstein Olsmo Sæbø, Lillestrøm, September 2010.

List of Acronyms

AUV	Autonomous Underwater Vehicle
AWGN	Additive White Gaussian Noise
CCF	Cross-Correlation Function
CRLB	Cramér-Rao Lower Bound
CUF	Cross-Uncertainty Function
DEM	Digital Elevation Map
DPCA	Displaced Phase Center Antenna
DTM	Digital Terrain Map
FFI	Norwegian Defence Research Establishment
IMU	Inertial Measurement Unit
INS	Integrated Navigation System
InSAR	Interferometric Synthetic Aperture Radar
InSAS	Interferometric Synthetic Aperture Sonar
LFM	Linear-Frequency Modulated
MBE	MultiBeam Echo sounder
MLP	Maximum Likelihood Phase-difference
MTI	Moving Target Indication
PCA	Phase Center Antenna
PDF	Probability Density Function
PGA	Phase Gradient Autofocus
REA	Rapid Environmental Assessment
RMS	Root-Mean-Square
SAR	Synthetic Aperture Radar
SAS	Synthetic Aperture Sonar
SNR	Signal-to-Noise Ratio
STD	STandard Deviation
TDOA	Time Difference Of Arrival
TOA	Time Of Arrival
WB	Weighted Bilateral
WM	Weighted Median
WS	Weighted Smoothing

List of symbols

a	Image amplitude [#]
B	Signal bandwidth [Hz]
c	Sound velocity [m/s]
d	Receiver along-track element size [m]
dr	Spatial sampling frequency in slant-range images [m]
D	Interferometric vertical baseline [m]
D_{crit}	Critical baseline [m]
f_{θ}	Fringe frequency [m^{-1}]
f_1	Signal recorded at interferometric receiver #1 [#]
f_2	Signal recorded at interferometric receiver #2 [#]
j	Imaginary unit [#]
k	Coherence [#]
k_D	Baseline dependent coherence [#]
k_{SNR}	SNR dependent coherence [#]
k_T	Temporal coherence [#]
K	Normalization factor in the normalized cross-correlation function [#]
L	Receiver along-track array size [m]
L_{SA}	Synthetic aperture length [m]
N	Number of independent samples in an estimator [#]
n_1	Noise part of signal recorded at interferometric receiver #1 [#]
n_2	Noise part of signal recorded at interferometric receiver #2 [#]
N_r	Number of elements in each receiver array [#]
P	Length of correlation window in sidescan bathymetry [m]
P_x	Length of coherence window along-track [m]
P_y	Length of coherence window in ground-range [m]
r	Slant-range [m]
R	Rotation matrix [#]
S	Scaling factor [#]
t	Signal two-way travel time [s]
T	Time interval of correlation window [s]
W	Spectral shift [Hz]
x	Along-track coordinate axis [m]
y	Cross-track (ground-range) coordinate axis [m]
z	Vertical coordinate axis [m]
\hat{z}	<i>A priori</i> relative seafloor depth [m]
z_{est}	Estimated relative seafloor depth [m]
z_{true}	True relative seafloor depth [m]

$z_{2\pi}$	Height ambiguity [m]
α	Interferometric dilation-factor [#]
β	Vehicle pitch [rad]
γ	Complex coherence / complex interferogram [#]
Γ	Gamma function [#]
δy	Spatial interferometric separation in ground-range images [m]
δ_r	Slant-range resolution [m]
δ_x	Along-track resolution [m]
δ_y	Ground-range resolution [m]
θ	Interferometric phase-difference [rad]
λ	Wavelength of transmitted sonar signal [m]
ν	Image phase [rad]
ρ	Signal-to-noise ratio [#]
σ_c	Standard deviation of coarse (magnitude-based) time delay estimate [m]
σ_f	Standard deviation of fine (complex-based) time delay estimate [m]
σ_i	Standard deviation of phase-difference based time delay estimate [m]
σ_z	Standard deviation of relative depth estimate [m]
σ_τ	Standard deviation of time delay estimate [s]
σ_θ	Standard deviation of phase-difference estimate [rad]
τ	Interferometric time delay [s]
τ_c	Coarse (magnitude-based) time delay estimate [s]
τ_f	Fine (complex-based) time delay estimate [s]
ϕ	Vehicle roll [rad]
Φ	Depression angle in sonar body frame [rad]
Φ_0	Interferometric array tilt-angle relative to vertical [rad]
ψ	Vehicle yaw [rad]
$f(r)$	Backscatter range function [#]
$p(\cdot)$	Probability density function (PDF) [#]
$I(x, y)$	Sonar intensity image (SAS or sidescan) [dB]
$R(\tau)$	Cross-correlation function [#]
$R[j]$	Discrete cross-correlation function [#]
$ R[j] $	Discrete coherence function [#]
$s(x, y)$	Seafloor reflectivity function [dB]
$\chi(\tau)$	Cross-uncertainty function [#]
$\chi[j]$	Discrete cross-uncertainty function [#]
$\frac{\partial \theta}{\partial z}$	Height sensitivity [rad/m]
$E\{\cdot\}$	Expectation operator
$\mathcal{I}\{\cdot\}$	Interpolation operator
$\mathcal{U}\{\cdot\}$	Unwrap operator
$\mathcal{M}\{\cdot\}$	Median operator

Contents

Abstract	i
Acknowledgments	iii
List of Acronyms	v
List of symbols	viii
Table of Contents	xi
1 Introduction	1
1.1 Motivation	1
1.2 Thesis scope	3
1.3 Thesis contribution	4
1.4 List of publications	5
1.5 Outline	8
2 Interferometric synthetic aperture processing	11
2.1 Imaging	11
2.2 Synthetic aperture processing	13
2.3 Interferometry	15
2.4 Synthetic aperture image statistics	16
2.5 Relation to radar	17
3 Geometry	21
3.1 Geometry in the vertical-plane	21
3.1.1 Interferometric time delay	22
3.1.2 Interferometric time dilation	26
3.1.3 Depth estimation in co-registered ground-range	30
3.1.4 Co-registration in slant-range	34
3.2 Geometry in the horizontal plane	35
3.3 Depth accuracy and baseline limitations	37

4	Time delay estimation	41
4.1	Cross-correlation of signals with relative delay	42
4.1.1	Locating the peak of the cross-correlation function	43
4.1.2	Magnitude-correlation	45
4.1.3	Accuracy of the time delay estimate	45
4.1.4	Center frequency shift correction	49
4.2	Cross-correlation of signals with relative delay and relative dilation	50
4.3	Wideband cross-uncertainty function of signals with relative delay and relative dilation	52
4.3.1	Implementation of a wideband CUF estimator	53
4.4	Numerical study of the time delay accuracy	53
4.5	Phase-differencing	59
4.6	Model errors	66
4.6.1	Uncorrelated additive white noise	66
4.6.2	Baseline decorrelation	66
4.6.3	Multipath	67
4.6.4	Correlated noise	67
4.6.5	Multiplicative noise	67
4.6.6	Dispersive scattering	68
4.6.7	Frequency-dependent noise	68
4.6.8	Phase ambiguities	68
4.7	Coherence estimation	70
5	Algorithms for depth estimation	73
5.1	Sidescan seafloor depth estimation	73
5.1.1	Re-gridding to ground-range	76
5.1.2	Cross-correlation of patches	77
5.1.3	Unwrapping the sidescan bathymetry estimates	78
5.1.4	Converting to relative depth	81
5.1.5	Sound speed correction	82
5.1.6	Exporting sidescan bathymetry digital terrain maps	84
5.2	Use of depth estimates in SAS processing	86
5.2.1	Integration of micronavigation estimates	86
5.2.2	Focus plane for imaging	87
5.3	SAS interferometry	88
5.3.1	Co-registration	89
5.3.2	Estimating the interferogram	89
5.3.3	Unwrapping the interferogram	90
5.3.4	Generating SAS images using SAS bathymetry	90

6	System description	93
6.1	The HISAS 1030 interferometric SAS	93
6.2	The HUGIN 1000-MR AUV	96
6.3	HISAS compared to selected InSAR systems	97
6.3.1	System descriptions	97
6.3.2	Synthetic aperture data collection	100
6.3.3	Interferometric processing	101
6.3.4	Summary	103
7	Resolution and precision assessment	105
7.1	Sidescan bathymetry	106
7.2	SAS bathymetry	113
7.3	Summary	117
8	Results and studies	121
8.1	The phase-difference filter size and shape	121
8.1.1	Summary	127
8.2	Performance of CUF in sidescan bathymetry	127
8.2.1	Summary	131
8.3	Split-bandwidth interferometry	131
8.3.1	Summary	136
8.4	Filtering of the depth maps	136
8.4.1	Maximum likelihood phase-difference filter	140
8.4.2	Weighted smoothing filter	140
8.4.3	Weighted median filter	140
8.4.4	Weighted bilateral filter	141
8.4.5	Experimental results	141
8.4.6	Summary	142
8.5	Comparison of HISAS and EM 3000	142
8.5.1	The EM 3000	146
8.5.2	Large scale comparison between EM 3000 and HISAS	147
8.5.3	Small scale comparison between EM 3000 and HISAS	149
8.5.4	Summary	152
8.6	Using HISAS for pipeline surveying	155
8.6.1	Tracks parallel to the pipeline	159
8.6.2	Track perpendicular to the pipeline	165
8.6.3	Estimating the shape of the pipeline	166
8.6.4	Summary	167
9	Summary and conclusions	169
9.1	Suggested future work	171
	Bibliography	183

Chapter 1

Introduction

Detailed seabed mapping plays an important role in a number of different areas such as offshore exploration, environmental surveillance and military applications. This thesis presents methods for relative depth estimation using interferometric sidelooking sonar. *Low resolution* depth maps (with a few meters horizontal resolution) can be generated fast and robust by cross-correlating real aperture (sidescan) sonar images. *High resolution* maps can be generated using interferometric synthetic aperture processing (InSAS). We present methods, limitations and performance of both types, along with results from sample applications.

1.1 Motivation

There are three important quantities in high resolution seabed mapping: Vertical accuracy, horizontal resolution and area coverage rate. Today, multibeam echo sounders (MBEs) are the most common sensor (de Moustier et al. (1990); Lurton (2002, pages 272-275)). A multibeam echo sounder can not achieve high area coverage rate and high horizontal resolution simultaneously, since both the range and the resolutions are proportional to range (de Moustier et al., 1990). Another approach is to use an interferometric sidelooking sonar (Denbigh, 1989; Bird and Kraeutner, 2001; Denbigh, 1994). A real aperture interferometric sidescan (swath bathymeter) has long range and therefore high area coverage rate (Sæbø and Langli, 2010). However, its horizontal along-track resolution is also limited, unless at very short range and very high frequency.

Synthetic aperture sonar (SAS) processing (Cutrona, 1975; Gough and Hawkins, 1997; Pinto, 2002; Hayes and Gough, 2009) produces sonar images which ideally are range- and frequency-independent. The resolutions are limited by bandwidth in range and sonar element size along-track, and can be as low as a few centimeters in both dimensions. Although SAS is very similar to synthetic aperture radar (SAR), it is only recently that commercial SAS systems have become available. Interferometric SAS (Griffiths et al., 1997; Bonifant Jr et al., 2000) can potentially produce depth maps with close to image resolution, but is still an active research field. Most results presented in the




Figure 1.1: SAS image (left panel) and SAS bathymetry (right panel) of a German WWII plane Focke Wulf 190 A-3. The plane was found by the Royal Norwegian Navy Mine warfare flotilla at 98 m water depth during a search mission with the HISAS 1030 interferometric SAS.

literature are from simulations or from experimental systems.

We believe that SAS systems will replace traditional sidescan sonars in applications where resolution is important. Such applications include offshore exploration, naval mine hunting and archaeological surveying. Furthermore, interferometry will be important on SAS systems for two reasons: It is essential for successful SAS imaging in non-flat terrain (Jakowatz et al., 1996, pages 176-186) and it provides important depth information which is not accessible from the SAS images alone (see Figure 1.1).

In order to form a well focused SAS image in non-straight synthetic apertures, the seafloor depth must be known (Jakowatz et al., 1996, pages 176-186). We believe that sidescan bathymetry is a good approach to obtain this information. While most sidescan bathymeters are narrow-band, and has more than two vertically arrays to resolve phase ambiguities, we argue that two vertical arrays are sufficient as long as we have large bandwidth. The phase ambiguities of a large baseline, large bandwidth sonar can be unwrapped using cross-correlations (Sæbø et al. (2007a); Lurton (2002, page 267)) or multi-band techniques (Lanari et al., 1996).

Since most interferometric SAS systems are at an experimental level there is a lack of well-documented systems, which have been tested on large amount of experimental data in varying conditions. We believe that most of the InSAS principles can be adapted from interferometric SAR (Bamler and Hartl, 1998; Hanssen, 2001). However, there are different challenges which will be revealed in the transition of the technologies. A thorough investigation of all aspects of interferometric SAS will advance the state-of-the-art of underwater depth estimation and will open for new applications.



Figure 1.2: Picture of the HISAS 1030 arrays (left panel) and the HUGIN 1000-MR with the HISAS (right panel).

1.2 Thesis scope

The main scope of this thesis has been to develop methods and algorithms for measuring relative depth using the HISAS 1030 interferometric sidelooking sonar; (see Figure 1.2). We have divided this subject into two parts: Real aperture sonar and synthetic aperture sonar.

For depth estimation using real aperture sonar (swath bathymetry) the literature is scarce, so we build upon methods from other research fields, like time delay estimation. The goal has been to develop, implement and verify a method which is fast, robust and accurate enough for mapping of large areas (i.e. a few meters horizontal resolution). The depth estimation performance when it comes to horizontal resolution, vertical precision and sustainability should be established through theory, and verified by simulations and real-data measurements. The scope has *not been* to develop a method which is optimum from a stand point of information theory.

For depth estimation using synthetic aperture sonar (interferometry) the scope has been to learn from the interferometric synthetic aperture radar literature, and transfer existing algorithms. This means adapting the algorithms to a new environment, e.g. a different phase velocity and different geometry. Again, the methods have to be robust and applicable on large amounts of data. The scope of this thesis also include advanced concepts for filtering of depth estimates, resolving phase ambiguities using bandwidth, and investigation of possible applications like pipeline surveying. However, these topics will be presented at a concept level, and will not be subject to the same extensive testing as the main interferometric method.

Another important part of this thesis has been to implement the methods on a specific sensor platform; in this case the HUGIN 1000-MR autonomous underwater vehicle (AUV) showed in Figure 1.2(b). This includes solving interference issues, importing navigation and environmental data, and calibrating lever-arms. However, as long as the geometry is suitable, the methods in this thesis should be applicable to almost any

interferometric sidelooking sensor on almost any platform. An important topic is a comparison of the performance of our depth estimation algorithms with the performance of an EM 3000 mounted on the same AUV. An EM 3000 is an advanced multibeam echo sounder and represents in this case an excellent validation sensor.

In summary, this thesis advances the state-of-the-art of underwater depth estimation by describing and verifying methods for swath bathymetry and interferometry using AUV based interferometric synthetic aperture sonar.

1.3 Thesis contribution

The contributions of this thesis to the research area of relative depth estimation using sidelooking sonar are the following (in order of appearance):

- The design and evaluation of the HISAS sensor and the development of the FOCUS software package for SAS processing (Hansen et al., 2010a). Much of the work during this period has contributed to the development of HISAS, which has become a commercially available product from Kongsberg Maritime (Fossum et al., 2008). This sensor is bundled with the FOCUS processing software, which among other things contains all methods described in this thesis. HISAS with FOCUS is now sold to a number of international costumers, including two Navies.
- A thorough mathematical description of the geometry in relative seafloor depth estimation. Although equivalent methods exist, our novel description provides an approximation-free and alternative solution. This description was first presented in Sæbø et al. (2007b) and appears here in Section 3.1.3.
- The refined estimate of the time delay from time series with asymmetrical spectra. This method was first presented in Sæbø et al. (2007b) and appears here in Section 4.1.4.
- The use of the cross-uncertainty function as a time- and dilation-estimator in swath bathymetry. This work was first presented in Sæbø et al. (2007a) and appears here in Section 4.3, 4.4 and 8.2.
- The use of a coherence weighted bilateral filter to smooth interferometric depth maps. This method was first presented in Sæbø et al. (2009) and appears here in Section 8.4.
- The in-depth comparison of a swath bathymeter with a multibeam echo sounder, on data collected simultaneously using the same sensor platform. This study was first presented in Sæbø and Langli (2010) and appears here in Section 8.5.

1.4 List of publications

During the writing of this thesis I have been author on three peer-review journal articles:

- T. O. Sæbø, R. E. Hansen and A. Hanssen, “Relative height estimation by cross-correlating ground-range synthetic aperture sonar images,” *IEEE J. Oceanic Engineering*, vol. 32, pp. 971-982, October 2007.
- J. Groen, R. E. Hansen, H. J. Callow, J. C. Sabel and T. O. Sæbø, “Shadow enhancement in synthetic aperture sonar using fixed focusing,” *IEEE J. Oceanic Engineering*, vol. 34, pp. 269-284, July 2009.
- T. O. Sæbø, R. E. Hansen and A. Hanssen, “Seafloor depth estimation by interferometric sonar using the wideband cross-uncertainty function,” to be submitted to *IET Radar, sonar & navigation*, 2010.

I have also been first author on ten articles published in conference proceedings:

- T. O. Sæbø and R. E. Hansen, “A comparison between interferometric SAS and interferometric SAR,” in *Proceedings of the Institute of Acoustics*, Lerici, Italy, September 2010.
- T. O. Sæbø and B. Langli, “Comparison of EM 3000 multibeam echo sounder and HISAS 1030 interferometric synthetic aperture sonar for seafloor mapping,” in *Proceedings of the Tenth European Conference on Underwater Acoustics*, Istanbul, Turkey, pp. 451-461, July 2010.
- T. O. Sæbø, H. J. Callow and P. E. Hagen, “Pipeline inspection with synthetic aperture sonar,” in *Proceedings of the 33th Scandinavian Symposium on Physical Acoustics*, Geilo, Norway, February 2010. Online (<http://www.iet.ntnu.no/en/groups/akustikk/meetings/SSPA/2010>).
- T. O. Sæbø, R. E. Hansen and Ø. Midtgaard, “Filtering of high resolution interferometric synthetic aperture sonar,” in *Proceedings of Underwater Acoustic Measurements 2009*, Nafplion, Greece, June 2009. CDROM (ISBN 978-960-98883-4-9).
- T. O. Sæbø, B. Langli, H. J. Callow, E. O. Hammerstad and R. E. Hansen, “Bathymetric capabilities of the HISAS interferometric synthetic aperture sonar,” in *Proceedings of OCEANS 2007 MTS/IEEE*, Vancouver, Canada, October 2007.
- T. O. Sæbø, R. E. Hansen, H. J. Callow and B. Kjellesvig, “Using the cross-ambiguity function for improving sidelooking sonar height estimation,” in *Proceedings of Underwater Acoustic Measurements 2007*, Crete, Greece, pp. 309-316, June 2007.
- T. O. Sæbø, H. J. Callow and R. E. Hansen, “Synthetic aperture sonar interferometry: Experimental results from SENSOTEK,” in *Proceedings of the Eight European Conference on Underwater Acoustics*, Carvoeiro, Portugal, pp. 829-834, June 2006.

- T. O. Sæbø, R. E. Hansen and H. J. Callow, “Height estimation on wideband synthetic aperture sonar: Experimental results from InSAS-2000,” in *Proceedings of OCEANS 2005 MTS/IEEE*, Washington DC, USA, September 2005.
- T. O. Sæbø and R. E. Hansen, “Bathymetric imaging using synthetic aperture sonar,” in *Proceedings of the Seventh European Conference on Underwater Acoustics*, Delft, The Netherlands, pp. 1157-1162, July 2004.
- T. O. Sæbø and R. E. Hansen, “Synthetic aperture sonar interferometry,” in *Proceedings of the 27th Scandinavian Symposium on Physical Acoustics*, Ustaoset, Norway, January 2004. CDROM (ISBN 82-8123-000-2).

In addition, I have contributed to 23 other conference articles related to the work presented in this thesis:

- Ø. Hegrenæs, T. O. Sæbø, P. E. Hagen and B. Jalving, “Horizontal mapping accuracy in hydrographic AUV surveys,” in *Proceedings of the IEEE Autonomous Underwater Vehicles 2010*, Monterey, USA, September 2010.
- H. J. Callow, R. E. Hansen, T. O. Sæbø and S. A. V. Synnes, “Autofocus for circular synthetic aperture imaging,” in *Proceedings of the Institute of Acoustics*, Lerici, Italy, September 2010.
- S. A. V. Synnes, H. J. Callow, R. E. Hansen and T. O. Sæbø, “Multipass coherent processing on synthetic aperture sonar data,” in *Proceedings of the Tenth European Conference on Underwater Acoustics*, Istanbul, Turkey, pp. 1307-1314, July 2010.
- R. E. Hansen, H. J. Callow, T. O. Sæbø and S. A. V. Synnes, “Challenges in sea-floor imaging and mapping with synthetic aperture sonar,” in *Proceedings of the Eight European conference on synthetic aperture radar*, Aachen, Germany, June 2010. CDROM (ISBN 978-3-8007-3272-2).
- R. E. Hansen, T. O. Sæbø, H. J. Callow and P. E. Hagen, “Interferometric synthetic aperture sonar in pipeline inspection,” in *Proceedings of OCEANS 2010 MTS/IEEE*, Sydney, Australia, May 2010. CDROM (ISBN 978-1-4244-5222-4).
- Ø. Midtgaard, T. O. Sæbø and R. E. Hansen, “Estimation of target detection/ classification performance using interferometric sonar coherence,” in *Proceedings of Underwater Acoustic Measurements 2009*, Nafplion, Greece, June 2009. CDROM (ISBN 978-960-98883-4-9).
- S. A. Synnes, R. E. Hansen and T. O. Sæbø, “Assessment of shallow water performance using interferometric sonar coherence,” in *Proceedings of Underwater Acoustic Measurements 2009*, Nafplion, Greece, June 2009. CDROM (ISBN 978-960-98883-4-9).

- R. E. Hansen, H. J. Callow, T. O. Sæbø, S. A. Synnes, P. E. Hagen, T. G. Fossum and B. Langli, “Synthetic aperture sonar in challenging environments: Results from the HISAS 1030,” in *Proceedings of Underwater Acoustic Measurements 2009*, Nafplion, Greece, June 2009. CDROM (ISBN 978-960-98883-4-9).
- H. J. Callow, R. E. Hansen, S. A. Synnes and T. O. Sæbø, “Circular synthetic aperture sonar without a beacon,” in *Proceedings of Underwater Acoustic Measurements 2009*, Nafplion, Greece, June 2009. CDROM (ISBN 978-960-98883-4-9).
- R. E. Hansen, H. J. Callow, T. O. Sæbø, P. E. Hagen and B. Langli, “High fidelity synthetic aperture sonar products for target analysis,” in *Proceedings of OCEANS 2008 MTS/IEEE*, Quebec, Canada, September 2008.
- T. G. Fossum, T. O. Sæbø, B. Langli, H. J. Callow, and R. E. Hansen, “HISAS 1030 - high resolution interferometric synthetic aperture sonar,” in *Proceedings of the Canadian Hydrographic Conference and National Surveyors Conference*, Victoria, BC, Canada, May 2008.
- Ø. Midtgaard, T. O. Sæbø and H. J. Callow, “Automated recognition of short-tethered objects with interferometric synthetic aperture sonar,” in *Proceedings of OCEANS 2007 MTS/IEEE*, Vancouver, Canada, October 2007.
- R. E. Hansen, H. J. Callow and T. O. Sæbø, “The effect of sound velocity variations on synthetic aperture sonar,” in *Proceedings of Underwater Acoustic Measurements 2007*, Crete, Greece, pp. 323-330, June 2007.
- H. J. Callow, R. E. Hansen and T. O. Sæbø, “Effect of approximations in fast factorized backprojection in synthetic aperture imaging of spot regions,” in *Proceedings of OCEANS 2006 MTS/IEEE*, Boston, MA, USA, September 2006.
- R. E. Hansen, H. J. Callow, T. O. Sæbø, P. E. Hagen and B. Langli, “Signal processing for the SENSOTEK interferometric SAS: Lessons learned from HUGIN AUV trials,” in *Proceedings of the Institute of Acoustics*, vol. 28, Lerici, Italy, pp. 47-56, September 2006. CDROM (ISBN 1-901656-79-9).
- R. E. Hansen, T. O. Sæbø and H. J. Callow, “The SENSOTEK interferometric synthetic aperture sonar: Results from HUGIN AUV trials,” in *Proceedings of the Eighth European Conference on Underwater Acoustics*, Carvoeiro, Portugal, pp. 623-628, June 2006.
- R. E. Hansen, T. O. Sæbø, H. J. Callow and B. Langli, “Synthetic aperture sonar imaging on HUGIN AUV: Results from SENSOTEK SAS,” in *Proceedings of the 29th Scandinavian Symposium on Physical Acoustics*, Ustaoset, Norway, January 2006. CDROM (ISBN 82-8123-001-0).

- B. Kjellesvig, R. E. Hansen and T. O. Sæbø, “Improved height estimation for synthetic aperture sonar interferometry,” in *Proceedings of the 29th Scandinavian Symposium on Physical Acoustics*, Ustaoset, Norway, January 2006. CDROM (ISBN 82-8123-001-0).
- R. E. Hansen, T. O. Sæbø, H. J. Callow, P. E. Hagen and E. Hammerstad, “Synthetic aperture sonar processing for the HUGIN AUV,” in *Proceedings of OCEANS 2005 MTS/IEEE*, Brest, France, June 2005.
- H. J. Callow, T. O. Sæbø and R. E. Hansen, “Towards robust quality assessment of SAS imagery using the DPCA algorithm,” in *Proceedings of OCEANS 2005 MTS/IEEE*, Brest, France, June 2005.
- J. Groen, B. A. J. Quesson, J. C. Sabel, R. E. Hansen, H. J. Callow and T. O. Sæbø, “Simulation of high resolution mine hunting sonar measurements,” in *Proceedings of Underwater Acoustic Measurements 2005*, Crete, Greece, June-July 2005.
- R. E. Hansen, T. O. Sæbø and P. E. Hagen, “Development of synthetic aperture sonar for the HUGIN AUV,” in *Proceedings of the Seventh European Conference on Underwater Acoustics*, Delft, The Netherlands, pp. 1139-1144, July 2004.
- R. E. Hansen, T. O. Sæbø, K. Gade and S. Chapman, “Signal processing for AUV based interferometric synthetic aperture sonar,” in *Proceedings of OCEANS 2003 MT/IEEE*, San Diego, CA, USA, pp. 2438-2444, September 2003.

1.5 Outline

The outline of this thesis is as follows. Chapter 2 gives an introduction to interferometric synthetic aperture sonar. The main topics of the chapter is imaging and interferometric processing, but it also covers statistics of synthetic aperture images. We conclude with a comparison of interferometric synthetic aperture radar and interferometric synthetic aperture radar (InSAR). The material presented here is a review of earlier work.

In Chapter 3 we give a thorough description of the geometry in interferometric side-looking sonar. We describe the interferometric time delay and time dilation and present a ground-range co-registration technique which corrects for both effects. We also discuss effects of aperture length and deduce the accuracy of the relative depth estimate.

In Chapter 4 we discuss time delay estimation and time dilation estimation. We introduce the cross-correlation function, the cross-uncertainty function and the phase-difference technique. We present a number of likely model errors, before we conclude with a short section on coherence estimation.

In Chapter 5 we describe our specific implementation of the methods presented in Chapters 3 and 4. We have divided the subject into depth estimation using sidescan sonar, depth estimation using SAS and use of relative depth estimation in SAS processing.

In Chapter 6 we present detailed specifications of the HISAS 1030 sonar and compare HISAS to a number of SAR systems. We also briefly discuss the HUGIN 1000-MR autonomous underwater vehicle.

In Chapter 7 we present a theoretical study of the horizontal resolution and vertical precision of the relative depth estimates. We distinguished between depth estimation using sidescan sonar and depth estimation using SAS. We verify our findings using simulated data.

In Chapter 8 we present results from a number of studies using experimental HISAS data. These topics are comparison of HISAS with a multibeam echo sounder, filtering of interferograms and depth maps, resolving phase ambiguities and using HISAS for pipeline inspection.

The conclusion of this thesis is presented in Chapter 9.

Chapter 2

Interferometric synthetic aperture processing

This chapter provides a summary of the appropriate background knowledge needed to follow subsequent parts of this thesis. For an extensive coverage on the topic see e.g. Hanssen (2001) or Franceschetti and Lanari (1999).

2.1 Imaging

Echo imaging is an inverse problem where the goal is to construct an image which represents the locations of discrete targets, or the spatial distribution of a continuous parameter related to the physical properties of a target (Soumekh, 1994, chapter 1). In sonar imaging, large-bandwidth acoustic pulses are transmitted and pulse compressed on receive. The pulse compression is a matched filtering which maximizes the signal-to-noise ratio (SNR) in white Gaussian noise (DiFranco and Rubin (1980, chapter 5); Levanon (1988, chapter 5); Richards (2005, pages 161-167)). Each target in the illuminated scene produces an echo which after pulse compression appears as a sinc-like pattern in the received time series, dependent on the transmitted pulse form (Franceschetti and Lanari, 1999, pages 15-24). We define slant-range to be the direction of the acoustic waves for a given range. Often, range and slant-range are used interchangeably. The 3 dB *range resolution* is a measure of the range-distance needed to separate two targets of equal strength by 3 dB (Franceschetti and Lanari, 1999, page 23)

$$\delta_r^{\text{ssi}} = \frac{c}{2B}, \quad (2.1)$$

where c is the speed of sound and B the signal bandwidth. To suppress sidelobes at the expense of increased resolution we apply a mild window after the matched filter.

In this thesis we define an imaging coordinate system, where x is the coordinate along the receiver array, r is the slant-range coordinate, z is pointing down towards the seafloor and y is ground-range orthogonal to x and z (the xyz system forms a right




Figure 2.1: Sketch of the imaging geometry. The along-track direction (or equivalently, the receiver array direction) is x , z points down towards the seafloor and y is cross-track.

hand coordinate system, see Figure 2.1). In sonar it is common to have an horizontal array of receiver hydrophones. In this case the received time series can be dynamically focused or beamformed in different directions. A *narrow-band* beamformer uses phase multipliers to phase-delay the time series relative to each other (Van Trees, 2002, page 34). A *broad-band* beamformer, time delays each time series in order to achieve a focused image (Johnson and Dudgeon, 1993, pages 112-119). A common definition of a narrow-band signal is that the signal bandwidth is significantly less than the frequency (Taylor, 1994, page 2). Modern imaging sonars can be broad-band.

The along-track 3 dB image resolution of a physical array with N_r elements of along-track size d is (Van Trees, 2002, pages 46-51)

$$\delta_x^{\text{ssi}} = 0.891 \frac{\lambda}{N_r d} r = 0.891 \frac{\lambda}{L} r \approx \frac{\lambda}{L} r, \quad (2.2)$$

where λ is the wavelength of the transmitted signal, r is the range and $L = N_r d$ is the receiver array-length. Since the frequency is defined as $f = c/\lambda$, the along-track resolution can be increased *either* by increasing the array-length or by increasing the frequency. As very large arrays are impractical, the along-track resolution is usually improved by increasing the signal frequency. However, high-frequency signals suffer

from attenuation (Brekhovskikh and Lysanov (1982, pages 9-11); Lurton (2002, pages 18-26)) and will therefore decay at a shorter range.

A sidelooking imaging sonar which is based on beamforming a physical array is sometimes referred to as a real aperture sonar (in contradiction to a synthetic aperture sonar). When such a sonar is moved along a path and repetitive pulses are transmitted, it is called a sidescan sonar.

2.2 Synthetic aperture processing

A real aperture sonar is limited by a range-dependent along-track resolution. In order to achieve high along-track resolution, one must have very high frequency and short range. This reduces the area coverage rate and makes the sonar impractical for surveying of large areas. A solution adapted from radar is to use synthetic aperture processing (Franceschetti and Lanari, 1999; Jakowatz et al., 1996). In synthetic aperture processing successive pings (or pulses in radar terminology) are coherently combined to synthesize a longer array. The range-resolution is the same as for real aperture imaging

$$\delta_r^{\text{sasi}} = \delta_r^{\text{ssi}} = \frac{c}{2B}, \quad (2.3)$$

However, the along track resolution becomes

$$\delta_x^{\text{sasi}} \approx \frac{\lambda}{2L_{\text{SA}}}r, \quad (2.4)$$

where L_{SA} is the length of the synthetic aperture, and the factor two in the denominator is caused by the motion of the transmitter along the synthetic aperture (Lurton, 2002, page 173). The length of the synthetic aperture is limited by the beamwidth of the receive elements

$$L_{\text{SA}} \leq \frac{\lambda}{d}r. \quad (2.5)$$

This means that the maximum along-track resolution achievable is given by (Curlander and McDonough (1991, page 16-21); Franceschetti and Lanari (1999, pages 24-31))

$$\delta_x^{\text{sasi}} \approx \frac{d}{2}. \quad (2.6)$$

The most important property of Equation 2.6 is that it is range-independent. This is quite counter-intuitive for people trained within classical beamforming; the shorter the element, the better the resolution. Note that the diffraction limit (Jakowatz et al., 1996, page 75) is given by $\lambda/4$, so there is no benefit in having elements smaller than $\lambda/2$. Similar to the pulse compression, we apply a mild window in the synthetic aperture processing to suppress along-track sidelobes in the synthetic aperture images. Figure 2.2 shows an illustration of the difference between real aperture and synthetic aperture imaging.

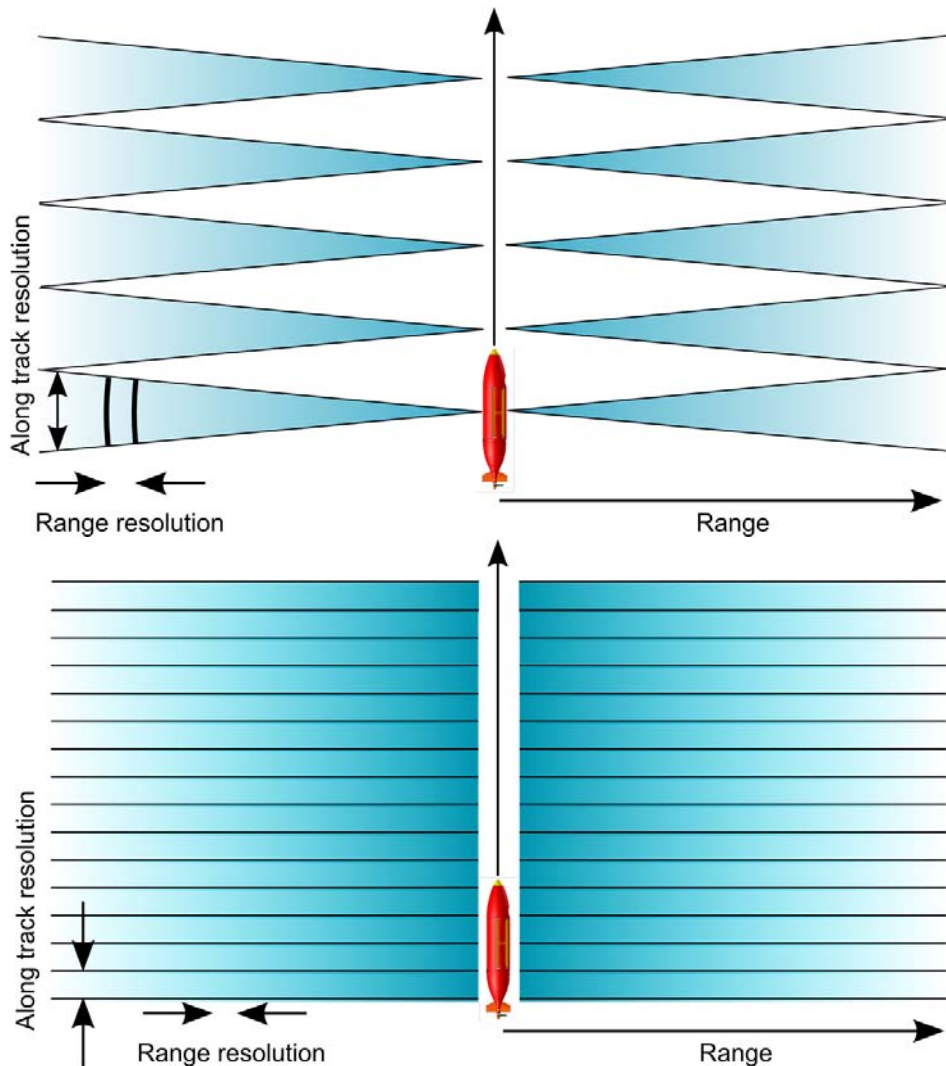


Figure 2.2: Illustration of sidescan sonar imaging geometry (upper panel) and SAS imaging geometry (lower panel).

In synthetic aperture processing one has to move less than half the receiver element size between pings, in order to avoid undersampling lobes (Franceschetti and Lanari, 1999, pages 28-29). Since the phase velocity is a factor of $2 \cdot 10^5$ lower for acoustic waves in seawater than for electromagnetic waves, this imposes an impractical limitation. It is therefore common to use a large number of elements along-track to increase the area coverage rate (Cutrona, 1975; Bruce, 1992).

Another serious constraint is the need for accurate navigation. Navigation errors larger than a fraction of a wavelength over the synthetic aperture will cause defocus in the synthetic aperture images (Jakowatz et al., 1996, pages 228-238). Since the length of the synthetic aperture increases with range, the navigation constraint becomes range dependent. Thus the image quality is often range dependent even if the theoretical image




Figure 2.3: The principle of interferometry is based on estimating the time difference of arrival of two vertically displaced receivers.

resolution is not. On small platforms such as AUVs, inertial navigation systems alone can not provide the desired navigational accuracy, so micrornavigation techniques which use redundancy in the data to estimate sensor translation has been developed. One of the most common methods is the displaced phase center antenna principle (DPCA) which uses cross-correlations on element data (Bellettini and Pinto, 2002).

Another approach adapted from SAR is autofocus, which is a method for blind correction of image degradations using the complex synthetic aperture image as input. The most common technique both in SAR and SAS is called phase gradient autofocus (PGA) (Jakowatz et al. (1996, pages 251-269); Carrara et al. (1995, pages 264-268); Callow (2003)).

2.3 Interferometry

Interferometry means to determine the angular direction of an arrival signal, by means of the time delay between the arrival of the signal at spatially separated receivers (Hanssen, 2001; Franceschetti and Lanari, 1999). Figure 2.3 shows a simple sketch of a typical interferometric sonar. A single transmitter and two vertically separated receivers are used to determine the depression angle of the arriving echo.

The distance between the interferometric receivers is called the *baseline*. Usually, one assumes that the baseline, D is small relative to the range so the arrival wavefronts can be considered parallel (Hanssen, 2001, page 36). The relative depth, z is then found from

$$z = r \left(\frac{c\tau}{D} \right). \quad (2.7)$$

where τ is the interferometric time delay between the arrival signals. The time delay is usually estimated from the phase-difference between the signals (Hanssen, 2001, page

15). The precision of the time delay estimate is a function of SNR, and the estimate can thus be very precise for high SNR. However, the phase-difference is ambiguous modulo 2π (Ghiglia and Pritt, 1998). A number of different approaches have been made to *unwrap* the phase. 2D phase unwrappers find the most likely phase assuming that the data are continuous (Ghiglia and Pritt, 1998), multi-receiver or multi-frequency systems use redundancy to resolve the ambiguities (Hanssen, 2001, pages 72-79) and cross-correlation based methods estimates the ambiguities at the expense of poorer horizontal resolution and increased processing time (Lurton, 2002, page 267).

The accuracy of the time delay estimate is proportional to the baseline (Hanssen, 2001, pages 35-38). However, increasing the baseline to much will reduce the coherence between the signals (Lurton, 2000) and also deteriorate the accuracy of the time delay estimate. Other limiting factors are

- Layover (Franceschetti and Lanari, 1999, page 37-41). In layover regions, there is a mixture of signals arriving from different directions. The different directions cannot be resolved and the coherence drops.
- Shadow (Franceschetti and Lanari, 1999, page 37-41). In shadow regions, there is a lack of signal energy and a time delay can not be estimated.
- Multipath (Brekhovskikh and Lysanov, 1982, chapter 9). Signals arriving from other directions than directly from the seafloor (e.g. via the sea surface or from an elevated object and via the seafloor) will deteriorate the time delay estimate.

In benign bathymetries, the interferometric performance is limited by baseline decorrelation at close range and SNR at long range (Lurton, 2000). In area with large bathymetric variations or with large man-made objects, layover, shadow and multipath will limit the interferometric performance.

2.4 Synthetic aperture image statistics

When the image resolution is significantly larger than the wavelength of the transmitted signal, many scatterers will contribute to the response for each resolution cell (Hanssen, 2001; Oliver and Quegan, 2004, page 89-91). It is then not possible to determine the response of individual scatterers within a resolution cell. The result is the characteristic speckle response (Hanssen, 2001, page 89-91). A common method is to model the measured reflection as a sum of many scatterers. Applying the central limit theorem, the observations can be considered complex (circular) Gaussian random variables. The following assumptions are made (Hanssen, 2001, page 89-91)

- No single scatterer should dominate the others in a resolution cell
- The phase of the individual scatterers must be uniformly distributed between $-\pi$ and π . This holds since a very large phase ($r \gg \lambda$) is wrapped into the $-\pi$ to π interval.

- The phases of the individual scatterers must be uncorrelated.
- The amplitude must be independent from the phase for every scatterer. This holds because the phase is a function of propagation length and is independent of the scattering mechanism

The joint probability density function (PDF) for the image amplitude, a and the image phase, ν can be written (Hanssen, 2001, page 89-91)

$$p(a, \nu) = \begin{cases} \frac{a}{2\pi\sigma^2} \exp\left(-\frac{a^2}{2\sigma^2}\right) & : \text{ for } a \geq 0 \text{ and } -\pi \leq \nu < \pi \\ 0 & : \text{ otherwise,} \end{cases} \quad (2.8)$$

where $\sigma = \sqrt{\mathbb{E}\{I\}}/2$. Here I is the intensity of the resolution cell and $\mathbb{E}\{I\}$ its expectation value. The marginal PDF of a is obtained by integrating ν from $-\pi$ to π .

$$p(a) = \begin{cases} \frac{a}{\sigma^2} \exp\left(-\frac{a^2}{2\sigma^2}\right) & : \text{ for } a \geq 0 \\ 0 & : \text{ otherwise.} \end{cases} \quad (2.9)$$

Equation 2.9 is the Rayleigh distribution (Hanssen, 2001, page 89-91). The marginal PDF of ν is found by integrating a from 0 to ∞

$$p(\nu) = \begin{cases} \frac{1}{2\pi} & : \text{ for } -\pi \leq \nu < \pi \\ 0 & : \text{ otherwise.} \end{cases} \quad (2.10)$$

Equation 2.10 describes a uniform distribution. Since $p(a, \nu) = p(a)p(\nu)$, a and ν are uncorrelated. The pixel intensity variation using the above model is known as speckle. The effect of speckle is often reduced by assuming ergodicity and averaging pixels incoherently. The resulting PDF of the intensity of N averaged pixels is the χ^2 -distribution with $2N$ degrees of freedom (Hanssen, 2001, page 89-91)

$$p(I; N) = \frac{I^{2N-1}}{\mathbb{E}\{I\}^N \Gamma(N)} \exp\left(-\frac{I}{\mathbb{E}\{I\}}\right), \quad (2.11)$$

where Γ is the Gamma function (Rottmann, 1995). For $N = 1$, Equation 2.11 reduces to the exponential PDF. For $N \rightarrow \infty$, it equals a Gaussian PDF.

2.5 Relation to radar

The principle of synthetic aperture radar and synthetic aperture sonar is the same, but there are fundamental differences (Hansen et al., 2010a):

- For electromagnetic signals in air, the phase velocity is typically 3×10^8 m/s. For acoustic waves in seawater, $c \approx 1.5 \times 10^3$ m/s, which limits the forward velocity in SAS. In practice, it is difficult to make a stable SAS-platform with a low enough velocity. The solution is to use multi-element receiver arrays (Cutrona, 1975; Bruce, 1992).

- The atmospheric attenuation of electromagnetic signals depends on the weather conditions, but is often considered a minor effect in SAR. In SAS, however, the seawater absorbs the acoustical signal energy through viscosity and chemical processes (Brekhovskikh and Lysanov (1982, pages 9-11); Lurton (2002, pages 18-26)). This limits the range for a given frequency, as the practical range is roughly constant measured in wavelengths.
- The phase velocity has to be known along the wave path. In SAR the speed of light is accurately known, but in SAS the speed of sound varies with depth (Brekhovskikh and Lysanov (1982, pages 2-9); Lurton (2002, pages [pages 39-41])). In coastal waters, there are also local horizontal and temporal variations. The variation may be as high as 2% along the wave path. The effect is two-fold: An error in the average sound speed leads to defocusing of the SAS images (Rolt and Schmidt, 1994; Hansen et al., 2007), while an error in the sound speed profile also causes position errors (Hegrenæs et al., 2010).
- The imaging geometry of existing SAS systems are very similar, with a swath reaching from nadir to roughly ten times the altitude. This geometry is very different from spaceborne SAR systems, which have a much more vertical geometry. The vertical geometry reduces the effect of shadowing, but increase the effect of foreshortening and layover (Franceschetti and Lanari, 1999, pages 37-42). An airborne SAR system usually has an imaging geometry somewhere between a SAS and a spaceborne SAR.
- To make a diffraction limited image, the sensor position has to be known within a fraction of a wavelength over the synthetic aperture. Satellite tracks are deterministic and accurately known within this limit, but on airborne SAR systems and SAS systems (which can not use GPS) the navigation is often a limiting factor.

While SAR, being available for decades, has reached a very high level of maturity, SAS has only recently become commercially available. This is partly due to the differences listed above. SAR interferometry is today very sophisticated, using techniques such as repeat-pass image collections over years and multi-baselines for tomographic (or 3D) imaging. SAS interferometry has been demonstrated successfully at numerous occasions, but has yet to reveal its full potential. It is likely that advanced methods in interferometric SAR will be adapted by the SAS community. Current technology trends in SAR interferometry are:

- Differential and repeat-pass interferometry for deformation monitoring, where multiple images are collected over a large time span (up to years). A major limitation is that the effect of the atmosphere has to be estimated and compensated for.
- Multi-baseline SAR tomography for 3D imaging, e.g. used in forest mapping (to estimate the average height of the trees).

- Single pass multi-platform interferometric SAR for increased baseline and mapping accuracy using several platforms in formation flying.
- Bistatic SAR using one moving antenna and one stationary antenna, or two moving antennas.
- Multi-frequency and ultra wideband SAR for characterization of areas and targets.
- Multi-channel along-track interferometry for moving target indication.

Chapter 3

Geometry

Interferometry can be looked upon as an advanced form of stereometri (Franceschetti and Lanari, 1999, pages 167-171). Geometrical differences in signals recorded at different positions are used to resolve the angle of arrival of the signal. In relative depth estimation using a bistatic sonar, the echo from the seafloor recorded at individual receivers is exploited (Lurton, 2002, pages 266-267). This method is closely related to radar interferometry (Hanssen (2001, pages 51-58); Franceschetti and Lanari (1999, pages 185-195); Jakowatz et al. (1996, pages 303-317)). However, the geometry can be quite different for sidelooking sonar. While a radar system may have a very large depression angle and baseline, a sonar is usually operated close to the seafloor and with a relatively small interferometric baseline (see Section 6.3). Figure 3.1 shows a schematic of the geometry for a sonar system on an autonomous underwater vehicle.

In more detail, the time delay for a location on the seafloor, τ , caused by travel path differences, is converted to a direction of arrival, or a *depression angle*, Φ . By combining this angle with the range to the seafloor we can determine the relative depth of the seafloor, z . A description of the relation between τ and z can be found from the vertical geometry alone, which means that a common set of equations apply for both relative depth estimation using sidescan sonar and interferometric SAS. However, SAS interferometry has some additional effects in the horizontal plane, due to the integration along-track. We will start by describing the common vertical geometry and then discuss the effects of synthetic aperture processing later in this chapter.

3.1 Geometry in the vertical-plane

Figure 3.1 shows an interferometric sonar with two receivers illuminating a seafloor with relative depth, z_0 . Without loss of generality, the coordinate system is chosen such that the transmitter is placed at origin. The upper receiver is at position (y_1, z_1) relative to the transmitter and the lower receiver at (y_2, z_2) . For a given position (y_0, z_0) at the seafloor, each receiver records a backscattered signal $f(r(y_0, z_0))$ where $r(y_0, z_0)$ is the two-way range from the transmitter, to the seafloor and back to the receiver. The re-




Figure 3.1: Schematic of the vertical plane of a sidelooking sonar with a transmitter at origin and two receivers at positions (y_1, z_1) and (y_2, z_2) . The signal is transmitted down to a position (y_0, z_0) at the seafloor and reflected back to the receivers. The z -axis points down toward the seafloor.

ceived signal is a geometrical transformation of a single realization of the underlying *reflectivity function* of the seafloor, $s(y, z)$.

3.1.1 Interferometric time delay

The basis of relative depth estimation and interferometry is an accurate description of the time delay caused by the bistatic configuration (Hanssen (2001, pages 51-58); Franceschetti and Lanari (1999, pages 185-195); Jakowatz et al. (1996, pages 303-317)). From Figure 3.1 we see that the backscattered signal for a reflector at (y, z) arrives at the two receivers with the time delay

$$\begin{aligned} \tau(y, z) &= t_1(y, z) - t_2(y, z) = \frac{1}{c}((r_0(y, z) + r_1(y, z)) - (r_0(y, z) + r_2(y, z))), \\ &= \frac{1}{c}(r_1(y, z) - r_2(y, z)), \end{aligned} \quad (3.1)$$

where $t_1(y, z)$ and $t_2(y, z)$ are the receive times for receiver #1 and #2, respectively, and c is the sound velocity. The time delay, τ , is in the literature sometimes referred to as

the time difference of arrival (TDOA), while $t_1(y, z)$ and $t_2(y, z)$ are the time of arrivals (TOA) (Falsi et al., 2006). The ranges can be written

$$r_0(y, z) = (y^2 + z^2)^{1/2}, \quad (3.2a)$$

$$r_1(y, z) = ((y - y_1)^2 + (z - z_1)^2)^{1/2}, \quad (3.2b)$$

$$r_2(y, z) = ((y - y_2)^2 + (z - z_2)^2)^{1/2}. \quad (3.2c)$$

From Equations 3.1 and 3.2 it is possible to find a solution for z . Notice that only the return paths are different in the time delay calculation. By relocating the origin to the lower receiver and rotating the coordinate system (see Figure 3.2) the ranges can be written as

$$r_1(y', z') = \left((y')^2 + (z' + D)^2 \right)^{1/2}, \quad (3.3a)$$

$$r_2(y', z') = \left((y')^2 + (z')^2 \right)^{1/2}, \quad (3.3b)$$

where D is the baseline between the receivers. The seafloor position (y', z') , relative to the lower receiver in the new rotated coordinate system is given by

$$y' = y \cos \Phi_0 + z \sin \Phi_0, \quad (3.4a)$$

$$z' = z \cos \Phi_0 - y \sin \Phi_0 - \frac{D}{2}, \quad (3.4b)$$

where Φ_0 is the rotation between the (y, z) -frame and the (y', z') -frame, or the tilt-angle of the sonar relative to vertical. For simplicity, we have assumed that the transmitter is centered between the receivers.

From Equations 3.1 and 3.3 we get the time delay on a functional form

$$\tau(y', z') = r_1(y', z') - r_2(y', z') = \tau(y, z) = \tau(r_2) = \frac{r_2}{c} \left(\left(1 + \frac{D(2z' + D)}{r_2^2} \right)^{1/2} - 1 \right). \quad (3.5)$$

The upper panel of Figure 3.3 shows the bistatic time delay as a function of range (in a non-rotated coordinate system). We see that for typical sidescan sonar of 100 kHz center frequency and higher (Lurton, 2002, page 264), the delay is multiple wavelengths for almost all ranges (a wavelength is equivalent to a delay of 0.01 ms for 100 kHz). This means that the phase-differences are wrapped modulus 2π and cannot be used as an estimate for the time delay without some sort of phase unwrapping (see Section 4.6.8). An alternative is to use short-time cross-correlations along the range axis.

Inverting Equation 3.5 gives the relative seafloor depth in the rotated coordinate frame as a function of the measured time delay $\tau = \tau(y', z')$, and the range $r_2 = r_2(y', z')$

$$z' = r_2 \frac{c\tau}{D} \left(\frac{c\tau}{2r_2} + 1 \right) - \frac{D}{2}. \quad (3.6)$$




Figure 3.2: Schematic of the vertical plane of a sidelooking sonar with one receiver at origin and one receiver at a distance D above origin. The coordinate system is rotated so the two receivers both lie along $y' = 0$. The signal is reflected from a position (y'_0, z'_0) at the seafloor.

Usually, the incoming sound paths can be assumed to be parallel with $2r_2 \gg c\tau$, which simplifies the solution to

$$z' \approx r_2 \frac{c\tau}{D} - \frac{D}{2}. \quad (3.7)$$

The effect of this approximation is illustrated in the lower panel of Figure 3.3.

The depth of the seafloor in a non-rotated coordinate system can now be calculated based on simple trigonometry


$$z = \sqrt{r_2^2 - (z')^2} \sin \Phi_0 + \left(z' + \frac{D}{2}\right) \cos \Phi_0 \quad (3.8)$$

Another common way of solving the geometry is to use angles. For a sonar with tilt-angle Φ_0 , the relative depth is given by


$$z = r \sin(\Phi + \Phi_0), \quad (3.9)$$

where the estimated depression angle relative to the tilt-angle Φ , is

$$\Phi \approx \sin^{-1} \left(\frac{c\tau}{D} \right). \quad (3.10)$$



(a) Time delay between the received signals



(b) Seafloor depth calculated from the time delay

Figure 3.3: Interferometric time delay and reconstructed seafloor depth. Panel (a) shows the time delay in milliseconds between signals recorded at the interferometric receivers, as a function of range. Panel (b) shows the corresponding calculated seafloor depth from the time delay, as a function of range. The solid line shows the result using Equation 3.6 and the dashed line shows the result using Equation 3.7. In both panels we have used a sonar with $c = 1500$ m/s, $z = 15$ m, $D = 30$ cm and $\Phi_0 = 22$ degrees.

Here we have assumed that the incoming sound ways are parallel, an approximation which is identical to assuming $2r_2 \gg c\tau$. For $\Phi_0 = 0$, this method gives the same answer as Equation 3.7 (ignoring the $\frac{D}{2}$ translation of the coordinate systems). However, it is more difficult to find an exact answer using this approach.

3.1.2 Interferometric time dilation

In this section we study details of the interferometric time delay and show how the delay can be interpreted as a dilation of the spatial geometry. A sidelooking sonar measures a realization of the underlying seafloor reflectivity function, $s(y, z)$, as a function of time, t . We assume that there is only one depth for each position, i. e. $s(y, z) = s(y)$. By replacing the bistatic sonar with a phase center antenna (PCA) (Bellettini and Pinto, 2002), time is related to the one-way range through the relation $t = 2r/c$. Since time is only a scaled version of range, we use the latter for convenience. We will show that the coordinate transform from y to r can be regarded as a *dilation* of the spatial y -coordinate. Without loss of generality we adopt the rotated coordinate system from the previous section and omit the apostrophe for simplicity.

A signal, $g(u)$, is dilated relative to a reference signal, $h(v)$ if $g(u) = h(\alpha u)$ for some dilation-factor α . Through the use of the chain-rule we see that the dilation-factor is given by

$$\alpha = \frac{du}{dv}. \quad (3.11)$$

At the receivers we measure a realization of the reflectivity function, $s(y)$, as a function of time or range, $r(y)$. From Equation 3.3, we differentiate $r(y)$ with respect to the spatial y -coordinate, which is equivalent to the Jacobian used in a one-parameter transform

$$\frac{dr_1(y)}{dy} = \frac{y}{r_1(y)} = \cos \Phi_1(y) = \alpha_1(y), \quad (3.12a)$$

$$\frac{dr_2(y)}{dy} = \frac{y}{r_2(y)} = \cos \Phi_2(y) = \alpha_2(y), \quad (3.12b)$$

where $\Phi_1(y)$ and $\Phi_2(y)$ are the depression angles for receiver #1 and #2, respectively, and $\alpha_1(y)$ and $\alpha_2(y)$ are the dilation-factors. The dilation-factors are functions of y , which induces a range-dependent dilation, equal to the usual cosine-transform between slant-range and ground-range (Jakowatz et al., 1996, pages 317-320).

We express the two time signals received at the interferometric array as

$$f_1(r_1) = s(\alpha_1(y) \cdot r_1), \quad (3.13a)$$

$$f_2(r_2) = s(\alpha_2(y) \cdot r_2), \quad (3.13b)$$

where $f_1(r_1)$ is the recorded signal at receiver #1 and $f_2(r_2)$ is the recorded signal at receiver #2. The two signals are time dilated or time scaled versions of the seafloor reflectivity function, $s(y)$. For a constant seafloor depth, z_0 , and a fixed baseline, D ,

$\alpha_1 = \alpha_2$ for $y \rightarrow \infty$. From Equation 3.13 we see that there is in fact no time delay between the two signals, only a different dilation of the seafloor reflectivity function. However, when comparing a small subset of two signals dilated differently from an origin outside the selected subset, one can perceive the dilation as a delay and a local dilation. The upper panel of Figure 3.4 shows the dilation of the seafloor reflectivity function for the two receivers. The dilation is almost identical for the two receivers since $z > D$.

In seafloor depth estimation, the relative dilation between $f_1(r)$ and $f_2(r)$ is more important than the dilation of the seafloor reflectivity function, $s(y)$, itself. By using the signal at receiver #2 as a reference signal and differentiating $r_1(y)$ with respect to $r_2(y)$ we get

$$\alpha(y) = \frac{dr_1(y)}{dr_2(y)} = \left(\frac{dr_1(y)}{dy} \right) \left(\frac{dr_2(y)}{dy} \right)^{-1} = \left(\frac{y^2 + z^2}{y^2 + (z + D)^2} \right)^{1/2}, \quad (3.14)$$

where $\alpha(y)$ is the relative time dilation between the received signals. Equation 3.14 shows that the dilation-factor is unity for $D = 0$ or for $y \rightarrow \infty$, which is consistent with our intuition. The lower panel of Figure 3.4 shows the relative dilation between the signals in a non-rotated coordinate frame. Notice the rapid decrease in the dilation-factor for close ranges.

To be able to resolve the seafloor depth at different positions, we divide the received time series into a number of time intervals. The relative dilation around a position y_0 becomes

$$\alpha(y_0 \pm \Delta y) = \left(\frac{(y_0 \pm \Delta y)^2 + z^2}{(y_0 \pm \Delta y)^2 + (z + D)^2} \right)^{1/2} \approx \left(\frac{y_0^2 + z^2}{y_0^2 + (z + D)^2} \right)^{1/2} = \alpha(y_0), \quad (3.15)$$


where the assumption

$$(y_0^2 + z^2) \gg \Delta y(2y_0 + \Delta y) \quad (3.16)$$


gives the approximate form. This assumption is approximately equivalent to $r_2 \gg 2\Delta y$. Hence, the time dilation can only be regarded as constant for ranges that are large compared to the patch length. Figure 3.5 shows the comparison between the approximated dilation and the exact dilation for patches at 30 and 100 meters. The patch length is two meters in both cases. While the approximation is very good at 100 meters, it is significantly less accurate at 30 meters.

We have now established the following properties for our interferometric model:

- The received signals are time dilated versions of the seafloor reflectivity function with $\alpha_1(y)$ and $\alpha_2(y)$ as dilation-factors for receiver #1 and #2, respectively (see Equation 3.12).
- There is a relative dilation given by $\alpha(y)$ between the signals received at receiver #1 and receiver #2 (see Equation 3.14).




(a) Time dilation between the seafloor and the received signals




(b) Relative dilation between the two received signals

Figure 3.4: Dilation of the reflectivity function. Panel (a) shows the time dilation of the seafloor reflectivity function seen by the two receivers as a function of range. Circles show the dilation seen by receiver #1 and squares the dilation seen by receiver #2. Panel (b) shows the *relative* dilation between the two signals as a function of range. In both panels we have used a sonar with $c = 1500$ m/s, $z = 15$ m, $D = 30$ cm and $\Phi_0 = 22$ degrees.



(a) Time dilation and approximate dilation at 30 m



(b) Time dilation and approximate dilation at 100 m

Figure 3.5: Dilation of the reflectivity function. Panel (a) shows the relative time dilation between the signals received at the two receivers for a two meters patch around 30 meters, as a function of range. The solid line shows the exact dilation and the dashed line an approximated dilation constant over the displayed range interval. Panel (b) shows the same for a patch around 100 meters. In both panels we have used a sonar with $z = 15$ m, $D = 15$ cm and $\Phi_0 = 22$ degrees.

- The relative dilation can be approximated as constant over a small patch, i.e., $\alpha(y \pm \Delta y) \approx \alpha(y)$ (see Equation 3.15).
- In effect, the relative dilation causes a time delay, $\tau(y)$, between the two receivers for an echo from the same seafloor location (see Equation 3.5).

3.1.3 Depth estimation in co-registered ground-range

As we saw in the previous section, interferometric signals are both time shifted (Lurton, 2000) and time dilated (Gatelli et al., 1994) relative to each other. In Section 4.2 we present an estimator suited for time delay estimation on such signals. A different approach is to use *a priori* information to eliminate most of the shift and dilation. This method is in many cases desirable because

- Pre-processing the data increases the coherence between the signals
- Standard time delay estimators (like cross-correlation or phase-differencing) can be used directly
- High quality *a priori* data are available from the previous estimate since the seafloor are varying slowly relative to the platform motion
- Fast 1D-interpolation along the range axis is usually sufficient

An effect we have to take into account when compensating for the dilation is a change of the effective frequency of the signals. We consider this effect in detail in Section 5.1.1.

In Figure 3.6 we demonstrate the potential improvement in SNR by beamforming onto an *a priori* bathymetry. The maximum cross-correlation coefficients in the time delay estimation is converted to SNR according to Equation 4.17. The results are from a Monte Carlo simulation of 5000 realizations, with the correlation performed both directly in slant-range and after beamforming onto an estimated seafloor depth (with five meters induced depth-error). Use of this (poor) *a priori* information provides a significant improvement in SNR, as can be seen in Figure 3.6. Another improvement is achieved by using the first estimate to correct for footprint shift, and thereafter correlate again (a simple beamforming). The correction co-registers the data to the nearest integer number of pixels, and provides a significant SNR improvement. Note that the oscillations in the iterated correlations are caused by the discrete grid applied in the co-registration, and may be reduced by increasing the sampling frequency or by using an improved sub-pixel co-registration technique (the SNR will then approach the peaks of the oscillations).

Figure 3.7 illustrates how the data is re-gridded down onto an *a priori* seafloor depth, \hat{z} . This method is a simple co-registration (Jakowatz et al., 1996, pages 288-302) where features are shifted onto the same positions in the two data-sets. If the seafloor is correct (i.e. $\hat{z} = z_0$ in Figure 3.7), $\hat{y} = y_0$ and $\delta y = 0$. When $\hat{z} \neq z_0$, \hat{y} differs from y_0 and the




Figure 3.6: Estimated SNRs for correlations in slant-range (thick and dashed) and ground-range (thick and solid) as a function of range. The thin lines are the corresponding SNRs after co-registering the data to the nearest integer number of pixels.

spatial lag, δy , becomes non-zero. For a given \hat{z} , δy is estimated for each y -position. We define the ground-range positions for the reflector at (y_0, z_0) to be

$$\hat{y}_1 = \hat{y} - \delta y/2, \quad (3.17a)$$

$$\hat{y}_2 = \hat{y} + \delta y/2, \quad (3.17b)$$

where \hat{y}_1 is the ground-range position for receiver #1 and \hat{y}_2 is the position for receiver #2. The iso-range ellipses (with the transmitter and a receiver as foci) describes \hat{y}_1 and \hat{y}_2 as a function of \hat{z}

$$\sqrt{y_0^2 + z_0^2} + \sqrt{(y_0 - y_1)^2 + (z_0 - z_1)^2} = \sqrt{\hat{y}_1^2 + \hat{z}^2} + \sqrt{(\hat{y}_1 - y_1)^2 + (\hat{z} - z_1)^2}, \quad (3.18a)$$

$$\sqrt{y_0^2 + z_0^2} + \sqrt{(y_0 - y_2)^2 + (z_0 - z_2)^2} = \sqrt{\hat{y}_2^2 + \hat{z}^2} + \sqrt{(\hat{y}_2 - y_2)^2 + (\hat{z} - z_2)^2}. \quad (3.18b)$$

These equations can be solved analytically for y_0 and z_0 . The problem can be substantially simplified by applying the PCA approximation (Bellettini and Pinto, 2002). This approximation replaces each transmitter-receiver pair with a common transmitter and receiver midway between them, at the expense of a small geometrical error. Note that the normal criterion for using PCA approximation, $r \ll D^2/\lambda$ (where λ is the wavelength of the transmitted signal) does not apply since we are considering spatial errors, not phase coherence. The error in the approximation (equivalent to approximating the ellipses with circles) is smallest for a sonar with a transmitter centered between the receivers. The PCA error decreases approximately as $1/r^2$, and increases with the error of the *a priori* depth.




Figure 3.7: Schematic of the vertical plane of a sidelooking sonar with a transmitter in the origin and two receivers at positions (y_1, z_1) and (y_2, z_2) . The signal is transmitted down to a position (y_0, z_0) at the seafloor and reflected back to the receivers. The z -axis points down toward the seafloor. The *a priori* depth is \hat{z} , which separates the iso-range lines with δy at the location \hat{y} .

The PCA position for receiver #1 is defined as $(y_{1P}, z_{1P}) = (y_1/2, z_1/2)$ and $(y_{2P}, z_{2P}) = (y_2/2, z_2/2)$ for receiver #2. Equations 3.18 are simplified to

$$(y_0 - y_{1P})^2 + (z_0 - z_{1P})^2 = (\hat{y}_1 - y_{1P})^2 + (\hat{z} - z_{1P})^2, \quad (3.19a)$$

$$(y_0 - y_{2P})^2 + (z_0 - z_{2P})^2 = (\hat{y}_2 - y_{2P})^2 + (\hat{z} - z_{2P})^2. \quad (3.19b)$$

The depth is found by subtracting the two equations and solving for z_0 . This solution is then inserted in any of the equations above and solved for y_0 . In general the solution for the relative depth, z_0 , and position, y_0 , is given by

$$z_0 = \frac{1}{2}D - Cy, \quad (3.20a)$$

$$y_0 = \frac{-F \pm \sqrt{F^2 - 4EG}}{2E}, \quad (3.20b)$$




Figure 3.8: Error ($z_{\text{est}} - z$) of depth estimate as a function of *a priori* depth, \hat{z} , for a reflector located at $z = 15$ meters using Equation 3.20a. In the figure we show the results for a reflector located at $y = 15$ meters (solid line), $y = 50$ meters (dashed line) and $y = 150$ meters (dash-dotted line).

where the constants are given by

$$\begin{aligned}
 C &= \frac{y_{2P} - y_{1P}}{z_{2P} - z_{1P}}, \\
 D &= \frac{A - B}{z_{2P} - z_{1P}}, \\
 E &= 1 + C^2, \\
 F &= 2z_{1P}C - DC - 2y_{1P}, \\
 G &= (D/2)^2 - z_{1P}D - A,
 \end{aligned}$$

where

$$\begin{aligned}
 A &= (\hat{y}_1 - y_{1P})^2 + (\hat{z} - z_{1P})^2 - y_{1P}^2 - z_{1P}^2, \\
 B &= (\hat{y}_2 - y_{2P})^2 + (\hat{z} - z_{2P})^2 - y_{2P}^2 - z_{2P}^2.
 \end{aligned}$$

In Figure 3.8 we show the difference $z_{\text{est}} - z$ (geometrical error) between the estimated depth z_{est} and the true depth z as a function of *a priori* depth for a reflector located at $z = 15$ meters. From the figure we see that for a sonar with $D = 30$ cm and $\Phi_0 = 22$ degrees, the depth-error due to the PCA approximation is less than 1 cm provided that the *a priori* depth is between 0 and 50 meters. This error is for most practical purposes insignificant, but to further reduce the error, one could either solve Equations 3.18 directly (which would give an exact solution) or run the depth estimation iteratively.




Figure 3.9: Illustration of a simple two-step co-registration technique. The data from one of the receivers are interpolated onto an *a priori* seafloor (1) and then interpolated back onto the geometry of the other receiver (2).

3.1.4 Co-registration in slant-range

An alternative co-registration method, extensively used in SAR, is to use image features to eliminate the relative shift and dilation (Fornaro and Franceschetti, 1995). Instead of using an *a priori* depth, common features are identified in the two images and one of the images (the *slave*) is interpolated to match the other (the *master*). Usually, a warping function describing shift, dilation and rotation is calculated from a set of *control points* or distinctive image features (Hong et al. (2006); Jakowatz et al. (1996, Pages 293-298)). Since the images are in slant-range, the phase-differences are directly proportional to the time delay in Equation 3.5, and Equation 3.6 can be used to estimate the relative seafloor depth. The advantage with this method is that *a priori* depth information is unnecessary. The drawback is that the warping function can be difficult to obtain if the SNR in the images is marginal.

By combining the *a priori* co-registration technique from Section 3.1.3 with slant-range geometry, it is easy to implement a slant-range co-registration technique based on *a priori* depth information. One intuitive approach is to interpolate the data from one of the receivers onto an *a priori* seafloor (step 1 in Figure 3.9) and then reverse interpolate the results back to the other receiver (step 2 in Figure 3.9)

$$\text{slave}_{\text{sr}} \xrightarrow[\text{slave geometry}]{\text{sr to gr}} \text{slave}_{\text{gr}} \xrightarrow[\text{master geometry}]{\text{gr to sr}} \text{slave}_{\text{sr}}^{\text{co-registered}}, \quad (3.21)$$

where sr means slant-range and gr ground-range. The geometrical conversions can be described by Equations 3.2. To reduce errors the two interpolations can be merged. However, we have chosen ground-range as our preferred coordinate system for depth estimation and will not pursue this subject further.

3.2 Geometry in the horizontal plane

So far we have discussed the geometry in the range-depth plane. The description is only dependent on the spatial locations of the receivers relative to the transmitter, and the relative depth of the seafloor. In the along-track horizontal plane, depth estimation may also be effected by aperture length.

A beamforming algorithm delays or migrates the received data onto a grid, which flattens the polar properties of the data. In the case of a full length synthetic aperture, the iso-phase curves are independent of along-track position, i.e. the vertical cross-section we considered in Section 3.1 is valid along the y -axis, independent of x . We identify three cases: An infinitesimal aperture (an omni-directional hydrophone), a small aperture (a physical aperture) and an infinite aperture (a synthetic aperture). Figure 3.10 illustrates the iso-phase curves for these three cases. An infinitesimal receiver has cylindrical symmetry with large x -dependence. An infinite aperture, on the other hand, has plane iso-phase curves (pick any point on the aperture and it looks the same). This means that somewhere between these extremities there is a transition. Figure 3.10 shows that for a slice along line A ($x = 0$) all the sonars behave identically, but for a slice along line B ($x \neq 0$) the phases are scaled relative to the broadside direction. This scaling may affect the time delay estimate of a bathymetric multibeam sidescan sonar with parallel lines.

A non-rigorous consideration of the importance of the scaling factor can be found from simple geometry. We start by assuming circular geometry and define a scaling factor $G = y_{\text{of}}/y_{\text{bs}}$, where y_{bs} is the y -coordinate of the iso-phase line at broadside and y_{of} is the y -coordinate of the same iso-phase line offset x_{of} from broadside. In Figure 3.10 this corresponds to where the lines A and B cross one of the iso-phase lines. From the left panel of Figure 3.11 we see that

$$y_{\text{of}} = (y_{\text{bs}}^2 - x_{\text{of}}^2)^{-1/2}, \quad (3.22a)$$

$$G = \left(\frac{y_{\text{bs}}^2 - x_{\text{of}}^2}{y_{\text{bs}}^2} \right)^{-1/2}, \quad (3.22b)$$

where G is bounded between zero and one, and is equal to one for $x_{\text{of}} = 0$. Clearly, the scaling is important for this geometry. For example, $x_{\text{of}}/y_{\text{bs}} = 0.1$ gives a five percent scaling of the time delay estimate, which easily translates to tens of cm in depth-error.

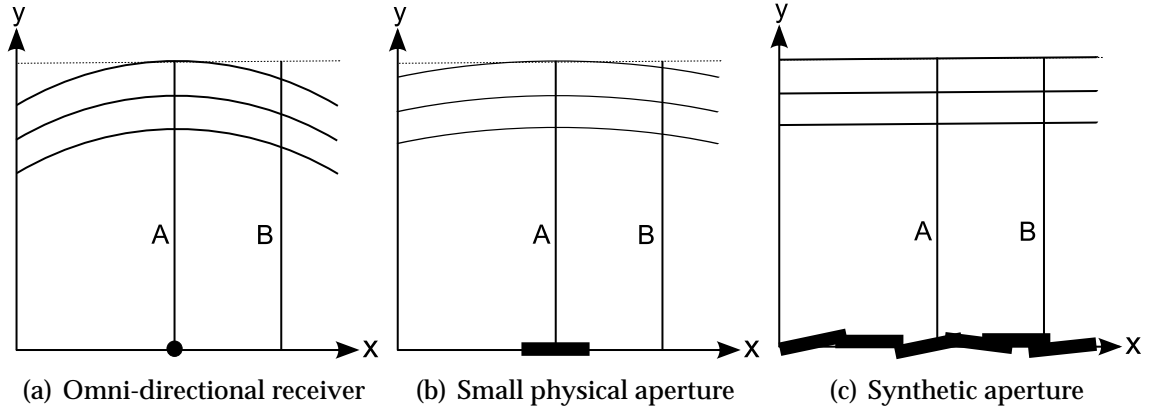


Figure 3.10: Iso-phase lines. Panel (a) illustrates a single omnidirectional receiver where the iso-phase lines follows circular arcs. Panel (b) illustrates a physical aperture. The iso-phase lines are now ellipses. In panel (c) we illustrate a synthetic aperture with straight iso-phase lines. Line A shows a broadside cross-section ($x = 0$) and line B an offset cross-section ($x \neq 0$).

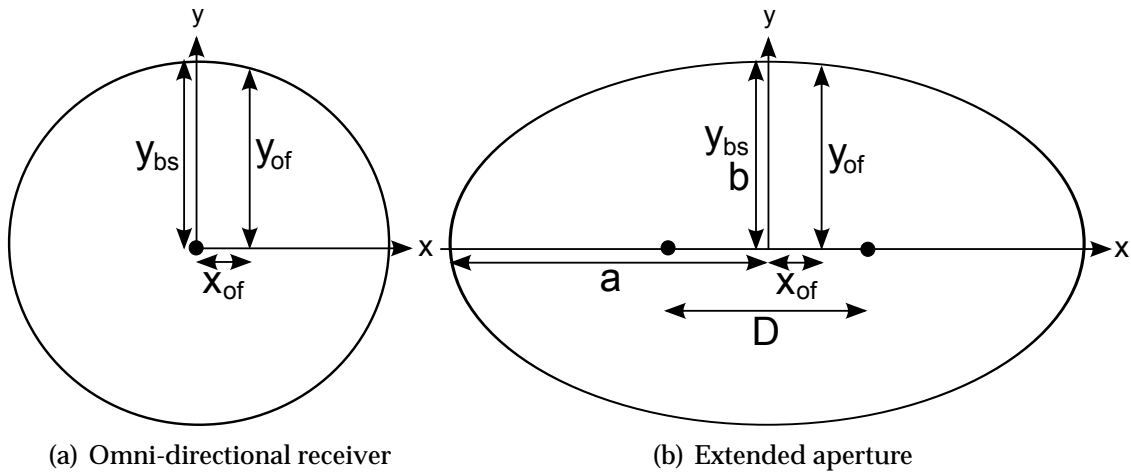


Figure 3.11: Imaging geometries. In panel (a) we describe a circular geometry. The circle has a radius of y_{bs} and (x_{of}, y_{of}) is a random point on the circle. In panel (b) we describe a ellipse geometry with a semi-major axis, a , and a semi-minor axis, b . The distance between the foci is D . The y -coordinate at $x = 0$ is given by y_{bs} , while (x_{of}, y_{of}) is a random point on the ellipse.

For an extended aperture, we model the iso-phase curves as ellipses with the aperture edges as foci. From the right panel of Figure 3.11 and the properties of ellipses (Rottmann, 1995) we now that $r_1 + r_2 = 2a$, where a is the semi-major axis. At the same time we know that distance between the foci equals the the aperture length, D , so at broadside $r_1 = r_2 = (y_{bs}^2 + (D/2)^2)^{-1/2}$. By inserting these equations into the description of an ellipse we get $a = (y_{bs}^2 + (D/2)^2)^{-1/2}$ and $b = y_{bs}$, where b is the semi-minor

axis. Thus the offset y -position and the scaling are

$$y_{\text{of}} = \frac{b}{a} (a^2 - x_{\text{of}}^2)^{-1/2} \quad (3.23a)$$

$$= \left(\frac{y_{\text{bs}}^2}{y_{\text{bs}}^2 + \left(\frac{D}{2}\right)^2} \left(y_{\text{bs}}^2 - x_{\text{of}}^2 + \left(\frac{D}{2}\right)^2 \right) \right)^{-1/2}, \quad (3.23b)$$

$$G = \left(\frac{y_{\text{bs}}^2 + (D/2)^2 - x_{\text{of}}^2}{y_{\text{bs}}^2 + (D/2)^2} \right)^{-1/2}. \quad (3.23c)$$

We see that Equation 3.23 reduces to Equation 3.22 for $D = 0$. The scaling is insignificant when $x_{\text{of}} = 0$ or when $y_{\text{bs}}^2 + (D/2)^2 \gg x_{\text{of}}^2$. Notice the similarity between the scaling and the dilation described in Equation 3.15. For real aperture depth estimation the scaling can be notable at close range, but usually becomes insignificant at long range. For SAS interferometry it is usually negligible.

3.3 Depth accuracy and baseline limitations

The theoretical accuracy of the depth measurement can be calculated by differentiating any of Equations 3.8, 3.9 or 3.20 (Franceschetti and Lanari, 1999, pages 167-170). Due to the simplicity of Equation 3.9, we choose that when we derive the theoretical accuracy. Performing the partial derivative, the standard deviation of the depth estimate, σ_z can be expressed as

$$\sigma_z = \frac{r \cos(\Phi + \Phi_0)}{D \cos \Phi} \sigma_\tau, \quad (3.24)$$

where σ_τ is the standard deviation of the time delay estimate. All other variables are assumed known. Notice that since an interferometer has a fixed *angular* resolution, the accuracy of the depth estimate is inversely proportional to the range. Other effects like the possibility for phase wraps, may be less challenging at long range. This is an effect related to the specific estimation method and does not effect the theoretical performance.

Equation 3.24 shows that the depth accuracy is increased proportionally to the interferometric baseline. A longer baseline results in less coherent signals, which again reduces the accuracy of the depth estimate. In Section 4.7 we discuss coherence and the estimation of coherence. In this section we will deduce the baseline dependent coherence. We start by noticing that the interferometric travel-path difference, $c\tau$ is equal to the $\frac{\theta}{2\pi}\lambda$, where θ is the phase-difference between the two signals and λ is the signal wavelength. The depression angle can be written

$$\Phi = \sin^{-1} \left(\frac{\lambda}{2\pi} \frac{1}{D} \theta \right). \quad (3.25)$$

The spectral shift, W , of the signal is (Bamler and Hartl, 1998)

$$W = \frac{c}{2\pi} \frac{\partial \theta}{\partial r} = -c \frac{z}{\lambda r^2} \frac{\cos \Phi}{\cos(\Phi + \Phi_0)}. \quad (3.26)$$


This expression gives the change in phase-difference as a function of range and can be used to calculate the maximum theoretical coherence between the receiver arrays. The *critical baseline*, D_{crit} is found by setting the spectral shift equal to the bandwidth, B

$$D_{\text{crit}} = \lambda \frac{B r^2 \cos(\Phi + \Phi_0)}{c z \cos \Phi}. \quad (3.27)$$


The baseline dependent coherence, k_D , can be approximated as (Bamler and Hartl, 1998)

$$k_D \approx 1 - \frac{D}{D_{\text{crit}}}, \quad \text{for } D \ll D_{\text{crit}}. \quad (3.28)$$

Panel (a) of Figure 3.12 shows the critical baseline as a function of y -position for a few different bandwidths. A higher bandwidth allows a larger baseline, which means better depth estimation accuracy. In SAR interferometry it is common to filter the spectra of the interferometric data (or even transmit slightly different frequency bands) to increase the coherence at expense of range resolution. Panel (b) of Figure 3.12 shows the baseline dependent coherence as a function of y -position for a few different baselines and a 30 kHz bandwidth. In order to achieve sufficient performance, the coherence should be above 0.9 for depression angles less than 45 degrees. In Figure 3.12 this corresponds to y -positions above 15 meters, and as the figure shows coherence filtering is unnecessary for the considered baselines.



(a) Critical baseline



(b) Coherence

Figure 3.12: Baseline limitations. Panel (a) shows the critical baseline as a function of y -position for bandwidths of 15 kHz (solid), 30 kHz (dashed) and 60 kHz (dotted). The sonar is at an altitude of 15 meters with $\lambda = 0.015$ cm and $\Phi_0 = 22$ degrees. Panel (b) shows the baseline dependent coherence as a function of y -position for baselines of 15 cm (solid), 30 cm (dashed) and 60 cm (dotted) for a 30 kHz sonar.

Chapter 4

Time delay estimation

In the previous chapter we showed that the time delay between the two interferometric signals is required in order to estimate the relative seafloor depth. In single-pass interferometry (Jakowatz et al. (1996, pages 280-282); Franceschetti and Lanari (1999, pages 31-37)), the accuracy of the estimated bathymetry is almost always limited by the accuracy of the time delay estimate. In repeat-pass interferometry (Hanssen, 2001, page 2) additional factors like co-registration and environmental corrections can be more important (Franceschetti and Lanari, 1999, Pages 173-178).

In synthetic aperture processing, time delay estimation occurs at numerous locations (Shippey et al., 1998). Micronavigation techniques use some form of time delay estimation. DPCA (Bellettini and Pinto, 2002) is based on time delay estimation on 1D time series of complex data, and very similar to the methods described in the first part of this chapter. Map-drift (Carrara et al., 1995, Pages 246-260) is based on image correlation methods, either on complex or magnitude data. Although the displacement of the images can be ascribable to several effects, the algorithms are often similar to interferometry (the difference is often that in map-drift you estimate a single offset between the images, valid for the whole image, while in interferometry you estimate a delay for each pixel). After the data are beamformed, techniques like shear averaging (Fienup, 1989) and PGA (Jakowatz et al. (1996, page 251); Callow (2003)) use phase-differences to autofocus the image. As we will show in this chapter, phase-differencing is a simple and efficient estimation method for interferometric systems.

In this chapter we discuss three different time delay estimators: the cross-correlation function (CCF), the cross-uncertainty function (CUF) and phase-differencing. We show that the cross-correlation function is a special case of the cross-uncertainty function and that phase-differencing is a special case of the cross-correlation function. We present simulation results showing situations where one has to use cross-uncertainty function in order to avoid strong decorrelation. However, by co-registering the data in advance (see Section 3.1.3) the cross-uncertainty function can be replaced by the (much faster) cross-correlation function without loss of accuracy. If the co-registration is accurate within a fraction of a resolution cell, a full correlation is unnecessary and phase-differencing will provide a fast solution.

4.1 Cross-correlation of signals with relative delay

Assume for now that the time dilation between the interferometric signals is negligible and can be interpreted as a time delay for small subsets of the data. The received signals can then be expressed as a function of time (or equivalently, a function of range)

$$f_1(t) = s(t) + n_1(t), \quad (4.1a)$$

$$f_2(t) = s(t - \delta t) + n_2(t), \quad (4.1b)$$

where we have ignored the time dilation and advanced $f_2(t)$ instead of delaying $f_1(t)$. The random processes $n_1(t)$ and $n_2(t)$ model the random noise recorded by the receivers. The statistical cross-correlation function of two random signals $f_1(t)$ and $f_2(t)$ is given by (Peebles (2001, pages 188-190); Shiavi (1999, page 355))

$$R_{f_1 f_2}(t, t - \tau) = E\{f_1(t)f_2^*(t - \tau)\}, \quad (4.2)$$

where $E\{\cdot\}$ denotes the statistical expectation or ensemble average operator. We further assume that the signals are jointly wide-sense stationary such that the cross-correlation function is a function of time difference, or local time τ , but not of absolute or global time t

$$R_{f_1 f_2}(t, t - \tau) = R_{f_1 f_2}(\tau). \quad (4.3)$$

By assuming that the noise processes $n_1(t)$ and $n_2(t)$ are jointly uncorrelated zero-mean processes, and that both noise processes are uncorrelated with the signal $s(t)$, we find that the cross-correlation function simplifies to

$$R_{f_1 f_2}(\tau) = E\{s_1(t)s_2^*(t - \delta t - \tau)\}. \quad (4.4)$$

If we now invoke the ergodic hypothesis, we may replace expectation operators with time averages, which yields the following useful formulation for the cross-correlation function

$$\begin{aligned} R_{f_1 f_2}(\tau) &= \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^T s(t)s^*(t - \delta t - \tau) dt \\ &= R_{ss}(\delta t + \tau). \end{aligned} \quad (4.5)$$

The auto-correlation function $R_{ss}(\tau)$ has its maximum at the origin, which means that the cross-correlation between $f_1(t)$ and $f_2(t)$ has its maximum for $\tau = -\delta t$, i.e.,

$$R_{f_1 f_2}(\tau = -\delta t) = R_{ss}(0). \quad (4.6)$$

The delay between the signals in additive zero-mean noise can therefore be found by locating the peak in the cross-correlation function between the two signals.

The echoes from the seafloor received by the two receivers in the interferometric sonar are *not* jointly wide-sense stationary. It is the non-stationarity we exploit when we

convert the time dependent time delay to a spatially dependent depth. The signals are slowly varying so we can adopt the idea of local stationarity and define a *local estimator* of the cross-correlation function as

$$\widehat{R}_{f_1 f_2}(\tau) = \frac{1}{2T} \int_{-T}^T f_1(t) f_2^*(t - \tau) dt, \quad (4.7)$$

where $2T$ is a finite time-interval around time t , equivalent to a range patch of a few meters.

4.1.1 Locating the peak of the cross-correlation function

A standard biased estimate of the *normalized* cross-correlation function of two discrete time complex valued *zero-mean* time series $a[i]$ and $b[i]$ of length N samples, is given by Shiavi (1999, chapter 9)

$$R_{ab}[j] = K \sum_{i=1}^{N-j} a[i+j] b^*[i], \quad j = 0, 1, \dots, N-1, \quad (4.8)$$

where $j \in \{-(N-1), \dots, -1, 0, 1, \dots, N-1\}$ is a discrete time lag index. Symmetry gives $R_{ab}[j] = R_{ba}^*[-j]$ for $j = -1, -2, \dots, -(N-1)$. Note that time index i (or lag index j) corresponds to the discrete time $t = i\Delta t$ (or time lag $\tau = j\Delta t$), where Δt is the temporal sampling interval. The normalization factor, K , is chosen so that it bounds the magnitude of the correlation coefficient between 0 and 1, and is given by

$$K = \left(\sum_{k=1}^N |a[k]|^2 \sum_{l=1}^N |b[l]|^2 \right)^{-\frac{1}{2}}. \quad (4.9)$$

It is well known that this correlation estimator is biased by a factor $1/(N - |j| - 1)$ for lag j . Although it is mathematically easy to construct an unbiased correlation estimator, see e.g., (Shiavi, 1999, chapter 9), the unbiased estimator has an unacceptably large variance, and it is not positive semi-definite, which means that it may imply negative power spectra for $y[i] = x[i]$. Since negative power spectra are physically impossible, we conclude that the unbiased correlation estimator is rendered unacceptable for physical applications.

Following Franceschetti and Lanari (1999, page 50) we define the *coherence function* as the magnitude of the normalized complex cross-correlation function $|R_{ab}[j]|$. The *coherence*, k is the maximum of the coherence function

$$k = \max_j |R_{ab}[j]|. \quad (4.10)$$

The coherence function and coherence in general are discussed further in Section 4.7.

A *coarse* estimate of the time delay can be derived from the coherence function as follows. First, we baseband (Glover and Grant, 1998, chapter 13) and cross-correlate

complex subset of the recorded data. Thereafter, a simple search gives a coarse time delay estimate $\hat{\tau}_c$ through

$$\hat{\tau}_c = j_c \Delta t, \quad \text{where } j_c = \underset{j}{\operatorname{argmax}}\{|R_{ab}[j]|\}. \quad (4.11)$$

Hence, the coarse time delay estimate corresponds to the peak of the discrete time coherence function. Notice that $k = |R_{ab}[j_c]|$. The estimate can be improved by increasing the sampling frequency of the data collection system, or by interpolation of the coherence function (or its logarithm) by means of a parabolic fit (Moddemeijer, 1991; Holm, 1993). Alternatively, the phase in the cross-correlation function can be used to improve the time delay estimate.

Since a correlation function at carrier frequency has the property that the phase is zero at the true coherence peak, we exploit this by mixing the phase of the cross-correlation at the location of the coherence peak up to carrier frequency again (i.e., we undo the basebanding by multiplying with a complex exponential). The measured phase at the coarse time lag estimate will in general deviate from zero. Hence, this measured phase deviation can be used to deduce the time delay deviation of the coarse time delay by

$$\Delta\tau = -\frac{\operatorname{arg}\{R_{ab}[j_c]e^{j2\pi f_0\tau_c}\}}{2\pi f_0}, \quad (4.12)$$

where f_0 is the center frequency of the signal. We now correct the coarse time delay estimate by this delay deviation, and based on this discussion, propose the following estimator for the *fine* time lag

$$\hat{\tau}_f = \hat{\tau}_c + \Delta\tau. \quad (4.13)$$

The probability of 2π wrap errors in the fine estimate is minimized, but may still occur since the 3 dB resolution of the coherence peak is $1/B$ while the wrap interval is $1/f_0$. B is here the bandwidth of the signal. For systems with $B < f_0$, *super-resolution* is in principle required to avoid wrap-errors. This can be achieved by upsampling the cross-correlation function (Mitra, 2002, pages 47-48), provided the SNR is sufficiently large. Even though the coarse estimate may be too inaccurate to determine the correct 2π -interval, the phase can still be reliable. A 2π -correction to the phase can therefore be applied, provided one has *a priori* knowledge to substantiate this. It should be noticed that the same performance can be achieved by cross-correlating the real signals at carrier frequency, but then the 2π problem is replaced by the problem of selecting the correct coherence peak (since the correlation function of two real signals looks like a high-frequency modulation of the magnitude of the corresponding complex correlation function).

To avoid interpolation errors, upsampling should always be performed on basebanded data rather than on data at carrier frequency (Hawkins, 1996, chapter 4). Therefore one should always correlate basebanded time series and correct the phase of the coarse correlation peak only, as showed in Equation 4.12, rather than correlating the time series at carrier frequency and directly read out the correct phase.

4.1.2 Magnitude-correlation

A change of the phase-differences between the input signals over the correlation window will cause decorrelation. The phases, which give us the high accuracy of the time delay estimate, thus reduce the robustness of the estimate. Therefore it can sometimes be advantageous to cross-correlate the magnitude of two complex signals (magnitude correlation) rather than cross-correlating the complex signal directly (complex correlation). One of these circumstances is when baseline decorrelation is significant (see Section 4.6.2). Note that the Nyquist-frequency doubles when taking the magnitude of the complex values, since the number of samples is halved.

An estimator of the magnitude cross-correlation function is given by

$$R_{ab}^m[j] = K \sum_{i=1}^{N-j} |a[i+j]| |b^*[i]|, \quad j = 0, 1, \dots, N-1. \quad (4.14)$$

Since the magnitude operator is nonlinear, the magnitude-correlation and the complex-correlation differ in general. Figure 4.1 shows the normalized magnitude- and complex-correlation function of a pulse compressed linear chirp. The figure illustrates how the complex-correlation function has a narrower peak than the magnitude-correlation function. This is due to the fact that the magnitude-correlation is a sum of magnitudes and the complex-correlation is a sum of magnitudes multiplied by exponential phase terms. Since $|A||B| \geq |AB|$ holds for complex numbers A and B , the same is valid for the sum. When two signals are completely coherent, the magnitude- and complex-correlation are identical, but when the phase starts to decorrelate, the complex-correlation has lower values. Therefore the pixels in the neighborhood of the peak are more suppressed in the complex-correlation, increasing the resolution.

On the other hand, the magnitude-correlation function is more robust to noise and has higher correlation when the SNR is low. Figure 4.2 shows $\mathbf{a}^H \mathbf{b}$ and $|\mathbf{a}|^T |\mathbf{b}|$, where $\mathbf{a} = [a_1, a_2, \dots, a_n]$ and $\mathbf{b} = [b_1, b_2, \dots, b_n]$ are complex time series consisting of a signal in uncorrelated additive white noise. The superscript H denotes the Hermitian operator defined as the complex conjugate of the transpose of the vector. The sum of the complex products approximate a random walk, while the sum of the magnitude products always contribute along one axis. The length of the straight lines in Figure 4.2 represents the correlation value. However, for positive SNR, we have found that the the more precise complex correlation always is preferable over the more robust magnitude correlation.

4.1.3 Accuracy of the time delay estimate

The accuracy of the coarse and fine delay estimates can be found from the SNR, ρ , of the correlation function. An important relation is therefore the relation between ρ and the peak of the normalized coherence function (Zebker and Villasenor, 1992). Again, consider the two time series $a[i]$ and $b[i]$. If we assume that they consist of a synchronized

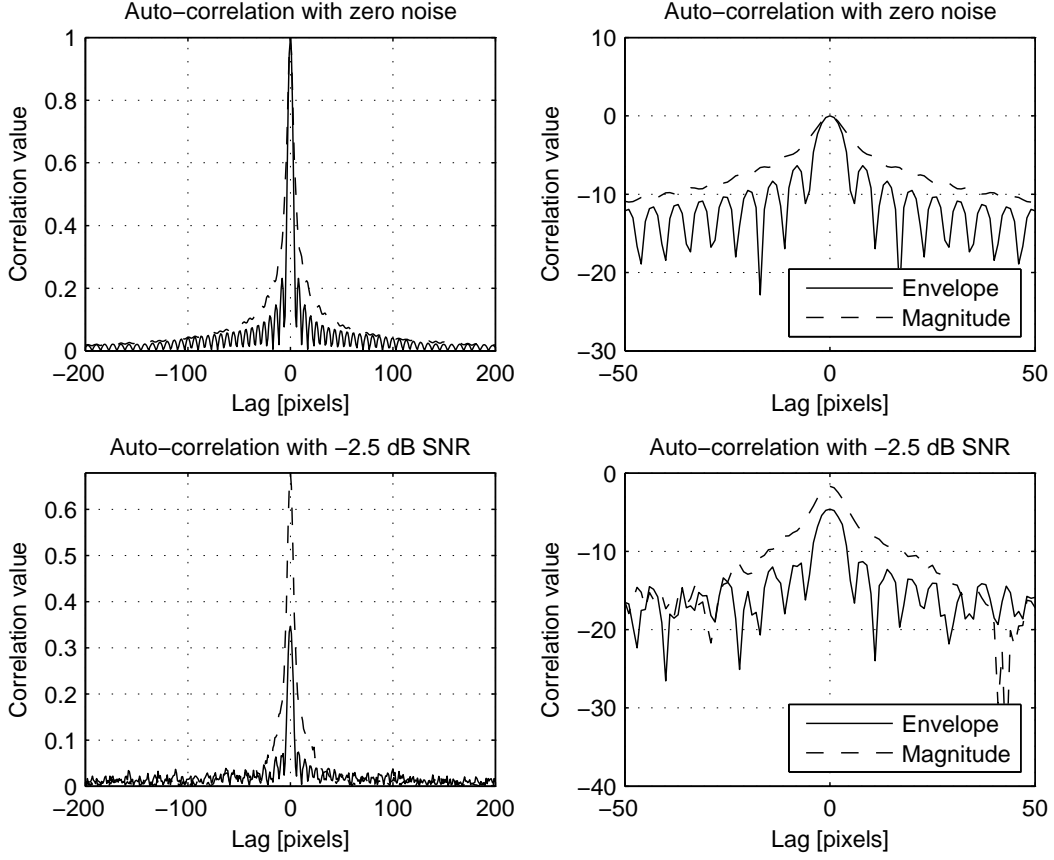


Figure 4.1: Normalized magnitude-correlation function (solid) and complex-correlation function (dashed) for a pulse compressed chirp. The panels on the right show blow-ups of the peaks. The upper panels show the zero-noise case and the lower panels the case with uncorrelated additive white noise and a SNR of -2.5 dB.

signal part and additive noise, we can write them in vector form as

$$\mathbf{a} = S_a \mathbf{s} + N_a \mathbf{n}_a, \quad (4.15a)$$

$$\mathbf{b} = S_b \mathbf{s} + N_b \mathbf{n}_b, \quad (4.15b)$$

where S_a , S_b , N_a and N_b are scalars, \mathbf{s} a discrete vector with unity energy, and \mathbf{n}_a and \mathbf{n}_b two discrete noise realizations with unity energy. The corresponding SNRs are $\rho_a = S_a^2/N_a^2$ and $\rho_b = S_b^2/N_b^2$. The coherence between \mathbf{a} and \mathbf{b} (at zero-lag) can now be formulated as

$$k = \frac{|\mathbb{E}\{\mathbf{s}_a^H \mathbf{s}_b\}|}{\sqrt{\mathbb{E}\{|\mathbf{s}_a|^2\} \mathbb{E}\{|\mathbf{s}_b|^2\}}}, \quad (4.16)$$

where the superscript H denotes the Hermitian operator and $\mathbb{E}\{\cdot\}$ the expectation operator. We now assume that the signal is uncorrelated with the noises, $\mathbb{E}\{\mathbf{s}^H \mathbf{n}_a\} =$

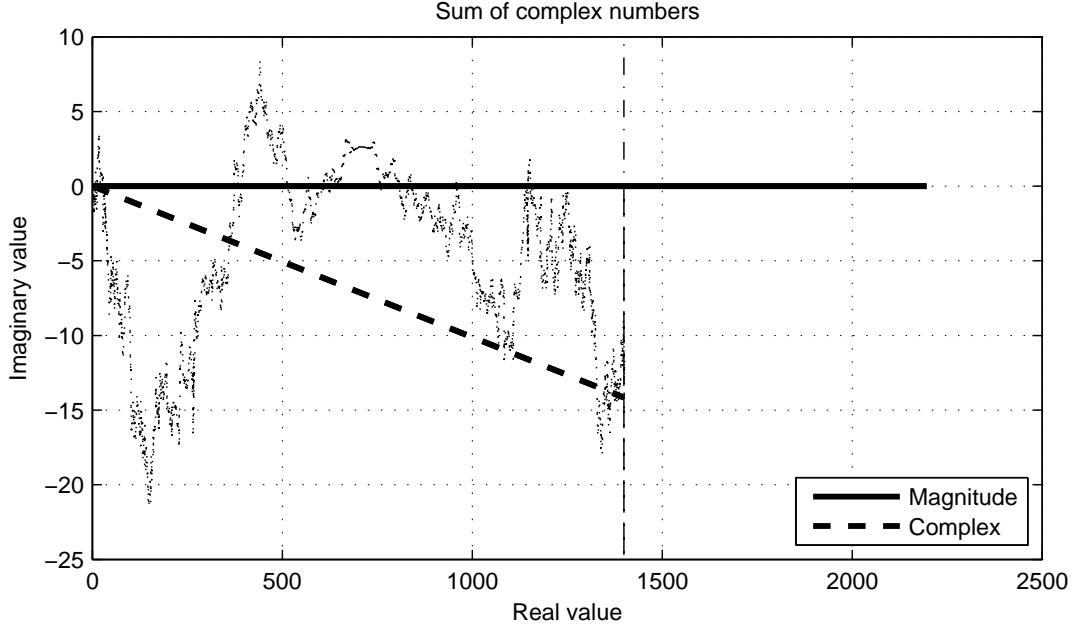


Figure 4.2: Sum of the products of the magnitude of two time series (solid) and the sum of the complex conjugate products (dotted). The dashed line shows the length of the complex sum and the dash-dotted line shows a circle with radius equal to the length of the complex sum.

$E\{s^H \mathbf{n}_b\} = 0$, and that the two noise realizations are uncorrelated, $E\{\mathbf{n}_a^H \mathbf{n}_b\} = 0$. Equation (4.16) then simplifies to

$$k = \frac{S_a S_b}{\sqrt{S_a^2 + N_a^2} \sqrt{S_b^2 + N_b^2}} = \frac{S^2}{S^2 + N^2} = \frac{\rho}{1 + \rho}, \quad (4.17)$$

where we assumed $S = S_a = S_b$ and $N = N_a = N_b$ and where $\rho = S^2/N^2$. This result yields an easy and important relation between the coherence and the SNR.

If one assumes that the transmitted signal and the noise are white Gaussian random processes, analytical expressions of the approximate theoretical Cramér-Rao Lower Bounds (CRLB) of the coarse and fine estimates can be expressed in terms of the SNR as (Quazi, 1981)

$$\text{std}\{\hat{\tau}_c\} \geq \sigma_c \equiv \frac{\sqrt{3}}{\pi} \frac{1}{B} \frac{1}{\sqrt{BT}} \sqrt{\frac{1}{\rho} + \frac{1}{2\rho^2}}, \quad (4.18)$$

$$\begin{aligned} \text{std}\{\hat{\tau}_f\} \geq \sigma_f &\equiv \frac{1}{2\pi f_0} \frac{1}{\sqrt{BT}} \frac{1}{\sqrt{1 + B^2/12f_0^2}} \sqrt{\frac{1}{\rho} + \frac{1}{2\rho^2}} \\ &\approx \frac{1}{2\pi f_0} \frac{1}{\sqrt{BT}} \sqrt{\frac{1}{\rho} + \frac{1}{2\rho^2}}, \end{aligned} \quad (4.19)$$

where, $\text{std}\{\cdot\}$ denotes the standard deviation, $\hat{\tau}_c$ and $\hat{\tau}_f$ are the coarse and fine time delay estimates, respectively, and σ_c and σ_f are the corresponding lower bounds for the

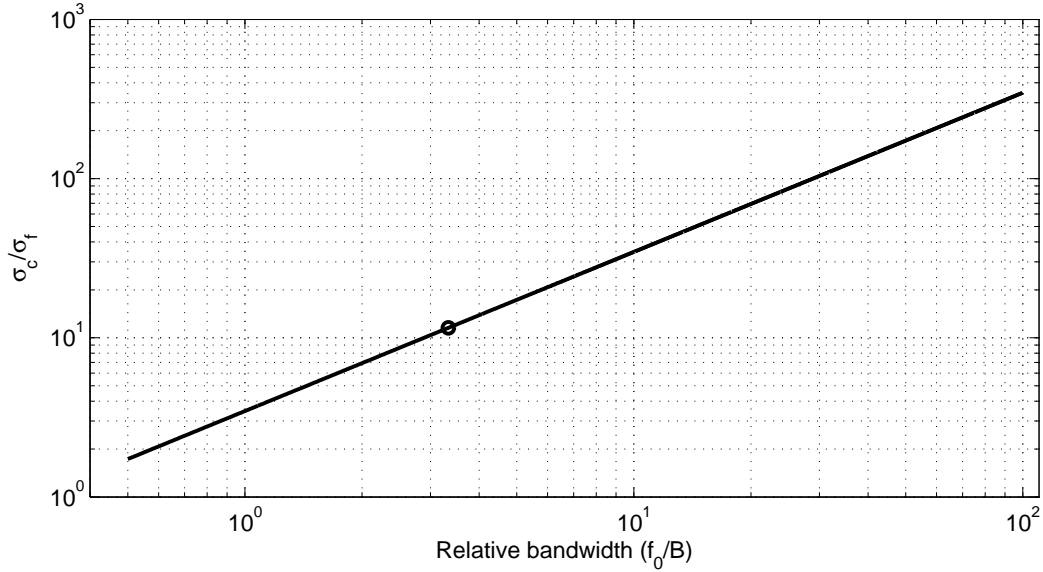


Figure 4.3: Ratio between the CRLBs of the fine and of the coarse estimate as a function of relative bandwidth. The circle shows the ratio and relative bandwidth for a sonar with $f_0/B = 3.33$.

standard deviations of these time delay estimates. Furthermore, T is the time interval of the correlation window, and BT is the number of independent samples in the correlation. The final step in Equation 4.19 involves the simplifying assumption $B^2 \ll 12f_0^2$. The ratio between the CRLBs of the fine and of the coarse estimate is a function of the relative bandwidth, f_0/B , only. Figure 4.3 shows this ratio as a function of relative bandwidth. Note that the ability to resolve the phase ambiguity is inversely proportional to f_0/B . Inserting HISAS parameters ($f_0 = 100$ kHz, $B = 30$ kHz, see Section 6.1) we obtain

$$\frac{\sigma_c}{\sigma_f} = 2\sqrt{3} \left(\frac{f_0}{B} \right) \approx 11.5, \quad (4.20)$$

which is indicated by a circle in Figure 4.3. Increasing the bandwidth reduces the difference between the CRLBs of the fine and of the coarse estimate. Since the coarse estimate is used to resolve the 2π -ambiguity in the fine estimate, large bandwidth is always preferable.

The expressions for the Cramér-Rao lower bounds are based on an idealized scenario. The accuracy of the delay estimate can be limited by several additional factors, e.g.:

- The reflectivity distribution is heavy-tailed (Hanssen et al., 2003; Lyons and Abraham, 1999).
- Geometrical decorrelation (which is caused by footprint shift (Lurton, 2000), dilation (Gatelli et al., 1994) and baseline decorrelation (Lurton, 2000)) can cause the

signal part of the echo to differ for the two interferometric receiver arrays.

- Reverberation (clutter) is unwanted echo (e.g. echo from living organisms or echo from a transducer sidelobe). These parasitic echoes can give rise to signals from unwanted directions (Lurton, 2002, page 122-125).
- Multipath, or coherent replicas reflected from medium boundaries is echo which do not have a direct path into the sonar. Multipath can smear the signal in time and frequency, due to the roughness of the boundaries (Brekhovskikh and Lysanov, 1982, chapter 9).
- Ocean variability (Dozier et al., 1991) can cause phase and amplitude fluctuations (Brekhovskikh and Lysanov, 1982, pages 213-216) of the signal (multiplicative noise).
- Additive noise, either acoustic background noise, acoustic interference or electrical noise, may be non-Gaussian and correlated (correlated noise may cause a nonphysically high coherence) (Lurton, 2002, pages 107-118).

For a more thorough discussion of model errors see Section 4.6. Baseline decorrelation is discussed further in Section 4.7.

4.1.4 Center frequency shift correction

A significant contribution to the bias in the delay estimate is a shift of the effective center frequency (i.e. f_0 is incorrect in Equation 4.12). The measured spectrum of an echo returned from a small part of the seafloor is a random realization of the transmitted spectrum. The received spectrum is affected by a number of different factors:

- The echo is a contribution from multiple scatterers, which interfere constructively and destructively.
- Scattering from some parts of the seafloor may be frequency-dependent (dispersive scattering) (Jackson and Richardson, 2007, pages 331-376).
- Absorption in the medium may be frequency-dependent (Lurton, 2002, pages 19-26).
- The antenna directivity is frequency-dependent (Van Trees, 2002, pages 23-37).

The latter can be measured and calibrated in a controlled environment, but our experience is that the first three factors listed above are dominant. The sum of these effects cause an uneven weighting of the received frequency components. For narrowband systems, unevenness in the spectrum is less significant, but for wideband systems this can change the centroid of the spectrum. In Shippey et al. (1998) it was discussed how a cross-correlation based time delay estimator strictly applies only when the received

amplitude spectra are symmetrical around the center frequency. We have found that a cross-correlation based method can still be a good estimator for the time delay if we model shift or unevenness of the spectra as changes in the *effective* center frequency.

A shift of the center frequency does not affect the coherence as it is the same for both arrays. The result is a linear trend as a function of lag on the basebanded phase of the correlation function. In Figure 4.4 we show an estimate of a cross-correlation function from real data collected during the InSAS-2000 trial (Wang et al., 2001), where the phase in the lower panel should, according to theory, be constant within $\pm 1/B$. The linear trend of the phase in the lower panel is caused by a *stochastic* frequency shift. It resembles, but is different from the *deterministic* wavenumber shift, which is a geometric effect (Gatelli et al., 1994).

As a result of this shift of the center frequency, errors are likely if the observed phase is used directly in the delay estimation. To compensate for the trend, one can estimate a corrected center frequency from the centroid of the measured spectrum. Note that even if the centroids of the two received time series differ, one common frequency should be used in all conversions between baseband and carrier frequency. If two different modulation frequencies are applied, significant errors will be induced in the delay estimate. However, the estimated centroid is usually not accurate enough to improve the accuracy of the delay estimate. We have instead implemented a method which utilizes the fact that the phase is still correct at the true correlation peak. We first perform a parabolic fit to the coarse estimate, thereafter we interpolate the phase linearly around this location, and finally, we use this location and the interpolated phase in Equation 4.13. Monte Carlo simulations have shown that this method gives improved accuracy for $\text{SNR} \leq 60$ dB, which is sufficient for most applications. A linear regression estimate of the frequency has potential to increase the accuracy further, but this remains to be tested.

4.2 Cross-correlation of signals with relative delay and relative dilation

In Section 3.1.2 we showed that there is a relative time dilation between the received signals, not a time delay. We also argued that for a small patch, a dilation causes an apparent time delay, but with the dilation still present. We can therefore model the received signals as delayed *and* time dilated relative to each other. We start with the model we developed in Section 3.1.2 and add a time delay to one of the signals

$$f_1(t) = s(\alpha_1(t) \cdot t) + n_1(t), \quad (4.21a)$$

$$f_2(t) = s(\alpha_2(t) \cdot (t - \delta t)) + n_2(t). \quad (4.21b)$$

In general, the dilation itself is time dependent, but as we saw in Section 3.1.2, the time dependence is weak for a small patch at relative long range, so in the subsequent calculations we ignore the time dependence. As before we only consider a relative dilation

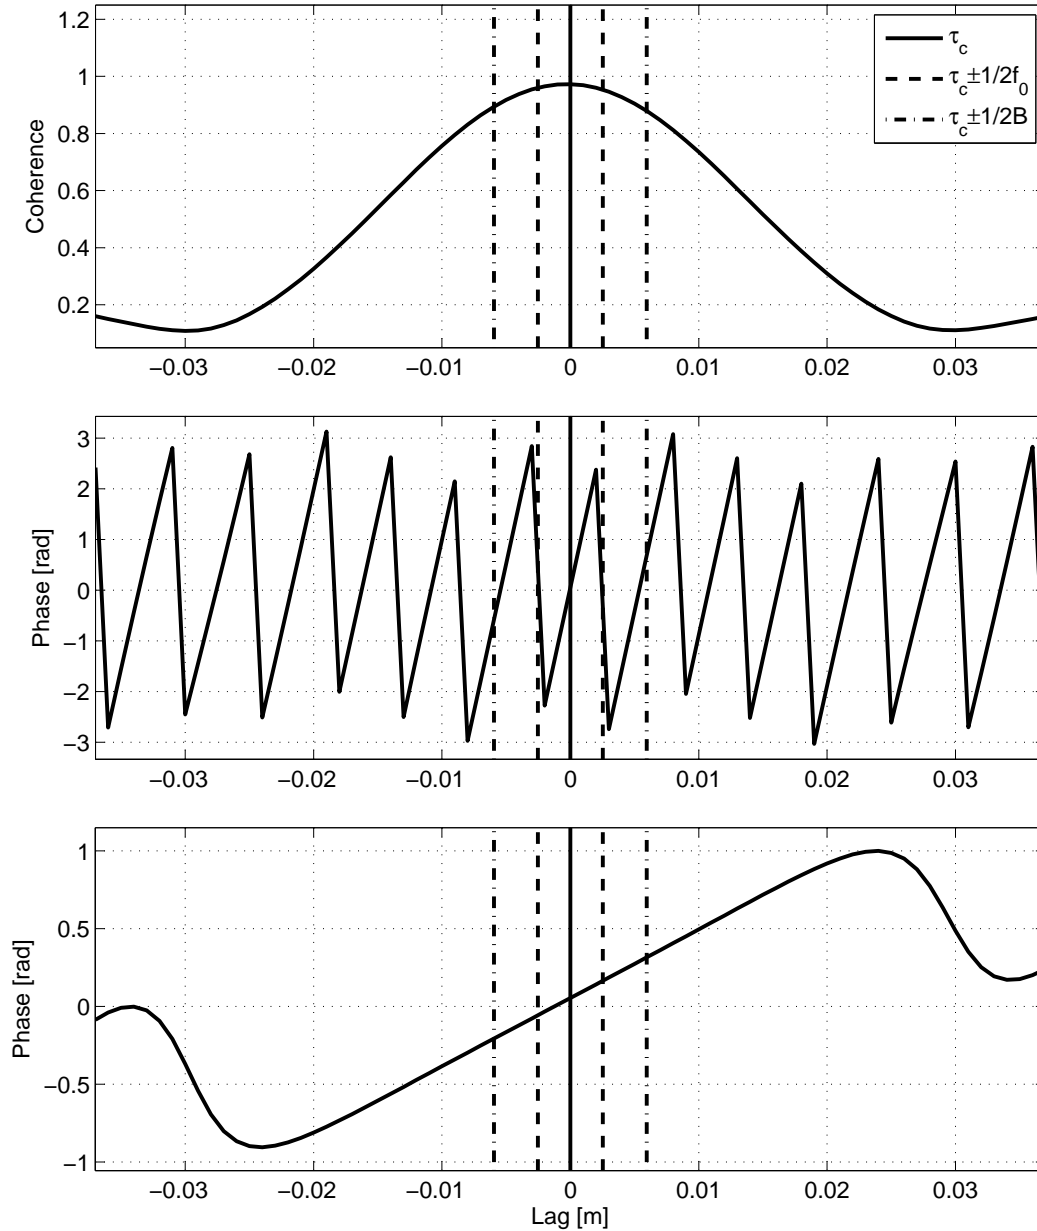


Figure 4.4: Cross-correlation function estimated from real data collected in Italy, 2000. The sonar had a center frequency of 150 kHz and a bandwidth of 64 kHz. The upper panel shows the magnitude (the absolute value of the envelope) of the estimated correlation function (thick), the center panel shows the estimated phase at the carrier frequency (thick) and the lower panel shows the basebanded phase (thick). The magnitude peak, the 2π wrap-interval and the 3 dB resolution are also shown.

$\alpha(t)$

$$f_1(t) = s(t) + n_1(t), \quad (4.22a)$$

$$f_2(t) = s(\alpha(t - \delta t)) + n_2(t), \quad (4.22b)$$

where $\alpha = \alpha(t_0) \approx \alpha(t)$ is the relative dilation-factor evaluated at the center of the patch. Again, we assume that the noises are uncorrelated with each other and with the signal. The cross-correlation function of two signals with relative delay and relative time dilation is equal to

$$R_{f_1 f_2}(\tau) = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^T s(t) s^*(\alpha(t - \delta t - \tau)) dt. \quad (4.23)$$

which is difficult to interpret unless $\alpha = 1$.

4.3 Wideband cross-uncertainty function of signals with relative delay and relative dilation

In the previous section we found that we need a more general mathematical tool than the cross-correlation function. The wideband cross-uncertainty function, $\chi_{f_1 f_2}(\beta, \tau)$, is a function of both time dilation (or time scaling) and delay, and it has the cross-correlation function as a special case (Altes (1973); Ricker (2003, pages 153-185)),

$$\chi_{f_1 f_2}(\beta, \tau) = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^T f_1(t) f_2^*(\beta(t - \tau)) dt. \quad (4.24)$$

The so-called cross-ambiguity function is given by $|\chi_{f_1 f_2}(\beta, \tau)|^2$, and is simply the modulus squared of the complex valued uncertainty function. The narrowband cross-uncertainty function is defined through a simple phase-shift instead of a time scaling. We will in this thesis refer to the *wideband* cross-uncertainty function unless otherwise explicitly stated.

We now consider the cross-uncertainty function of the two received signals modeled with relative delay and relative dilation, which yields

$$\chi_{f_1 f_2}(\beta, \tau) = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^T s(t) s^*(\alpha(\beta(t - \tau) - \delta t)) dt.$$

Again, we assume that the noises are uncorrelated with each other and with the signal. The cross-uncertainty function therefore reduces to the auto-correlation function for $\beta = 1/\alpha$ since

$$\chi_{f_1 f_2} \left(\beta = \frac{1}{\alpha}, \tau \right) = R_{ss}(\tau + \alpha \delta t). \quad (4.25)$$

Since $R_{ss}(0) \geq \|R_{ss}(\tau)\|$ for all $\tau \neq 0$, we can find the relative delay and the relative dilation by searching for the maximum of the cross-uncertainty function

$$\chi_{f_1 f_2} \left(\beta = \frac{1}{\alpha}, \tau = -\alpha \delta t \right) = R_{ss}(0). \quad (4.26)$$

4.3.1 Implementation of a wideband CUF estimator

We have implemented an estimator for the cross-uncertainty function using a non-linear least-squares data-fitting technique (Kay, 1993, pages 254-260). The objective function for the search function is

$$J(\beta) = \left(1 - \max_j \frac{|\sum_n f_1[n] f_2^*[\beta(n-j)]|}{\sqrt{\sum_n |f_1[n]|^2 \sum_n |f_2[\beta n]|^2}} \right)^2, \quad (4.27)$$


where n denotes discrete time, j denotes discrete time delay, and $f_1[n]$ and $f_2[n]$ are the discrete time representations of the signals $f_1(t)$ and $f_2(t)$. In other words, for each dilation-factor, we scale one of the signals with a spline interpolator, we cross-correlate the signals with each other, and then we find the peak correlation value. The estimated dilation-factor is thus the dilation-factor which matches the highest correlation value (Kjellesvig, 2006). This method has its weaknesses in that there are errors in the interpolation, and that a termination tolerance has to be specified by the user.

A more complex method would be to perform a two-dimensional grid search to maximize both β and j simultaneously, but this is significantly more computationally intensive. In Chan and Ho (2005), a faster implementation is suggested, which utilizes an approximated sinc-interpolator in an iterative scheme. However, the definition of the cross-uncertainty function in Chan and Ho (2005) differs from the definition most commonly found in the literature and the implications thereof are currently not investigated.


4.4 Numerical study of the time delay accuracy

To evaluate the performance of the cross-correlation function and the cross-uncertainty function on delayed and dilated signals, we have performed a set of Monte Carlo simulations. The simulated data are generated by a point scatter simulator and then run through a matched filter (see the introduction of Chapter 7). The transmitted signal is a 4 ms chirp with 40 kHz bandwidth and 100 kHz center frequency. The sonar has a vertical baseline of 30 cm.


Figure 4.5 shows the results for 1000 Monte Carlo realizations where the seafloor is located at the symmetry axes of the two receivers (i.e. $z_0 = -\delta z/2$). Point scatterers with Gaussian amplitude distribution are uniformly distributed on the seafloor along the y -axis. For each realization we have inserted different levels of additive white Gaussian




(a) STD of the error of the time delay estimate



(b) RMS-value of the error of the time delay estimate



(c) Estimated SNR



(d) Estimated dilation-factor

Figure 4.5: Results from time delay estimation on 1000 Monte Carlo simulations in AWGN. We have selected a 3 meters patch at 50 meters range with the seafloor located at the symmetry-axis of the interferometer (i.e. $z_0 = -0.15$ m). The plots show results from CCF estimates, co-registered CCF estimates (CCF2), CUF estimates and co-registered CUF estimates (CUF2) as a function of SNR. The different Panels show the following: (a) STD of the error in time delay estimates with the CRLB included; (b) RMS-values of the same error with the CRLB included; (c) estimated SNRs (from the maximum correlation coefficients) for the four different techniques; (d) estimated dilation-factors for the CUF estimates.

noise (AWGN) into the data, selected a patch of three meters, and estimated the time delay from the cross-correlation function and the cross-uncertainty function. We have then co-registered the data based on these estimates and repeated the described method. The second iteration reduces the effect of scallop loss (Harris, 1978).

Figure 4.5(a) shows the standard deviation (STD) of the error in the time delay estimate for the four estimates mentioned above. The Cramér-Rao lower bound (Quazi,

1981) is also included. The cross-correlation function and the cross-uncertainty function performs equally well on time delayed signals with no time dilation. They are both close to the theoretical Cramér-Rao lower bound. At signal-to-noise levels above 20 dB co-registration is needed in order to achieve theoretical performance. This is due to the footprint shift effect (Lurton, 2000). At negative SNR levels, the estimates deviates from the theoretical curve since 2π wrap-errors are introduced as an additional source of error. This deviation could have been removed by replacing the CRLB with the Ziv-Zakai lower bound (Ziv and Zakai, 1969; Bell et al., 1997).

Figure 4.5(b) shows the root-mean-square (RMS) value of the four estimates. Clearly there is no visible bias in the estimates and they both perform equally well in this scenario.


Figure 4.5(c) shows the SNRs estimated from the peak of the CCF and the CUF (see Section 4.1.3), as a function of the inserted SNR levels. The curves match theory very well down to zero dB, were the peak correlation value asymptotically approaches a lower threshold. Also notice that the estimates before co-registration suffers a small SNR loss at high SNRs.

Figure 4.5(d) shows the estimated relative dilation-factors for the CUF estimates and the co-registered CUF estimates. Since the signals are not dilated, this factor should be equal to one. For positive SNRs the estimates are indeed very close to this prediction, but at negative SNRs there is a mismatch between the theoretical and estimated dilations. This may deteriorate the CUF estimates in this region, but as we have seen from Figures 4.5(a) and 4.5(b), 2π wrap-errors already render these estimates useless.


In sum, Figure 4.5 shows that the cross-uncertainty function performs equally well as the cross-correlation function for time delayed signals with no time dilation. They both suffer from the fact that 2π wrap-errors deteriorate the time delay estimates at negative SNRs. At SNRs above 20 dB both methods need a co-registration in order to achieve theoretical performance.

Figure 4.6 shows the same results as Figure 4.5, but now the seafloor is located 15 meters below the sonar. The center y -position of the patch is at 50 meters which gives a dilation-factor of approximately 0.999. Again a random realization of the reflectors and a random displacement of the scatterers for the second receiver is inserted.


In Figure 4.6(a) we see that the time delay estimates using CCF deviates from the CRLB for almost all SNR levels. This is expected since the signal model is incorrect. However, the corresponding results when applying CUF follows the theoretical CRLB closely. There is a small deviation from theoretical performance, which increases for high SNRs. There are at least two explanations for this deviation: Firstly, we assume a constant dilation over each patch. In reality, the dilation is changing non-linearly over the patch (see Figure 3.5). This error in the model will limit the accuracy when the SNR is sufficiently high. Secondly, we have tested this method on band-limited linear frequency-modulated (LFM) signals with realistic sidelobe performance. This means that a reflector at 50 meters range will have sidelobes spread out over a region around the reflector. The sidelobes have the same dilation as the reflector, but will be treated as signals originating from other locations. This effect is inherent when the point-spread




(a) STD of the error of the time delay estimate



(b) RMS-value of the error of the time delay estimate



(c) Estimated SNR



(d) Estimated dilation-factor

Figure 4.6: Results from time delay estimation on 1000 Monte Carlo simulations in AWGN. We have selected a 3 meters patch at 50 meters range with the seafloor located 15 meters below the sonar (i.e. $z_0 = 15$ m). The plots show results from CCF estimates, co-registered CCF estimates (CCF2), CUF estimates and co-registered CUF estimates (CUF2) as a function of SNR. The different Panels show the following: (a) STD of the error in time delay estimates with the CRLB included; (b) RMS-values of the same error with the CRLB included; (c) estimated SNRs (from the maximum correlation coefficients) for the four different techniques; (d) estimated dilation-factors for the CUF estimates.

function differs from impulse signals and may start to limit the time delay accuracy for high SNRs.

In Figure 4.6(b) we see a clear bias in the time delay estimate using CCF, but only a negligible bias when using CUF. However, the CCF has such a low overall performance, that an attempt to correct the bias is without interest. Figure 4.6(c) shows that the estimated SNR also significantly improves when using CUF relative to CCF. The SNR calculated from the peak correlation value in the CCF deviates from the theory almost immediately, while the SNR calculated from CUF follows the theory relatively close. Figure 4.6(c) also shows that the results are improved by reducing the footprint shift effect. In Figure 4.6(d) we show that the estimate of the relative dilation-factors again is very stable, but now at a value of around 0.999. This is a moderate, but quite realistic dilation factor for sidelooking sonars.

Figure 4.6 illustrates that for a sidelooking sonar geometry, the cross-correlation based time delay estimate is sub-optimum and will limit the accuracy of the relative depth estimate. The cross-uncertainty function better matches the signal model and improves the accuracy of the time delay estimate significantly. The standard deviation for high SNRs is reduced by an order of magnitude and a small bias present in the CCF estimate is removed. The estimated SNR from the peak correlation value is also increased significantly, indicating a better match of the signals after CUF.

In synthetic aperture radar it has been suggested to eliminate the change of frequency in the signals received from different angles by trimming the spectra at the edges (Gatelli et al., 1994). This frequency change is the same as we have described as a time dilation. The trimming is performed to ensure full overlap of the spectra at the expense of a reduced bandwidth. This method works well for a theoretical spectrum with zero amplitude in the stop band and at a flat response inside the passband. However, due to the constructive and destructive interference of the scatterers, the received spectrum is a random realization of the transmitted spectrum. Figure 4.7 shows an example of the variations in such a spectrum. From the properties of the Fourier transform we know that dilation or time scaling in time domain also dilates the spectrum (McClellan et al., 2003, pages 322-323). For a small patch of the spectrum we therefore see an apparent frequency shift. This shift will not be addressed by cutting the spectra at the edges. However, the re-gridding performed in the CUF-algorithm will, as we can see in Figure 4.7, remove most of this effect and increase the correlation between the signals. An equivalent approach suggested by Gatelli et al. (1994) is to apply a range- and geometry-dependent frequency filter.

An alternative method of estimating the dilation in the received time signals is to estimate the shift in the spectra of the data and relate that to a dilation. This can be performed by a center of mass estimation or by running a cross-correlation function directly on the spectrum.

Another important question is the performance of the CCF and CUF estimators as a function of dilation for a given SNR. In our sidelooking sensor geometry we have varied the range (and thereby the dilation) for a fixed seafloor depth, and then performed a Monte Carlo simulation similar to the simulations presented above. Figures 4.8 and 4.9




Figure 4.7: Zoom of an oversampled spectrum from one of the received time signals as a function of data samples. The signal is selected from one of the realizations of the data presented in Figure 4.6. The solid line is the spectrum for array one and the dot-dashed line the spectrum for array two. The dashed line is the spectrum for array two after the interpolation performed in the CUF algorithm (almost on top of the solid line).

show the results from such a study.

The suppression angle in Figures 4.8 and 4.9 is defined as $\cos^{-1}(z_0/r)$ where $z_0 = 15$ m and r is the range from receiver #2 to the center of the seafloor patch. Figure 4.9 shows that CUF is superior for all suppression angles larger than five degrees. In this region and for this SNR, the dilation is the limiting factor when using the CCF estimator. From Figure 4.9(a) we see that at 35 degrees suppression angle, a CCF based method gives a correlation coefficient of 0.5, equivalent to 0 dB SNR, while CUF has a correlation coefficient corresponding to > 20 dB. Note the effect of the co-registration showed in Figure 4.9(b). Using co-registered CUF is important even for very high SNRs if the predicted SNR is to be used as a quality measure.

Figure 4.9(c) shows that CUF reduces the standard deviation of the error in the time delay estimate between one and two orders of magnitude compared to CCF. From Figures 4.9(c) and 4.9(d) we also see that the effect of scallop loss is strongly related between the dilation-factor estimate and the time delay estimate. The co-registration eliminates the scallop loss *and* increases the estimated SNR significantly.

The results presented in this section indicate that CUF is a better time delay estimator for scenarios where the effect of dilation overshadows the effect of additive noise. For a standard sidelooking sonar system this is almost always the case at near range. Only at far range where the dilation is less important, the CUF has the same performance as CCF. It should be noted that the effect of the dilation on CCF performance is a function

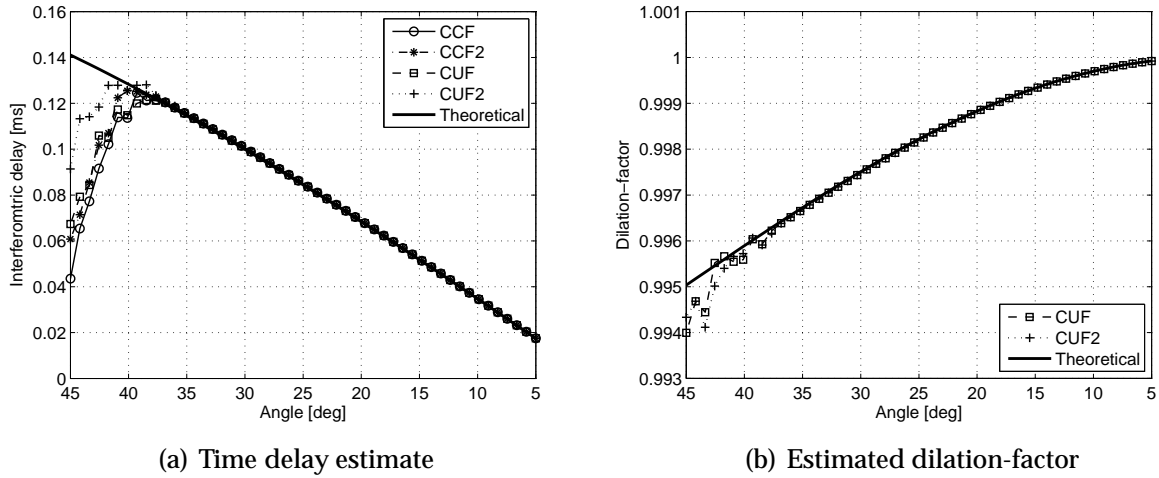


Figure 4.8: Results from 1000 Monte Carlo simulations in AWGN. The SNR in the simulations is 25 dB. We have selected a 3 meters patch at varying range, with the seafloor located 15 meters below the sonar (i.e. $z_0 = 15$ m). The plots show results from CCF estimates, co-registered CCF estimates (CCF2), CUF estimates and co-registered CUF estimates (CUF2) as a function of suppression angle. The different panels show the following: (a) theoretical and estimated interferometric time delay; (b) theoretical and estimated dilation-factor.


of the patch (or window) size. A larger patch will cause more decorrelation due to dilation. However, a larger patch does also increase the accuracy in regions limited by additive noise. As such, CUF provides the possibility to use larger patches, and thereby better performance, than CCF.

As described in Section 3.1.3, the coherence-loss due to time-shift and time dilation can be reduced by co-registering the signals in ground-range. This method does not assume a linear dilation and hence is a higher order model than the CUF estimator, but it requires accurate knowledge of the geometry of the environment. A CUF-based time delay estimation can then be applied on the co-registered data, in case there are any residual dilation left. In cases where the residual dilation is insignificant, CCF is preferable to CUF due to simplicity and time-consumption.


4.5 Phase-differencing

In the SAR literature the interferometric delay is usually found by means of phase-differencing. The phase-difference is estimated using a maximum likelihood estimator (Hanssen, 2001, Page 92), which only differs from the complex cross-correlation in that it is two-dimensional and that it does not shift the data. The phase-difference estimator is a 2D zero-lag complex cross-correlation.


The phase-difference, θ , between two vertically separated receivers, is directly linked




(a) Maximum Correlation coefficient



(b) Estimated SNR



(c) STD of the error of the time delay estimate



(d) STD of the error of the dilation estimate

Figure 4.9: Results from 1000 Monte Carlo simulations in AWGN. The SNR in the simulations is 25 dB. We have selected a 3 meters patch at varying range, with the seafloor located 15 meters below the sonar (i.e. $z_0 = 15$ m). The plots show results from CCF estimates, co-registered CCF estimates (CCF2), CUF estimates and co-registered CUF estimates (CUF2) as a function of suppression angle. The different Panels show the following: (a) maximum correlation coefficients; (b) estimated SNR from the maximum correlation coefficients; (c) standard deviation of the error in the time delay estimates; (d) standard deviation of the error in the dilation-factor estimates.

to the time delay estimate through the relation

$$\tau = \frac{\theta}{2\pi f_0}, \quad (4.28)$$

which is essentially the same as Equation 4.12. In principle, the time delay can be estimated for each sample-pair in the input data, but to reduce the noise we follow the SAR literature and introduce a maximum likelihood estimator over a local neighborhood.

The *complex* coherence (or zero-lag complex correlation coefficient) is given by

$$\gamma_{f_1 f_2}(t) = \frac{\mathbb{E}\{f_1(t)f_2^*(t)\}}{\sqrt{\mathbb{E}\{|f_1(t)|^2\}\mathbb{E}\{|f_2^*(t)|^2\}}}, \quad (4.29)$$

where $f_1(t)$ and $f_2(t)$ are the two interferometric signals. Assuming the signals are wide-sense stationary over the subset, the time dependence disappears. An estimate on discrete signals, $a[i]$ and $b[i]$, is found by a weighted sum of the complex phase-differences

$$\gamma = \frac{\sum_{i=1}^N a[i]b^*[i]}{\sqrt{\sum_{i=1}^N |a[i]|^2 |b^*[i]|^2}}. \quad (4.30)$$

Rodriguez and Martin (1992) showed that the argument of γ is the maximum likelihood estimate of the phase-difference $\theta = \arg\{\gamma\}$. This estimate is unbiased modulo 2π . We also see that $\theta = \arg\{R_{ab}[j=0]\}$, where $R_{ab}[j]$ is defined in Equation 4.8. We define the *interferogram* to be a matrix of θ -estimates (some refer to the complex coherence as the interferogram and θ as the interferometric phase, other use complex interferogram instead of complex coherence). The magnitude of the complex coherence is the coherence, $k = |\gamma|$. In this thesis we usually do not distinguish between this coherence and the coherence estimated from the cross-correlation function (see Equation 4.10), but we emphasize the difference where needed.

Since the interferogram is the zero-lag samples in the cross-correlation functions, the data are assumed co-registered. If the data are co-registered within one wavelength, the coherence is usually close to the maximum value given by the SNR, and the phase estimate is unambiguous. If the time delay is larger than one wavelength, the coherence starts to drop (in Figure 4.4, this corresponds to moving the peak away from the origin and using the sample at origin), and the phase will wrap modulus 2π . If the time delay is larger than the peak-to-null distance of the correlation function, the phase will jump π and the coherence will have drop close to zero (see Figure 4.4). In this case the estimate will be rendered useless.

The probability density function of θ , given the true phase-difference, θ_0 , can be calculated by assuming that the imaginary and real parts of the data are uncorrelated random zero mean distributions with equal variance. This means that the magnitude of the signal is Rayleigh distributed, and the phase uniformly distributed (Hanssen, 2001, pages 89-91). The PDF for a single sample (single look) is given by Hanssen (2001, pages 94) and Just and Bamler (1994)

$$p(\theta; k) = \frac{1 - k^2}{2\pi} \frac{1}{1 - k^2 \cos^2(\theta - \theta_0)} \left(1 + \frac{k \cos(\theta - \theta_0) \cos^{-1}(-k \cos(\theta - \theta_0))}{\sqrt{1 - k^2 \cos^2(\theta - \theta_0)}} \right). \quad (4.31)$$

Figure 4.10 shows the PDF for a few different coherence values. The figure also shows the estimated SNRs using Equation 4.17. It is evident that a high coherence (i.e. a high SNR) is crucial for an accurate phase-difference estimate.




Figure 4.10: PDF of the phase-difference, θ for $k = 0.9$ (solid line), $k = 0.8$ (dashed line), $k = 0.66$ (dash-dotted line) and $k = 0.25$ (dotted line). The equivalent SNRs are 9.5 dB, 6.0 dB, 2.9 dB and -4.8 dB.

The standard-deviation of the single-look phase-difference estimate, $\hat{\theta}$ is

$$\text{std} \{ \hat{\theta} \} \geq \sigma_{\theta}(k) \equiv \sqrt{\int (\theta - \theta_0)^2 p(\theta; k) d\theta}. \quad (4.32)$$

The solution to this equation was found by Bamler and Hartl (1998)

$$\sigma_{\theta}^2(k) = \frac{\pi^2}{3} - \pi \arcsin(k) + \arcsin^2(k) - \frac{\text{Li}_2(k^2)}{2}, \quad (4.33)$$

where Li_2 is the Euler's dilogarithm defined as

$$\text{Li}_2(k^2) = \sum_{n=1}^{\infty} \frac{k^{2n}}{n^2}. \quad (4.34)$$

Figure 4.11 shows the standard-deviation as a function of coherence. Clearly, the coherence needed to achieve $\sigma_{\theta}(k) \leq 10$ degrees is difficult to obtain. The standard deviation is, however, reduced by averaging over N independent samples. Often, it is assumed that the standard deviation is proportional to $N^{-1/2}$. Oliver and Quegan (2004, page 344) and Llorc-Pujol et al. (2006) present the marginal PDF for the phase-difference when es-




Figure 4.11: Standard deviation of the phase-difference as a function of $1 - k$, where k is the coherence.

estimated using a filter window of N independent samples

$$\begin{aligned} \bar{p}(\theta; k, N) = & \frac{(1 - k^2)^N}{2\pi} \left(\frac{(2N - 2)!}{((N - 1)!)^2 2^{2(N-1)}} \left(\frac{(2N - 1)K}{(1 - K^2)^{N+\frac{1}{2}}} \cos^{-1}(-K) + \frac{1}{(1 - K^2)^N} \right) \right. \\ & \left. + \frac{1}{2(N - 1)} \sum_{r=0}^{N-2} \frac{\Gamma(N - \frac{1}{2})}{\Gamma(N - \frac{1}{2} - r)} \frac{\Gamma(N - 1 - r)}{\Gamma(N - 1)} \frac{1 + (2r + 1)K^2}{(1 - K^2)^{r+2}} \right), \end{aligned} \quad (4.35)$$

where $K = \cos(\theta - \theta_0)$ and Γ is the Gamma function (Rottmann, 1995). The standard-deviation becomes a function of coherence and number of independent samples.

$$\text{std} \left\{ \hat{\theta} \right\} \geq \bar{\sigma}_\theta(k, N) \equiv \sqrt{\int (\theta - \theta_0)^2 p(\theta; k, N) d\theta}. \quad (4.36)$$

Figure 4.12 shows the standard deviation as a function of coherence for a few filter-sizes. Clearly, $\bar{\sigma}_\theta(k, N)$ does not scale as $N^{-1/2}$. We define a *scaling factor*, S as

$$S = S(k, N) = -\frac{\log(\sigma_\theta(k)) - \log(\bar{\sigma}_\theta(k, N))}{\log(10)}, \quad (4.37)$$

in such a way that $\bar{\sigma}_\theta(k, N) = N^{-S} \sigma_\theta(k)$. Notice that $\bar{\sigma}_\theta(k, 1) = \sigma_\theta(k)$. For a Gaussian PDF, S is equal to 0.5. Figure 4.13 shows S as a function of k and N . In the useful region ($k > 0.66$) the reduction in standard-deviation as N is increased is larger than for




Figure 4.12: Standard deviation of the phase-difference as a function of coherence, k , for filtering over different independent number of samples, N .




Figure 4.13: Scaling factor, S , as a function of coherence, k , for filtering over different independent number of samples, N .

Gaussian distributions. For $k = 0.9$ and $N = 81$, $S = 0.66$ and the standard-deviation is reduced by a factor $81^{0.66} = 18.2$. The corresponding factor is $\sqrt{81} = 9$ for a Gaussian distribution.

By converting phase-differences to time delays and noticing that $N = BT$, we can




Figure 4.14: Comparison of σ_c (solid line), σ_f (dashed line) and σ_i (dash-dotted line) as a function of SNR, ρ . The dotted line shows σ_i under the approximation that the scaling factor $S = 0.5$. The left panel shows the case for $N = 5$ and the right panel for $N = 81$. In the right panel, σ_i is indistinguishable from σ_f .

compare the three expressions for the standard deviation of the time delay estimate directly

$$\sigma_c(\rho, N, B) = \frac{\sqrt{3}}{\pi} \frac{1}{B} \frac{1}{\sqrt{N}} \sqrt{\frac{1}{\rho} + \frac{1}{2\rho^2}}, \quad (4.38)$$

$$\sigma_f(\rho, N, f_0) \approx \frac{1}{2\pi f_0} \frac{1}{\sqrt{N}} \sqrt{\frac{1}{\rho} + \frac{1}{2\rho^2}}, \quad (4.39)$$

$$\sigma_i(\rho, N, f_0) = \frac{1}{2\pi f_0} \frac{1}{N^S} \sqrt{\int (\theta - \theta_0)^2 p(\theta; k(\rho)) d\theta}, \quad (4.40)$$

where σ_i is the standard deviation of the time delay estimate based on the interferogram. Figure 4.14 compares the three quantities as a function of ρ using $B = 30$ kHz and $f_0 = 100$ kHz. For $N = 5$, σ_c differs from σ_i even though the same data and almost the same method is considered. The difference lies in the underlying *assumption*. In Section 4.1.3, σ_c and σ_f were based on constant deterministic signals in additive white Gaussian random noise. In this section σ_i were based on a signals with Rayleigh distributed magnitude and uniformly distributed phase (Hanssen (2001, pages 89-91); Just and Bamler (1994)). According to Hanssen (2001, page 95) the CRLB is a better model in the presence of point scatterers while the marginal probability density function is best suited in the presence of speckle – the case for a general seafloor. As Figure 4.14 shows, σ_f is too optimistic in speckle statistics, while σ_i is too pessimistic on point scatterers. As we see in the right panel of Figure 4.14, the difference between σ_c and σ_f becomes negligible for large estimation windows.

4.6 Model errors

In the previous sections we have considered signals in uncorrelated additive white noise. In general there are several other types of noise which limit the accuracy of time delay estimation. Lurton (2002, page 103) groups the noise contribution into four

- Ambient noise – noise which originates from outside the system. It can be natural (e.g. waves or animals) or man-made (e.g. shipping).
- Self-noise – noise which originates inside the system. It can be caused by the supporting platform (e.g. electrical interference) or by its own electronics (thermal noise).
- Reverberation – unwanted echoes generated by the sonar's own signals.
- Acoustic interference – noise generated by other acoustic platforms in the vicinity.

Instead of grouping noise sources as Lurton (2002, page 103) does, we will try to give a brief mathematical description of different signal models. We start by reviewing signals in uncorrelated additive white noise and then expand to more complicated scenarios.

4.6.1 Uncorrelated additive white noise

We assume that we receive two time series of a signal in uncorrelated additive white noise

$$f_1(t) = s(t) + n_1(t), \quad (4.41a)$$

$$f_2(t) = s(t - \delta t) + n_2(t), \quad (4.41b)$$

The signals are time-shifted with δt and the SNR is the same in both time series. Uncorrelated additive white noise is always present in measurement, but the SNR can vary significantly. The simplest case is additive white *Gaussian* noise, which is assumed when establishing the CRLB of the time delay estimate (see Equation 4.18).

4.6.2 Baseline decorrelation

In general the signal-part of the received time series is not identical in the two time series. Increasing the interferometric baseline increases this difference

$$f_1(t) = s_1(t) + n_1(t), \quad (4.42a)$$

$$f_2(t) = s_2(t - \delta t) + n_2(t). \quad (4.42b)$$

True baseline decorrelation is mismatched spectra (Lurton (2000); Hanssen (2001, page 102)). This effect is irreversible and will lower the coherence and add noise to the estimate unless the spectra are filtered (which will reduce the bandwidth). Geometrical decorrelation (Gatelli et al., 1994) is simply a dilation of the signals inside the correlation window. As we showed in Section 4.3, the cross-uncertainty function is a better estimator than the cross-correlation function in the presence of baseline decorrelation.

4.6.3 Multipath

In a reverberant environment the received signal is not a pure replica of the transmitted signal. Reflections from medium boundaries and the roughness of the boundaries will spread the signal in time (Brekhovskikh and Lysanov, 1982, chapter 9). If the boundary is moving (like the sea surface), a Doppler shift will be added to the scattered field and the signal spread in frequency (Brekhovskikh and Lysanov, 1982, chapter 9). These effects can be incorporated in a transfer function, $h(\delta t)$

$$f_1(t) = h(\delta t) \odot s(t) + n_1(t), \quad (4.43a)$$

$$f_2(t) = h(\delta t) \odot s(t - \delta t) + n_2(t), \quad (4.43b)$$

where \odot is the convolution operator. To resolve these ambiguities one can for example beamform the data in a given direction, use a parametric method to locate the echo or use a generalized cross-correlation function with a filter in frequency domain (Knapp and Carter, 1976). Time spread and frequency spread is treated in detail in wireless communications (Benedetto and Biglieri, 1999, chapter 13).

4.6.4 Correlated noise

Uncorrelated noise and reverberant noise are both environment induced. Another noise-source is system self noise that may be correlated over the receiver array (Lurton, 2002, pages 116-117). This noise can be coherent and synchronized in the time series, especially if the receivers have a common electrical system. The estimated time-shift is usually dominated by *either* the time-shifted signal, or the synchronized noise. The received time series are

$$f_1(t) = s(t) + c(t) + n_1(t), \quad (4.44a)$$

$$f_2(t) = s(t - \delta t) + c(t) + n_2(t), \quad (4.44b)$$

where $c(t)$ is the correlated noise. This noise (if it in fact is caused by the electrical system) can be estimated by running the system without transmitting a signal. It can then be removed from the time series, although this can be challenging in practice.

4.6.5 Multiplicative noise

The medium (seawater) can also induce multiplicative noise from random fluctuations (Brekhovskikh and Lysanov (1982, chapter 10); Flatte et al. (1979, chapter 8)). One way to model this is to introduce a random multiplicative component

$$f_1(t) = \alpha_1(t)s(t) + n_1(t), \quad (4.45a)$$

$$f_2(t) = \alpha_2(t)s(t - \delta t) + n_2(t). \quad (4.45b)$$

In general $\alpha_1(t)$ and $\alpha_2(t)$ can be complicated and highly different, decorrelating the received time series. If the two time series have passed through the same medium, $\alpha_1(t)$

and $\alpha_2(t)$ will be similar, but they can bias the time delay estimation. A common way to model propagation in weak fluctuating media is to divide into phase and amplitude perturbations. This can be obtained by using our model with complex valued α_1 and α_2 .

4.6.6 Dispersive scattering

Scattering from the seafloor may be frequency-dependent in several ways. Scattering from a rough seafloor is dependent of the roughness scale relative to the acoustic wavelength (Ogilvy (1991); Jackson and Richardson (2007, chapter 13)). Reflection from an absorbing seafloor is frequency-dependent (Lurton (2002, page 65); Jackson and Richardson (2007, pages 331-376)). If the frequency response is dependent on the incident angle on the object, the spectra of the two received signals may differ

$$f_1(t, f) = s_1(t, f) + n_1(t), \quad (4.46a)$$

$$f_2(t, f) = s_2(t - \delta t, f) + n_2(t). \quad (4.46b)$$

For seafloor depth estimation, the incident angles are usually very similar. This means that $s_2(t - \delta t, f) \approx s_1(t - \delta t, f)$. The dominant effect of the dispersive scattering is a change of the effective center frequency of the signals (see Section 4.1.4).

4.6.7 Frequency-dependent noise

Acoustic or electrical noise can be frequency-dependent. Ambient noise is frequency-dependent (Urick, 1983, chapter 7). Self noise or acoustic interference may be very frequency-dependent (Lurton, 2002, page 115-117).

$$f_1(t, f) = s(t) + u_1(t, f) + n_1(t), \quad (4.47a)$$

$$f_2(t, f) = s(t - \delta t) + u_2(t, f) + n_2(t), \quad (4.47b)$$

where $u_1(t, f)$ and $u_2(t, f)$ are frequency-dependent noises. They may or may not, be similar depending on the location of the source relative to the sonar receivers. Narrow-band noise can be removed by filtering in the frequency domain (Mitra, 2002, chapter 7) at the expense of a reduces SNR.

4.6.8 Phase ambiguities

In the previous sections we have discussed different signal models. We saw that in the general case the interferometric signals may differ in both time and frequency, and the noises may be correlated and complicated. Now we return to a problem which occurs even for time delayed signals in noise free environment: Phase ambiguities. When we consider two complex signals, it is the phase-difference which provides accurate information of the time delay between them. Think of a continuous sine-wave: A shift can

easily be detected by matching the oscillations. However, there are an infinite number of periods which match. The same apply to broadband sonar signals. The true phase-difference, θ_0 is

$$\theta_0 = \theta + 2\pi n, \quad (4.48)$$

where θ is the estimated phase-difference and n is an integer. Phase unwrapping is a large and complicated research topic (Ghiglia and Pritt, 1998) and outside the scope of this thesis, so we have tried to keep the effort on the subject to a minimum. However, some work has been invested into simple and standard approaches.

Phase-ambiguities using the cross-correlation function

Phase-ambiguities may occur in a complex cross-correlation based time delay estimate in Equation 4.12. In Section 4.1.1 we described how the *probability* for a phase wrap is minimized by using bandwidth to estimate a coarse time delay. But since the 3 dB width of the coarse estimate is $1/B$, while the wrap interval is $1/f_0$, a very high SNR is needed to eliminate all wrap-errors. To resolve any residual 2π -errors we have developed the following *ad-hoc* approach

- We use the coarse estimate to minimize the probability of phase ambiguities (as described in Section 4.1.1).
- We discard all estimates with a coherence below 0.66. Since the probability of a phase-wrap is highest at low SNR (corresponding to low coherence) this further reduces the probability of a phase ambiguity.
- We assume that there are more wrap-free than wrapped estimates and that the probability of a large number of consecutive wrapped estimates is low. We have used these assumptions to develop a new, but simple algorithm which unwraps the estimates in a small neighborhood around the sample of interest. This algorithm is described in Section 5.1.3.

The advantage with this method is that it can not corrupt large amount of data in the same manner as a standard 2D method. Any remaining 2π -errors can simply be considered wild-points in the estimation algorithm.

2D phase unwrappers

Since the interferometric phase-difference for a typical sidelooking sonar geometry is multiple wavelengths (see Figure 3.3) Equation 4.29 provides an estimate which can be in almost any 2π -interval. The problem is largely reduced by the *a priori* ground-range mapping of the SAS images. The images are rendered in ground-range such that the phase-difference is zero if the *a priori* depth is correct. If we assume that the *a priori* depth estimate is unbiased with Gaussian probability, the probability of wrap-errors is also Gaussian, centered at the correct wrap-interval. In practice, this is not the case. Most

of the samples are, in fact, wrap-free, but due to a slowly varying *a priori* depth, sharp elevated features will have wrap-errors. This means that samples with wrap-errors are grouped together, but they should still be in minority. A standard 2D unwrapper works very well in such cases.

To solve the residual phase-wraps we have implemented a standard Goldstein's Branch cut algorithm as described in Ghiglia and Pritt (1998). We have not tested this method against other established methods, nor have we performed any thorough study of the performance. However, the method works well in most cases. The exception is in the most challenging scenarios, which typically is large, complicated man-made objects with vertically features. In such cases we have performed a preliminary study into using bandwidth to resolve the phase ambiguities (see Section 8.3). We believe that ultra wideband signals reduce the need for advanced 2D unwrappers. Notice the similarity between such a method and the cross-correlation based method described in the previous section.

4.7 Coherence estimation

In this chapter, as well as in most literature on interferometry, the focus has been on the time delay estimate. The coherence estimate can be almost equally important (Synnes et al., 2009; Midtgaard et al., 2009). In this section we give a short review of the statistics of the coherence estimate.

The coherence can be defined as the magnitude of the complex coherence (see Equation 4.29). Ideally, the estimate should be obtained using a very large number of observations for each pixel-pair in the SAS images. In practical situations, each pixel-pair is observed only once and we assume ergodicity to exchange ensemble averages with spatial averages. This assumption leads to a maximum likelihood estimator over N samples

$$k = \left| \frac{\sum_{i=1}^N a[i]b^*[i]}{\sqrt{\sum_{i=1}^N |a[i]|^2 |b^*[i]|^2}} \right|. \quad (4.49)$$

This estimator is clearly not unbiased. Two signals of uncorrelated noise should have $k = 0$, but Equation 4.49 will in this case result in a sum over vectors with random directions, which clearly have an expectancy > 0 unless $N \rightarrow \infty$. However, for completely coherent signals in zero noise, Equation 4.49 will give out the expected answer of one. It therefore appears that the estimator is biased towards higher values for low coherence and small estimation windows (Joughin and Winebrenner, 1994; Touzi et al., 1999).

Touzi and Lopes (1996) found that for $N > 2$, the probability density function of the coherence estimator can be expressed as a function of the true coherence, k_0 and the

number of independent samples, N

$$p(k; k_0, N) = 2(N - 1)(1 - k^2)^N k_0 (1 - k_0^2)^{N-2} {}_2F_1(N, N, 1, k^2 k_0^2), \quad (4.50)$$

where the hypergeometric function is defined as (Hanssen, 2001, page 94)

$${}_2F_1(N, N, 1, k^2 k_0^2) = \sum_{i=0}^{\infty} \frac{(N)_i (N)_i (k^2 k_0^2)^i}{(1)_i i!}, \quad (4.51)$$

and the Pochhammer symbol is defined as

$$(N)_i = N(N + 1)(N + 2) \cdots (N + i - 1), \quad (N)_0 = 1. \quad (4.52)$$

The expectation of k is derived as (Hanssen, 2001, page 96)

$$E\{k; N\} = \frac{\Gamma(N)\Gamma(3/2)}{\Gamma(N + 1/2)} {}_3F_2(3/2, N, N, N + 1/2, 1, k^2) (1 - k^2)^N, \quad (4.53)$$

where ${}_3F_2(3/2, N, N, N + 1/2, 1, k^2)$ is the generalized hypergeometric function.

Figure 4.15 shows the bias in the coherence estimate as a function of true coherence, for different number of independent samples. We discard all time delay estimates with coherence less than 0.66. For $N = 81$ and an estimated coherence of 0.66, the true coherence is 0.63. This is only an error of 1 dB. For higher coherences, the bias gets smaller. However, we also try to estimate the coherence for samples where the time delay estimate is useless. In Synnes et al. (2009) and Midtgaard et al. (2009) we present two applications for sonar coherence mapping. Clearly, the bias can be important for such topics. It has been suggested to correct the biases coherence estimate, since the bias is analytically known (Touzi et al., 1999). Hanssen (2001, pages 96-98) shows that at low coherence, the standard deviation of the coherence estimate is high. This means that a correction of the coherence can lead to unphysical values (e.g. the coherence can be greater than unity).




Figure 4.15: The expectation value of k , as a function of the true coherence, k_0 for a number of different samples in the estimator, N . The solid line shows the bias free estimate.

Chapter 5

Algorithms for depth estimation

In this chapter we describe how the methods from previous chapters are implemented in an interferometric SAS processing chain. Estimation and use of relative seafloor depth can be found at several levels in SAS processing, each step being crucial for the steps which follow. Thus, in order to make a high quality SAS bathymetry, all of the preceding calculations have to be implemented as accurately as possible.

As mentioned in Chapter 1, a large portion of this thesis has consisted of contributing to the development of the FOCUS toolbox (Hansen et al., 2005). FOCUS is a SAS processing toolbox which takes raw sonar (or radar) data as input and generates a wide variety of outputs. Except for a few 1D simulations, all results presented in this thesis are generated using the FOCUS toolbox. Figure 5.1 shows a flowchart of the main processing chain in FOCUS. Processing steps involving relative depth estimation or use of relative depth estimates, are colored red in Figure 5.1. These can be grouped into three categories:

- Ping-by-ping sidescan seafloor depth estimation (sidescan bathymetry)
- Use of estimated sidescan bathymetry in SAS processing
- SAS interferometry

5.1 Sidescan seafloor depth estimation

In this section we describe the implementation of relative depth estimation using single beam sidescan images. The sonar geometry can be found in Figure 3.1. Each of the receivers in Figure 3.1 may represent a receiver *array* perpendicular to the page, but when the distinction is not important, we refer to them as receivers. Figure 5.2 shows an overview of the relative seafloor depth estimation algorithm.

The first step of sidescan bathymetry is to generate sidescan images. For both interferometric receiver arrays, each ping is dynamically focused (Soumekh, 1994, pages

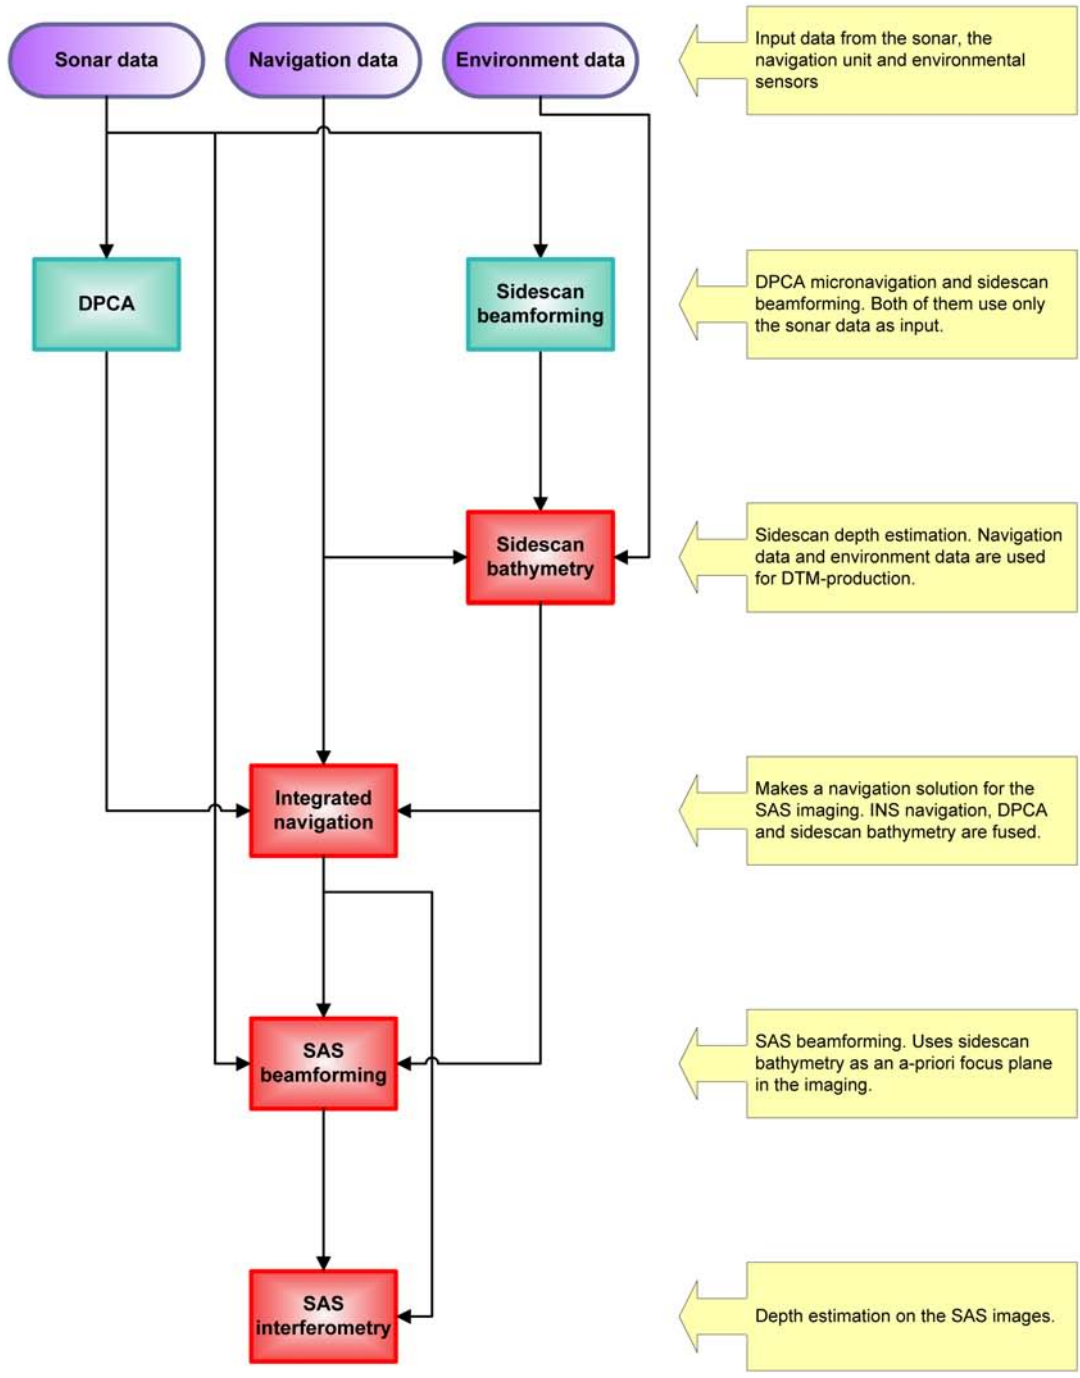


Figure 5.1: Flowchart of the main processing chain in the FOCUS SAS processing toolbox.

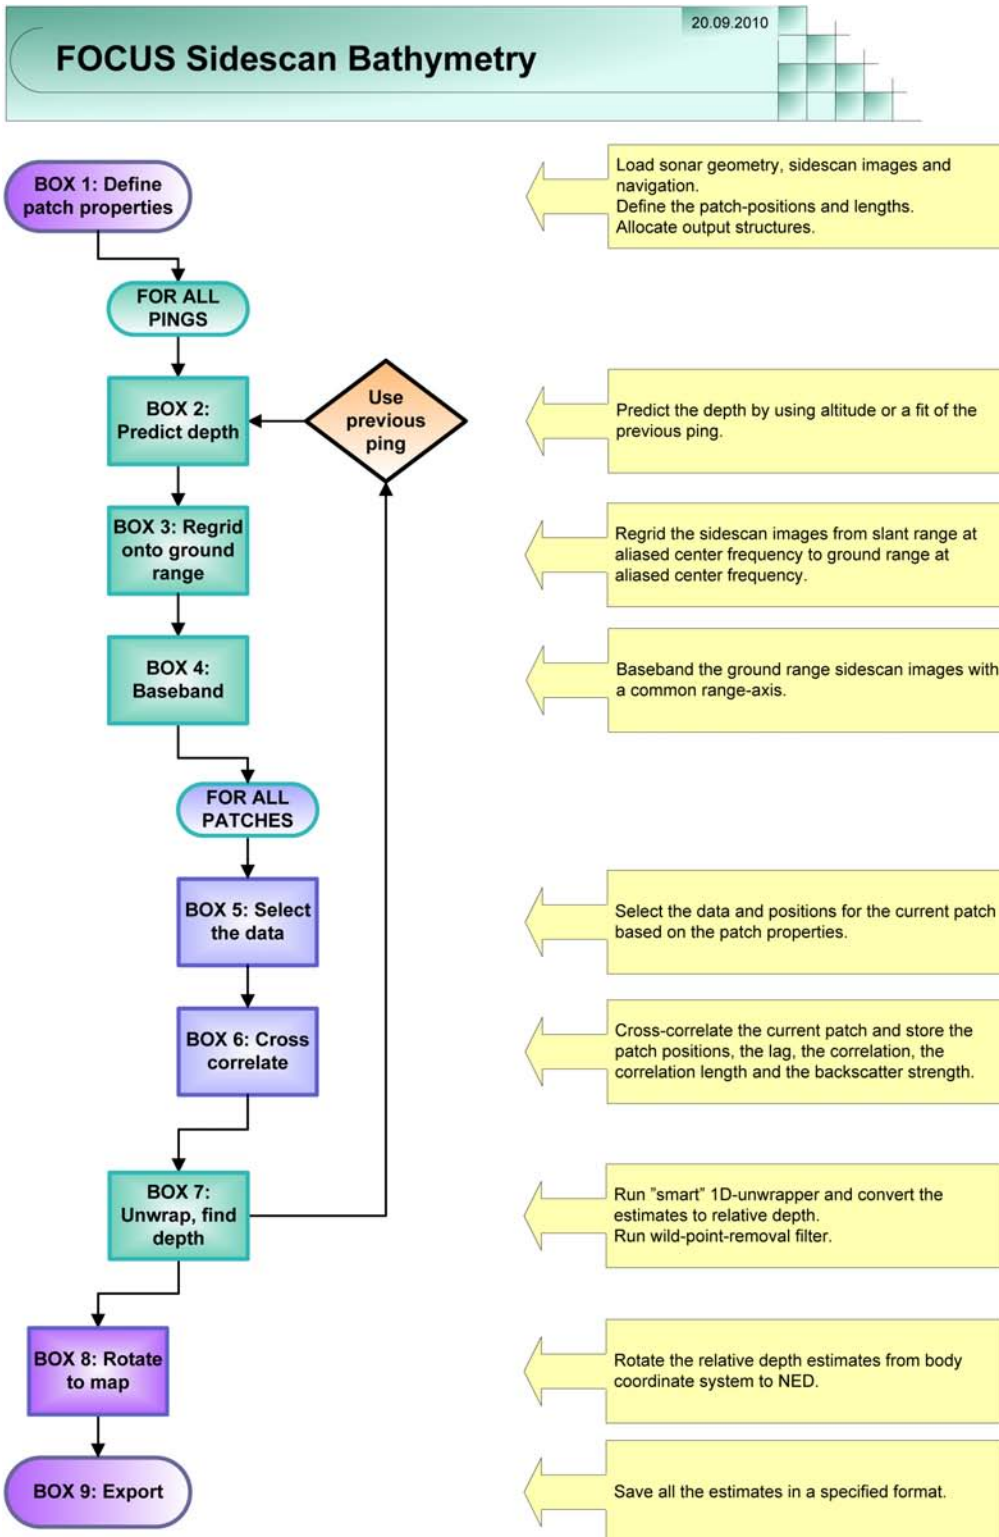


Figure 5.2: Flowchart of the processing steps used in relative depth estimation on sidescan images.

196-209) in slant-range, forming a beam perpendicular to the transmitter-receiver baseline. The beams from the two receivers are not in general parallel, this only occurs when the transmitter is positioned on a line between the two receivers. The origin of the beams can be the receivers, the PCA locations or the transmitter. The choice of origin is not important, as long as the preferred choice is carried properly through the processing. We have chosen PCA locations as origin.

The beamforming results in two sidescan images, $I_1^{\text{sr}}[p, n]$ and $I_2^{\text{sr}}[p, n]$, where p is the ping counter and n the sample counter along range. The subscript indicates receiver # and the superscript, sr indicates that the images are in slant-range. These images, together with the the sonar geometry and the sample distance, dr , are the main inputs to the depth estimation routine.

5.1.1 Re-gridding to ground-range

As we described in Section 3.1.3, we co-registrate the sidescan images by re-gridding them onto an *a priori* seafloor. The seafloor is calculated using the relative depth estimates from the previous ping. For the first ping we *kick-start* the processing using a flat seafloor at the altimeter depth, and then re-estimate the ping using the estimated depth as *a priori* information. The same kick-starting is performed when the estimates from the previous ping are rendered invalid due to low coherence.

Let the two-way range from the transmitter to each sample in the sidescan images and back into the receiver be

$$r^{\text{sr}}[n] = 2\sqrt{r^2[n] + \left(\frac{D}{2}\right)^2}, \quad (5.1)$$

where $r[n] = n \cdot dr$ and D is the interferometric baseline between the transmitter and receiver. This equation describes the dynamic focusing onto the PCA positions.

The range axes from the transmitter to the *a priori* seafloor and back to the receivers, are given by

$$r_1^{\text{gr}}[m] = \sqrt{y^2[m] + z^2[m]} + \sqrt{(y[m] - y_1)^2 + (z[m] - z_1)^2}, \quad (5.2a)$$

$$r_2^{\text{gr}}[m] = \sqrt{y^2[m] + z^2[m]} + \sqrt{(y[m] - y_2)^2 + (z[m] - z_2)^2}. \quad (5.2b)$$

Here $y[m] = m \cdot dy$ is a vector defined by a input parameter, dy , and $z[m]$ is the *a priori* depth corresponding to the positions $y[m]$. Figure 3.1 defines the receiver positions, y_1, y_2, z_1, z_2 . The sidescan images are critically sampled in slant-range so we have to be careful to avoid undersampling in the new coordinate frame. For a flat *a priori* seafloor, $dy = dr$ is sufficient, since the sampling frequency is dilated by $\cos \theta$ in ground-range relative to slant-range (see Equation 3.12). For a general $z[m]$, $dy = dr$ may be insufficient. By assuming a slowly varying *a priori* seafloor and decreasing dy by a small factor, undersampling is avoided for most practical geometries.

The re-gridding from slant-range to ground-range is implemented using a spline interpolation (Kincaid and Cheney, 1991, pages 374-392). The sidescan images in slant-range are at aliased center frequency f_0 . We mix the slant-range data to baseband, interpolate and mix the data up using the new ground-range sampling

$$I_1^{\text{gr}}[p, m] = \mathcal{I} \left\{ I_1^{\text{sr}}[p, n] \exp \left(\frac{-j2\pi f_0}{c} r^{\text{sr}}[n] \right) : r^{\text{sr}}[n] \Rightarrow r_1^{\text{gr}}[m] \right\} \exp \left(\frac{j2\pi f_0}{c} r_1^{\text{gr}}[m] \right), \quad (5.3a)$$

$$I_2^{\text{gr}}[p, m] = \mathcal{I} \left\{ I_2^{\text{sr}}[p, n] \exp \left(\frac{-j2\pi f_0}{c} r^{\text{sr}}[n] \right) : r^{\text{sr}}[n] \Rightarrow r_2^{\text{gr}}[m] \right\} \exp \left(\frac{j2\pi f_0}{c} r_2^{\text{gr}}[m] \right), \quad (5.3b)$$

where $\mathcal{I}\{A[x] : x \Rightarrow y\}$ is an interpolation operator which interpolates $A[x]$ from x onto y , and j is the imaginary unit. The ground-range sidescan images are now at aliased center frequency, f_0 , with non-linear range-axes (see Section 3.1.2).

We want to minimize interpolation errors in the subsequent processing, so we baseband the ground-range images (Hawkins, 1996, chapter 4). We have to decide on a common range-axis for the basebanding and the natural choice is the average of $r_1^{\text{gr}}[m]$ and $r_2^{\text{gr}}[m]$.

$$r^{\text{gr}}[m] = \frac{1}{2}(r_1^{\text{gr}}[m] + r_2^{\text{gr}}[m]). \quad (5.4)$$

The basebanded ground-range sidescan images are given by

$$I_1^{\text{gr,b}}[p, m] = I_1^{\text{gr}}[p, m] \exp \left(\frac{-j2\pi f_0}{c} r^{\text{gr}}[m] \right), \quad (5.5a)$$

$$I_2^{\text{gr,b}}[p, m] = I_2^{\text{gr}}[p, m] \exp \left(\frac{-j2\pi f_0}{c} r^{\text{gr}}[m] \right). \quad (5.5b)$$

5.1.2 Cross-correlation of patches

For each ping of the ground-range sidescan images, a set of patches along range is selected. The locations of the patches and the number of samples in each patch is independent parameters, and can be varied throughout the ping. Our default setting is one meter between each patch in ground-range, and a fixed patch-length of 3.2 meters. This gives us a regular gridding of the estimates in an earth-fixed coordinate system and $128 \cos \Phi$ individual samples in each correlation (for HISAS with 2.5 cm range resolution), where Φ is the depression angle. The number of individual samples is largest at far range, due to the scaling between ground-range and slant-range. It can be beneficial to increase the patch length at far range since baseline decorrelation is less of a problem and the SNR may be low.

For each of the selected patches we cross-correlate N complex samples from sample m_1 to sample m_2 (see Section 4.1)

$$R_{I_1 I_2}^{\text{gr,b}}[p, j] = K \sum_{m=m_1}^{m_2-m_1+1-j} I_1^{\text{gr,b}}[p, m+j] \left(I_2^{\text{gr,b}}[p, m] \right)^*, \quad j = 0, 1, \dots, N-1, \quad (5.6)$$

where K is the normalization factor defined in Equation 4.9. The peak correlation value is used to find the coarse pixel lag, j_c

$$j_c = \underset{j}{\operatorname{argmax}}\{|R_{I_1 I_2}[j]|\}, \quad (5.7)$$

where the superscripts and the ping dependence are omitted for simplicity. The fine delay is given by Equations 4.11, 4.12 and 4.13, but since the input data are sampled regularly in spatial ground-range coordinates instead of in time, the equations look different.

5.1.3 Unwrapping the sidescan bathymetry estimates

From the properties of the cross-correlation function of two time delayed signals, the phase should be zero at the true correlation peak. We have experienced from sidescan bathymetry estimates that the phase of the peak of the sampled cross-correlation function does not strictly behave in accordance with this property. There are a number of reasons why theory does not match experience

- The signals are sampled so the phase at the sampled peak differs from phase at the true peak, unless the sampling frequency is infinite.
- The signals have noise on both the magnitude and the phase, so the selected peak may be incorrect or the phase may be noisy.
- The signals are not only time delayed relative to each other, they are also dilated.

Figure 5.3 shows the phase distribution of the valid estimates for a simulated seafloor with a non-linear trajectory and a varying topography. It shows that the distributions are zero-mean, and that the spread of the phase distribution decreases when the cross-correlation function is upsampled (see Section 4.1.1). Note that if the data is adequately sampled one can upsample the cross-correlation function rather than the input sidescan images.

Figure 5.4 illustrates how the distribution is widened when the bathymetry of the seafloor becomes more challenging. This is due to an increased dilation and a less accurate *a priori* seafloor depth estimate. Comparison of the simulated data in Figure 5.3 with the experimental data in Figure 5.4 also reveals the widening of the phase distribution due to noise.

From Figures 5.3 and 5.4, and as mentioned in Section 4.1.1, 2π -wrap errors will occur even for cross-correlation based methods. These wrap-errors occurs rarely, but are randomly distributed throughout the sidescan image. A simple 1D unwrapper (Ghiglia and Pritt, 1998, pages 16-26) will successfully unwrap most of these samples, but in a few cases it can contribute to a large error: If the unwrapper is run across a shadow region or across pixels with low coherence, all subsequent samples for that ping may be wrapped out of the correct 2π -interval. Figure 5.5 shows an example where this has




Figure 5.3: Distribution of the phases of the peak of the cross-correlation functions in sidescan bathymetry, for 1000 simulated pings with 150 patches each. Pixels with a correlation value below 0.66 have been discarded. The bathymetry is slowly varying over the scene. The left panel shows the results when the sidescan data are critically sampled, while the right panel shows the results when the sidescan data are oversampled with a factor of ten in the re-gridding to ground-range. The solid lines indicate 90% of the data and the dashed line 99%.




Figure 5.4: Distribution of the phases of the peak of the cross-correlation functions in sidescan bathymetry, for 1000 pings with 200 patches each. The data are real data collected using HISAS. The left panel shows the results for an approximate flat seafloor, while the right panel shows the results for a seafloor with heavy topography. The solid lines indicate 90% of the data and the dashed line 99%.




Figure 5.5: Normalized lag, $\delta y/c/2f_0^{\text{gr}}$ for a ping of real data collected with HISAS. The patch numbers are increasing with range and are spaced with one meter separation in ground-range (e.g. patch 100 is at 100 meters ground-range). The raw estimates are showed with dots. The estimates unwrapped with a standard 1D unwrapper are shown with squares and the estimates unwrapped with our local 1D unwrapper are shown with circles. Notice that patch 148 and 149 originally seem to be in the wrong 2π -interval.

happened on real data using HISAS. All the estimates from patch 112 and upward are wrapped up to a normalized lag of approximately one. In this case, the correct wrap interval is closer to zero. The error is caused by the region of invalid samples from patch 102 to patch 111. This is a region with shadow in the sidescan images where the correlation values are below 0.66. The geometry changes through the shadow and it is not possible to know where to continue the unwrapping.

If we use phase differencing on the ground-range sidescan images instead of cross-correlations, the phases would vary smoothly with large spatial correlation between the samples. However, when we cross-correlate and select the peak of the correlation function, the phases will be zero-mean with no spatial correlation. In other words: a single sample will with high probability be in the correct 2π -interval and it will be uncorrelated to the neighboring samples. The wrapped samples will be few, and they will occur randomly in the data. We therefore define a new *local* unwrapper which unwraps each sample to a local median of valid pixels

$$\tilde{\delta y}[m] = \mathcal{U}\{\delta y[m], \mathcal{M}_u\{\delta y[m-u : m+u]\}\}, \quad (5.8)$$

where u is a small number of samples, $\mathcal{U}\{x, y\}$ unwraps the estimate x onto the same wrap-interval as y and $\mathcal{M}_u\{\delta y[m-u : m+u]\}$ is the median value of the u valid samples closest to the sample m in the interval $[m-u, m+u]$ (excluding m itself). This means for

example if all samples in $[m - u, m + u]$ is valid, median is performed on the interval $[m - u/2, m + u/2]$. If only the samples up to and including m is valid, median is performed on the interval $[m - u, m - 1]$. If there are fewer than u valid samples, fewer than u samples will be used in the median. This *ad-hoc* method has proven to be robust on large amounts of experimental data.

The unwrapping algorithm described above defines a small neighborhood in where the median wrap-interval is found, and then the estimate is wrapped into that interval. This reduces the probability of propagating large errors through the ping as it only unwraps small groups of samples. The circles in Figure 5.5 show the results after our local unwrapper; the region from patch 112 and upward is correctly unwrapped (in this case the samples are left as they were) and the two samples (148 and 149) which were wrapped are now unwrapped correctly. As mentioned before there should be no spatial correlation between samples, so the probability of two neighboring samples are wrapped as in Figure 5.5 is very low. However, this applies to *independent* samples. Since we use a 3.2 meter patch size and one meter patch separation, there is a spatial correlation up to four samples. That is why we chosen u to cover a larger neighborhood than the patch size.

5.1.4 Converting to relative depth

Using the geometry of the sonar, the phase unwrapped lags, $\tilde{\delta}y$, are converted to a *relative body depth* using Equations 3.17 and 3.20. This depth is in the body-frame of the platform, i.e. the receiver positions, y_1, y_2, z_1 and z_2 (see Figure 3.1) are the positions for zero roll (ϕ), pitch (β) and yaw (ψ).

The next step is to convert relative body depth estimates to relative depth estimates. We have chosen to only apply roll and pitch compensation. Since the yaw-rotation is the last rotation, it can be ignored at this stage. The relative depth estimates are stored in a coordinate system which changes from ping to ping and is defined as

- The z -axis points down toward the seafloor
- The y -axis is perpendicular to the heading of the platform at transmit time
- The x -axis is defined as orthogonal to the z - and y -axes such that the xyz -axes define a right-hand coordinate system.

Each estimate is represented by a xyz -vector from the transmitter position at transmit time, to the seafloor. In the relative body depth coordinate system this vector is \mathbf{v}_{rbd} . The x -component of this vector is by definition equal to zero. The relative depth after the pitch- and roll-compensation, \mathbf{v}_{rd} , is

$$\mathbf{v}_{\text{rd}} = \mathbf{R}_{\text{rbd} \rightarrow \text{rd}} \mathbf{v}_{\text{rbd}} \quad (5.9a)$$

$$\mathbf{R}_{\text{rbd} \rightarrow \text{rd}} = \mathbf{R}_{\text{rbd} \rightarrow \text{rd}}(\phi, \beta, \psi = 0), \quad (5.9b)$$

where ϕ and β are the roll and pitch of the platform at transmit time and $R_{\text{rbd} \rightarrow \text{rd}}(\phi, \beta, \psi)$ is the rotation matrix from body to the roll- and pitch-compensated coordinate system. These *broadside* depth estimates are a suitable intermediate data format since they can be used for ping based processing and easily converted into earth-fixed coordinate systems. Figure 5.6 shows an example of the ping-based relative depth estimates and a mosaic into a local earth-fixed coordinate system. Notice that the platform depths are added to the relative depths in the mosaic. Also notice that estimator at long range successfully estimates the seafloor depth even when it is above the sonar (the areas with negative depth in Figure 5.6).


In Figure 5.6(a) there are blank areas where data seem to be missing. These are areas where the correlation coefficients are less than the threshold. All sidescan bathymetry estimates in this thesis are presented with a threshold of 0.66 (approximately equal to a SNR of 3 dB). Invalid samples typically appears in the following regions:

- In close-range regions the baseline decorrelation effect can reduce the correlation coefficients
- In far-range regions the SNR may drop due to lack of signal energy
- In shadow regions caused by topography, there is no signal energy
- In areas with interference and multipath the correlation may drop (although sometimes the interference can be coherent between the interferometric receivers)


5.1.5 Sound speed correction

One of the effects which makes sidelooking depth estimation more challenging than downlooking depth estimation (e.g. single beam echo sounding and multibeam echo sounding (Lurton, 2002, chapter 8)) is the effect of a vertically varying sound speed profile. A signal traveling through a medium with a varying sound speed will not follow a straight path (Brekhovskikh and Lysanov, 1982, Pages 33-35). Only directly vertical rays will be straight. The degree of ray bending increases with how close you get to horizontal, so a sonar which looks 200 meters to the side in 20 meters water depth will be much more susceptible to this effect than a sonar which only sees 45 degrees and steeper. Sound velocity correction is common to do for high precision mapping sonars (Lurton, 2002, pages 278-281)

The effect of ray-bending is that the position and depth of the estimate both can be incorrect. What we really measure with the interferometer is the incident angle of the reflection into the sonar, along with the travel time. These two values can be found from the estimated depths and are sufficient to correct for ray-bending. We have chosen to use a standard constant gradient raytracer (Brekhovskikh and Lysanov (1982, Pages 37-38); Synnes (2008)) as a correction algorithm. The principle is to divide the sound speed profile into a number of segments. Within each segment a constant gradient is assumed.



(a) Relative depth estimates



(b) Mosaic of depth estimates in a earth-fixed coordinate system

Figure 5.6: In the upper panel we show the relative depth estimates for around 2000 sidescan pings, where we have used 200 patches with one meter separation. The recorded range of the sonar varies between 150 and 200 meters. Areas with no coverage or with a correlation value below 0.66 is colored white. In the lower panel we show the same data in a earth-fixed coordinate system where the data are mosaicked into a half-meter grid. The track of the AUV is showed as a black line.

The ray within such a segment will then follow a circular arc (Ziomek, 1985, Pages 220-240). The algorithm consists of estimating an incident angle from the uncompensated depth estimates

$$\hat{\Phi}(r) = \tan^{-1} \left\{ z(r), \sqrt{r^2 - z^2(r)} \right\}, \quad (5.10)$$

where z is the uncompensated depth estimate and r the range. A ray is then sent out from the sonar in the $\hat{\Phi}(r)$ direction and traced through the different segments until it has traveled a length r . This results in a new position and depth on the seafloor.

Figure 5.7(a) shows an example dataset where we have compared depth estimates before and after sound speed correction. Before correction there is a pronounced trend in the data as a function of range. The seafloor depth seems to bend downward on both sides of the sonar. This is a typical result of uncompensated sound speed variations. In the middle panel we have removed the artificial trend by sound speed correcting the estimates. There is a difference of up to five meters, which is significantly larger than the precision of the depth estimates. Figure 5.7(b) shows that the sound speed profile for this dataset has a relatively strong variation.

5.1.6 Exporting sidescan bathymetry digital terrain maps

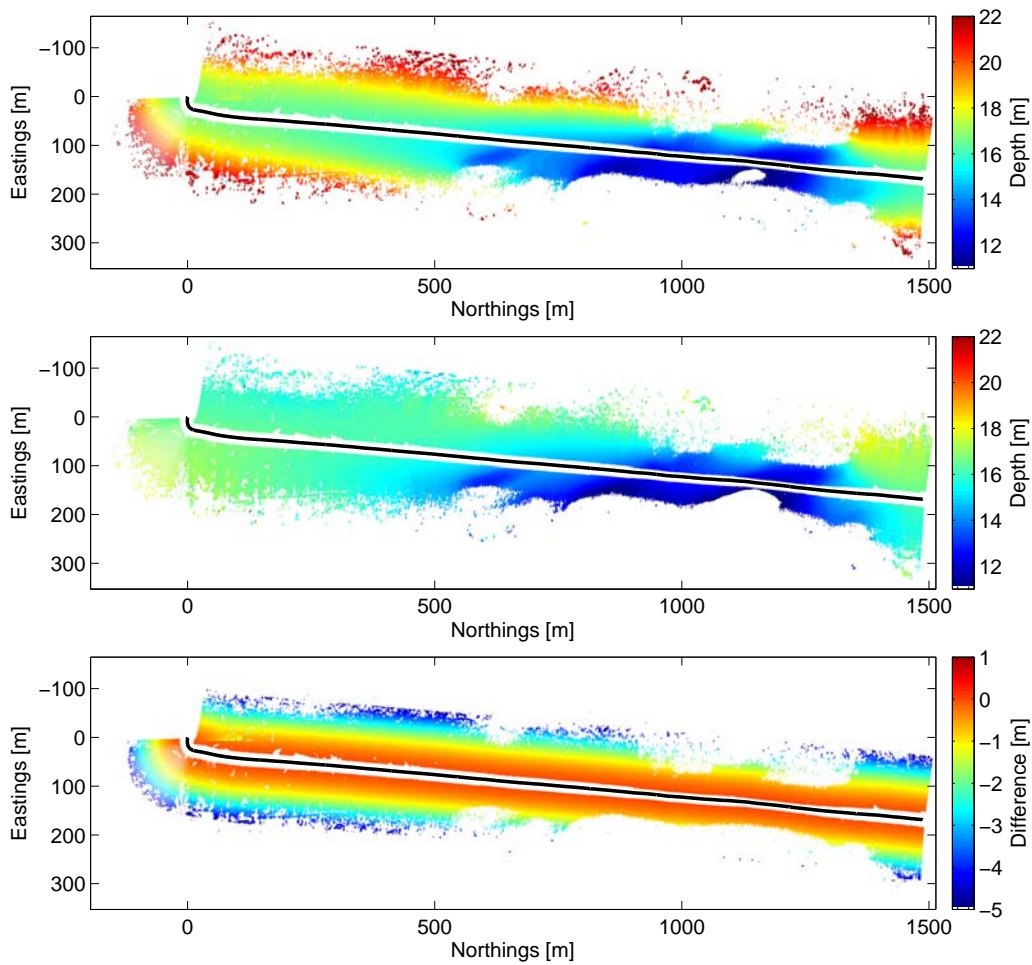
The sidescan bathymetry estimates are a vital component in the SAS processing and have to be easily accessible. In addition, the estimates are an excellent product for fast, low resolution assessment of the seafloor depth. We have therefore implemented three export formats

- MATLAB .mat format (Mat, 2010) for internal use in the processing
- Kongsberg .all format (Tri, 2008) for viewing with Kongsberg products
- Triton .xtf format (EMD, 2010) for viewing with Triton software


The exporting to .all and .xtf is performed in a stand-alone routine which reads the internal format and converts it to appropriate form. This makes it easy to extend to other formats if needed.

The details of the exporting is fairly straightforward, but there are a few remarks: Navigation information has to be provided along with the lever-arm from the navigation reference location to the transmitter. Since there is one sonar on each side there are two separate lever-arms and it is convenient to store the estimates from each side in its own file. The xyz -estimates, the reflectivity and the coherence should all be stored in the file to provide a quality estimate with the measurement.

In order to produce a proper digital terrain map (DTM) one should correct for environmental effects like tide and surface waves (Ambrose and Geneva (1995); de Jong et al. (2002, pages 12-39)). This should be done in an external software. The same goes for the geo-referencing of the data. A multitude of post-processing softwares designed for these purposes exist already, but it is outside the scope of this thesis to evaluate them.



(a) Relative depth estimates and differences



(b) Sound speed profile

Figure 5.7: In the upper panel we show the mosaicked depth estimates (in a half-meter grid) for around 3300 sidescan pings, where we have used 200 patches with one meter separation. The recorded range of the sonar varies between 150 and 200 meters. Areas with no coverage or with a correlation value below 0.66 is colored white. In the upper plot we show the estimates before sound speed correction, in the center plot after correction and in the lower plot we show the differences. The track of the AUV is shown as a black line in all three plots. The bottom panel show the sound speed profile for the same data set.

5.2 Use of depth estimates in SAS processing

In the previous section we discussed how we have implemented relative depth estimation using sidescan images. The method provided us with fast and robust knowledge of the bathymetry, with a resolution of a few meters. Although this is a valuable product for rapid environmental assessment (REA), it is even more important as *a priori* information in the SAS processing. The depth estimates provides the correct geometry for integration of micronavigation estimates and for imaging. The coherence of the sidescan bathymetry estimates is an estimate of the sonar SNR, which is very important for quality assessment.

5.2.1 Integration of micronavigation estimates

The HUGIN AUV has a high-grade inertial measurement unit (IMU) with a number of aiding sensors integrated in an error-state Kalman filter (Jalving et al., 2003). The attitudes estimates from the integrated navigation system (INS) are accurately enough for SAS imaging. However, the INS position-estimates are limited by a poor along-track velocity (surge) estimate. Synthetic aperture sonar processing, on the other hand, requires extremely accurate cross-track motion (sway) estimates. Depending on the type of error, the demand may be as high as better than $\lambda/16$ position-error over the synthetic aperture (Cutrona, 1975). By combining measurements using DPCA with the navigation solution from the integrated navigation system, the accuracy of both surge and sway may be improved to a level where defocus is avoided (Belletini and Pinto, 2002).

In order to combine DPCA surge and sway measurements with the INS navigation-solution the depression angle Φ (or equivalently: the relative depth), has to be known for each ping and each DPCA patch. In principle, a depth estimate should be an average over all depths within the sonars beamwidth and within the patch interval, weighted by the reflectivity. In order to achieve this all depth estimates have to be considered for each ping. We have implemented and tested this approach, but found that it was unpractical due to the time consumption.

A fast approximation is to consider each ping separately by interpolating the relative body depth estimates, v_{rbd} , into the appropriate range values. The depression angle, Φ is then defined as

$$\Phi = \begin{cases} \arcsin \left(\frac{v_{\text{rbd}}}{|v_{\text{rbd}}|} \cdot \hat{\mathbf{z}} \right) & : \text{starboard} \\ \pi \arcsin \left(\frac{v_{\text{rbd}}}{|v_{\text{rbd}}|} \cdot \hat{\mathbf{z}} \right) & : \text{port}, \end{cases} \quad (5.11)$$

where $\hat{\mathbf{z}}$ is the unit-vector $[0,0,1]$ and Φ is the angle between the vector to the seafloor and the xy body plane of the AUV. A depression angle, $\Phi = \pi/2$ means a direction directly down in the body coordinate system of the AUV.

In addition to the depression angle, the along-track slope of the seafloor relative to the body coordinate system of the AUV is used in the DPCA integration. The slope is

estimated for each patch and each ping by a simple weighted linear fit over the neighboring pings which are within the element beamwidth. Either reflectivity or coherence or a combination of both can be used as weights.

5.2.2 Focus plane for imaging

It is well known that synthetic aperture images defocus if the focus plane is incorrect *and* the sensor path is non-linear (Jakowatz et al. (1996, pages 176-186); Hansen et al. (2009)). For a linear path, cylindrical symmetry ensures that the two-way travel path is independent of depression angle, Φ . For a non-linear path, the travel path becomes a function of Φ , which means that the coherent summation of pings can decorrelate if Φ is incorrect. It has been claimed (Pinto, 2003) that the integration of micronavigation estimates into the navigation solution will compensate for errors in the focus plane, but we believe that this only applies in special cases: In general, both the navigation and the focus plane have to be known.

In the FOCUS processing suite there are two different beamformers: A time-domain interpolation beamformer (Johnson and Dudgeon, 1993, pages 112-119) which works on any form of trajectory and generates SAS images in ground-range, and a wavenumber algorithm (Soumekh, 1999, pages 198-206) which is much faster, but demands close to linear trajectories and generates slant-range images. The sidescan bathymetry depth estimates are used as *a priori* input to both algorithms, but the actual implementations differ strongly.

Depth plane for time-domain imaging

The time-domain interpolation beamformer migrates data from the sensor position given by the navigation system, down onto an earth-fixed focus plane. The plane is made from sidescan bathymetry estimates by mosaicking the estimate into the same coordinate system as defined by the navigation solution. Mathematically, we transform each roll- and pitch-compensated sidescan bathymetry estimate, v_{rd} , into the earth-fixed navigation coordinate system as following

$$v_{ef}^x = v_{rd}^x \cos(\psi) + v_{rd}^y \sin(\psi) + w_{ef}^x, \quad (5.12a)$$

$$v_{ef}^y = v_{rd}^y \cos(\psi) - v_{rd}^x \sin(\psi) + w_{ef}^y, \quad (5.12b)$$

$$v_{ef}^z = v_{rd}^z + w_{ef}^z, \quad (5.12c)$$

where v_{rd}^{xyz} are the xyz -components of v_{rd} , w_{ef} the 3D position-vector of the transmitter and v_{ef} are the depth estimates in the earth-fixed coordinate system. These estimates are scatter points (i.e. they are unevenly distributed in x and y) and cannot be used directly in the imaging.

The time-domain beamformer migrates the sonar data to a regular grid in x and y . The focus plane is therefore divided into rectangular cells with size given by the xy grid-spacing, and the depth for each cell is estimated as a weighted average of a number of

samples in v_{ef} . We have chosen to use a truncated Gaussian spatial weighting combined with coherence weighting. The size of the truncated Gaussian disk is a processing parameter. If it exceeds the cell-size in any dimension, a single sidescan bathymetry depth estimate will contribute to multiple z -estimates. This is a simple way of smoothing the data.

Slant-range directions for wavenumber imaging

The wavenumber algorithm beamforms sonar data in slant-range rather than ground-range, and does not use seafloor depth information. However, the beamforming assumes a linear sensor-trajectory, so we have to motion-compensate the data in advance (Callow et al., 2009). The motion-compensation transforms data collected on a non-linear track to data corresponding to data collected on a linear track. This is in fact a low-resolution time-domain beamforming, which needs seafloor depth information as input.

There are a number of different approaches to the motion-compensation. In wide-beam motion-compensation seafloor depth, not only as a function of range, but also as a function of horizontal angle are taken into account. In narrow-beam motion-compensation it is assumed that the broadside beam is representative for the full horizontal beamwidth. We have chosen a narrow-beam motion-compensation since it is much faster than the wide-beam. With 32 elements in each HISAS receiver array, this means that the depth is only estimated inside 1/32-part of the actual footprint of the element data.

Ideally we should use multiple sidescan bathymetry estimates as depth information for a single ping. This means making an earth-fixed depth map (as described in the previous section) and ray-tracing the beams of each element down onto the depth map. This method uses all available information, but it is very time-consuming. We have implemented this method as an option, but use a simpler approach as default. Sidescan bathymetry estimates from a single ping is used to predict the depths for a single ping in the motion-compensation. Missing estimates (due to too low coherence) are interpolated from neighboring ranges. Only in the case where the sidescan bathymetry estimates for an entire ping are missing, we copy information from neighboring pings.

5.3 SAS interferometry

Even though SAS interferometry is the end-product in the interferometric processing, it is also one of the easier steps to describe. Most of the processing is based on the same principles as the sidescan bathymetry algorithm. In particular, the conversion from delay to relative depth, the sound speed correction and the exporters are all more or less identical to the methods described in Section 5.1. However, other parts are significantly different. In this section we will discuss the details of the co-registration of the SAS images, the estimation of the interferogram, and the unwrapping of the interferogram.

We will also show how we can iteratively estimate the SAS bathymetry by using the first estimate as input to a new co-registration step. This potentially increases the coherence and the accuracy of the estimated bathymetry.

5.3.1 Co-registration

Co-registration of synthetic aperture images is in general a complicated problem. As we mentioned in Section 3.1.4, a method commonly used in SAR is to identify common features in the SAR images and interpolate one of the images to match the other (Fornaro and Franceschetti, 1995). Usually, a warping function describing shift, dilation and rotation is calculated from a set of *control points* or distinctive image features (Hong et al. (2006)); Jakowatz et al. (1996, Pages 293-298)). This is necessary in repeat-pass interferometry (Hanssen, 2001, page 2). For a single-pass interferometer like HISAS, the interferometric images are not rotated relative to each other, so instead of a warping function we use 1D interpolations along range.

Time-domain beamformed images are generated directly in ground-range and do not need to be co-registered. Slant-range SAS images are defined in the same coordinate system as sidescan images, so we can use the same mathematics as in Section 5.1.1. The main difference is that the data is motion-compensated to a straight-line and PCA-corrected to appear as collected from the PCA positions of the interferometric receivers. This means that $r^{\text{sr}}[n] = n \cdot dr$, and Equation 5.2 should be replaced by

$$r_1^{\text{gr}}[m] = 2\sqrt{(y[m] - y_1^{\text{PCA}})^2 + (z[m] - z_1^{\text{PCA}})^2}, \quad (5.13a)$$

$$r_2^{\text{gr}}[m] = 2\sqrt{(y[m] - y_2^{\text{PCA}})^2 + (z[m] - z_2^{\text{PCA}})^2}, \quad (5.13b)$$

with $y_1^{\text{PCA}} = y_1/2$, $y_2^{\text{PCA}} = y_2/2$, $z_1^{\text{PCA}} = z_1/2$ and $z_2^{\text{PCA}} = z_2/2$. Figure 3.1 defines the receiver positions y_1 , y_2 , z_1 , z_2 . The actual interpolation is the same as in Equation 5.3 and the basebanding is given by Equation 5.5. The *a priori* seafloor estimates $z[m]$ for each $y[m]$ is found by defining an earth fixed coordinate system and applying the methods described in Section 5.2.2.

5.3.2 Estimating the interferogram

Once we have co-registered the images in ground-range, the estimation of the interferogram is straightforward. We implement the maximum likelihood estimator of the complex coherence described in Equation 4.29 (see Section 8.1). The argument of the complex coherence is the interferogram, while the magnitude is the coherence. Note that the frequency used when converting the interferogram to relative depth is the scaled *ground-range frequency* (see Section 5.1.2).

5.3.3 Unwrapping the interferogram

Since the interferogram is wrapped modulo 2π , we have to unwrap the phases before we can estimate the relative depth. The actual implementation follows the description in Section 4.6.8. As we discussed, the phase wraps are minimized by the ground-range co-registration, and the residual unwrapping is performed using a standard Goldstein's Branch cut algorithm as described in Ghiglia and Pritt (1998).

5.3.4 Generating SAS images using SAS bathymetry

In this section we have described how FOCUS relies on relative seafloor depth estimates, made using the sidescan bathymetry algorithm. The performance of preceding methods rely on the accuracy of these estimates. SAS bathymetry is an estimate of the relative seafloor depth with much higher horizontal resolution than sidescan bathymetry. An iterative scheme can therefore be made by using SAS bathymetry as an improved *a priori* input to the SAS processing. This can be implemented throughout the processing chain, but we have chosen to focus on an improved estimation of a focus plane. An improved *a priori* input will provide a better co-registration and less occurrence of 2π phase wraps. This again will give us depth estimates with better precision and accuracy. Figure 5.8 shows the interferograms using flat seafloor, sidescan bathymetry and SAS bathymetry as *a priori* focus plane. The results are generated using a point scatter simulation of a bumpy seafloor (see the upper left panel of Figure 5.8). The interferograms get closer to zero, the better the *a priori* depth information is. Since the simulated seafloor is very bumpy, there is still a residual phase-difference left in the interferogram after iterative SAS processing (see the lower right panel of Figure 5.8).

Figure 5.9 shows the SNR calculated from the estimated coherence (using Equation 4.17). The figure illustrates the increase in SNR due to improved co-registration. The average improvement by using iterative SAS with SAS bathymetry as focus plane instead of a flat seafloor (which has maximum 1 meter error), is in this case 4 dB. In most realistic cases, the seafloor depths are sufficiently smooth so sidescan bathymetry will be accurate enough as depth information. However, on rock-formations or large man-made objects iterative SAS may significantly improve the results.

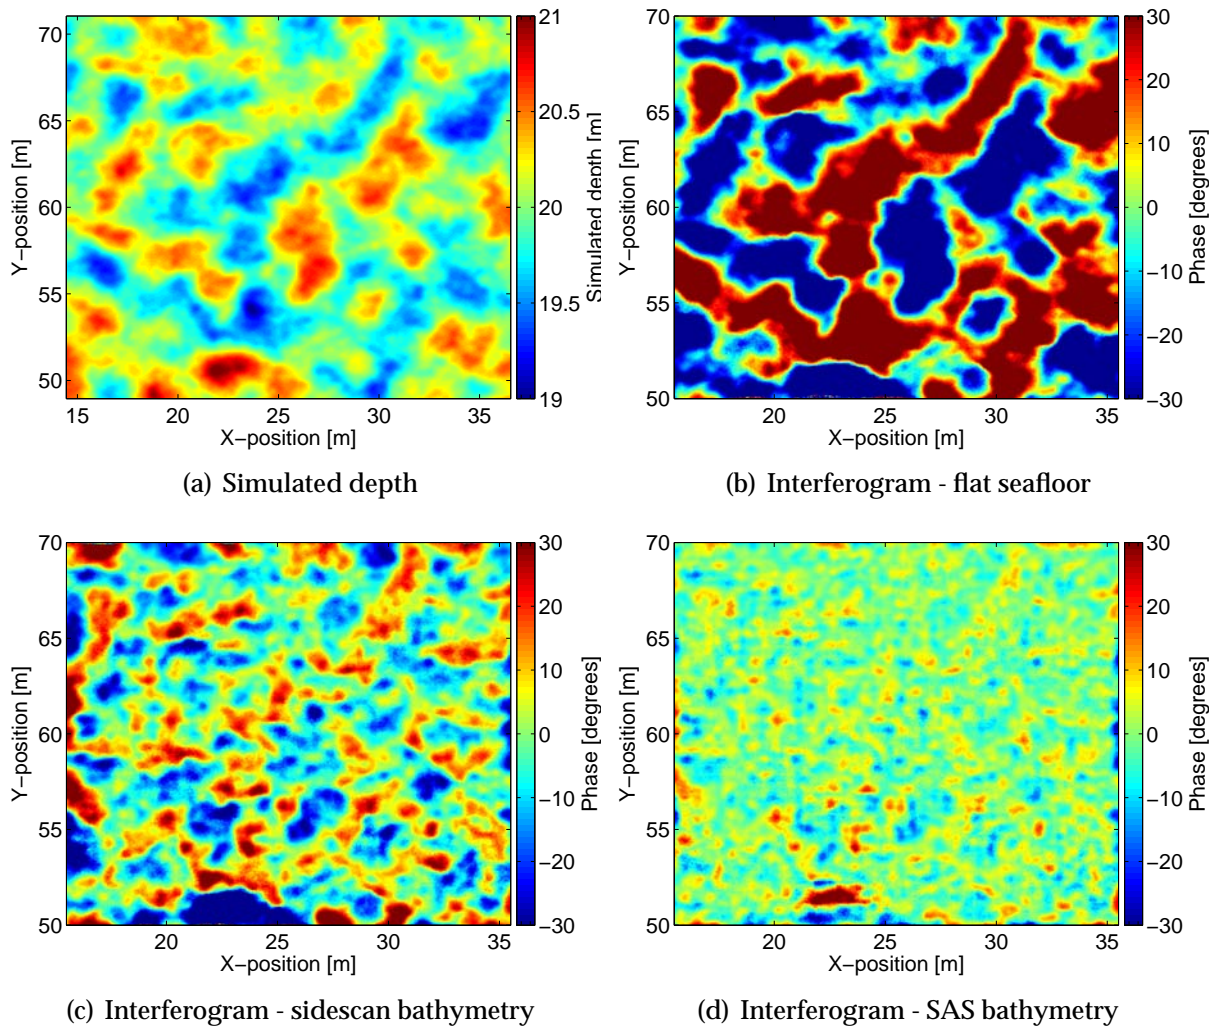


Figure 5.8: The figure shows results using different prediction methods for the focus plane. The results are generated using a point scatter simulation of a bumpy seafloor. The simulated depth is shown in the upper left panel. The upper right panel shows the interferogram assuming a flat seafloor at 20 meters as a *priori* depth. The lower left panel shows the same interferogram using sidescan bathymetry estimates as a *priori* depth, and the lower right panel show the interferogram using iterative SAS with the focus depth predicted from SAS bathymetry estimates.




Figure 5.9: The figure shows interferometric SNR using different prediction methods for the focus plane. The results corresponds to the results showed in Figure 5.8. The left panel shows the estimated SNR assuming a flat seafloor in the imaging and the right panel shows the SNR using iterative SAS with the focus depth predicted from SAS bathymetry estimates. The average increase in SNR is 4 dB.

Chapter 6

System description

The methods we have described in this thesis are applicable to any SAS- or SAR-system. However, in addition to the raw sonar (radar) data, there are some on-site measurements we assume are available:

- A navigation solution with positions and attitudes. Preferably, the navigation sampling frequency should be at least ten times higher than the ping frequency.
- Altitude measurements. Height above the seafloor is used as initialization in sidescan bathymetry. This initialization provides a more accurate estimate of the depth for the first pings.
- An estimate of the local sound speed and the sound speed profile. An incorrect local sound speed will cause a bias in the estimated bathymetry. As we have seen in Section 5.1.5 large systematic errors can occur in the bathymetric estimates if the sound speed profile is unknown. The bathymetry can be corrected in post processing provided the applied sound speed is known.

The *system parameters* of the sonar and the platform have to be known with high accuracy. This includes lever-arms, frequencies, bandwidth etc. Almost all the experimental data in this thesis is from the HISAS 1030 interferometric synthetic aperture sonar on the HUGIN 1000-MR autonomous underwater vehicle. Data from other systems is included for the purpose of illustration.

6.1 The HISAS 1030 interferometric SAS

The HISAS 1030 is a two-sided synthetic aperture sonar (Fossum et al., 2008; Hagen et al., 2006). It has two full-length receiver arrays on each side, which is a necessity for interferometric processing with image resolution. Each receiver array consists of 32 individual elements of size 3.75 cm along-track (see Figure 6.1). The transmitter is a vertical phased array of 16 $\lambda/2$ elements (at 100 kHz), which means that the transmitted




Figure 6.1: Schematic of one of the HISAS 1030 sonars. Each receiver array consists of 32 elements of size 3.75 cm along-track. The vertical baseline between the receivers is 30 cm or 20λ at 100 kHz. The transmitter is located approximately 20 cm in front of the midpoint between the receiver arrays. The transmitter consists of 16 elements in a vertical array, each of size $\lambda/2$ at 100 kHz.




Figure 6.2: Picture of a receiver array and a transmitter from the HISAS 1030.

pulse can be vertically steered to a preferable direction. Since the transmitter also records during reception, it can be used as a short interferometric array with $\lambda/2$ element spacing.

HISAS can operate between 50 and 120 kHz, typically using 30 kHz bandwidth at 100 kHz. With 30 kHz bandwidth the data rate is approximately 85 Gb/h (included re-




Figure 6.3: Picture of a HISAS 1030 mounted on a HUGIN 1000-MR AUV.

ording at the transmitter). The waveforms are LFM pulses. The theoretical image resolution in standard mode is 1.9 cm along-track and 2.5 cm cross-track, The maximum range of the system is around 400 meters due to attenuation and self noise (Lurton, 2002, pages 19-23). The recorded data are matched filtered and dynamically focused in real time, and both the individual receiver data and the sidescan data are stored. The system is sold to multiple Navies with a specification of better than 5 by 5 cm up to 200 meters range, providing possibility for windowing and filtering of the data.

With a center frequency of 100 kHz, the horizontal beamwidth is approximately 23.6 degrees. This provides multi-aspect capability, which for example can be used in shadow detection (Midtgaard et al., 2007). The reduction of synthetic aperture length in multi-aspect processing can also reduce the effect of shadow blurring inherent in SAS (Groen et al., 2009). Alternatively, *fixed focusing* proposed by Groen et al. (2009) can be used to restore the shadows.

Attenuation causes SNR to become marginal at ranges above 200 meters range, but the incident angle at the seafloor may also limit the backscattered energy. Therefore, it has been suggested by Hagen and Hansen (2007) that the maximum practical range for a SAS system is 10 times the altitude. A conservative estimate of the blind zone is from nadir and 45 degrees out. The resulting area coverage rate is around 2.3 km³/h (Hagen and Hansen, 2007). Table 6.1 summarizes all key features of the HISAS.

Table 6.1: Key features for the HISAS 1030 interferometric synthetic aperture sonar

Quantity	Value
Receiver length	120 cm
Number of receive elements on each side	2x32
Vertical baseline	30 cm
Transmitter size	7.5x15 cm
Number of transmitter elements	8x16
Typical mounting roll angle	20 degrees
Frequencies	50-120 kHz
Typical operational bandwidth	30 kHz
Typical pulse type	LFM
Typical pulse length	6 ms
Typical operational range	200 m
Theoretical resolution	1.9x2.5 cm
Practical resolution	5x5 cm
Theoretical area coverage rate	3.25 km ² /h
Practical area coverage rate	2.3 km ² /h

6.2 The HUGIN 1000-MR AUV

The HUGIN autonomous underwater vehicles have been developed by the Norwegian Defence Research Establishment (FFI) and Kongsberg Maritime (Hagen et al., 1999, 2004, 2005). The development started in 1991 with a technology demonstrator, and the first sea trial took place in 1996. Originally, the HUGIN was designed for surveying in deep waters. With the HUGIN 1000-MR, the aim has been to develop an AUV especially suited for military purposes.

The HUGIN 1000-MR is modular in size (see Figure 6.4). The vehicle frame is made of carbon fiber laminate and high performance syntactic foam. During operation, the vehicle is filled with water so the electronics is sheltered inside two spherical containers. Heavy equipment such as batteries and transducers can be used due to the buoyant frame. Typical dry weight of the vehicle is 600-800 kg.

The energy is delivered from two rechargeable Lithium polymer battery modules and a high efficiency, high torque synchronous motor allows direct drive of a long bladed propeller. Four individually controlled rudder blades are used for heading, pitch and roll control. The vehicle is controlled by a distributed software system featuring three or more main computers. The navigation processor runs the real-time integrated




Figure 6.4: Schematic of the different parts of the HUGIN 1000-MR AUV. Notice the modular design with three separate sections.

inertial navigation system. After recovery, the navigation is improved by using non-causal smoothing (Gade, 2004).

HISAS is the primary sensor on the HUGIN 1000-MR (see Figure 6.3). However, it is also the most energy consuming sensor. With HISAS the endurance of the HUGIN 1000-MR is approximately 18 hours.

6.3 HISAS compared to selected InSAR systems

In this section, we consider four different interferometric synthetic aperture radar systems and compare them with the HISAS interferometric SAS. The basic idea is to shed light on the main differences in InSAS compared to InSAR. The goal of this section is to highlight possible areas where InSAS can be improved based on state-of-the-art in InSAR. The results and presentation in this section is based on Sæbø and Hansen (2010).

6.3.1 System descriptions

There are noticeable differences between SAR and SAS:

- Almost all SAS systems (including HISAS) have a single along-track element transmitter and a multiple along-track receiver array, and use stripmap mode. The multi-channel receiver is primarily used to obtain large swaths.

Table 6.2: Sensor parameters for the airborne STAR-4 and PAMIR, the satellite based SRTM-X and TanDEM-X and HISAS.

	STAR	PAMIR	SRTM	TanDEM	HISAS
Frequency [MHz]	9580	9450	9600	9650	0.1
Bandwidth [MHz]	270	1820	9.5	150	0.03
Wavelength [cm]	3.1	3.2	3.1	3.1	1.5
Element size [m]	1.25	0.20	12	4.8	0.038
Array size [m]	1.25	0,60	12	4.8	1.2
Baseline [m]	0.96	1	60	500	0.28
Roll angle [deg]	81	0	45	0	14
Beamwidth [deg]	1.4	9.1	0.15	0.37	23
Range resolution [m]	0.55	0.08	15	1	0.02
Along resolution [m]	0.6	0.10	6	2.4	0.02

- Advanced SAR systems have 2D phased arrays and use electronic steering of the transmitter and receiver arrays for different SAR modes. ScanSAR (Franceschetti and Lanari, 1999, chapter 5) and *Terrain Observation by Progressive Scans* (TOPS) (Gebert et al., 2010) are used to increase the area coverage rate while lowering the resolution. Spotlight mode (Jakowatz et al., 1996) is used to increase the resolution while lowering the area coverage rate.

Table 6.2 summarizes the key parameters for the different sensors. The parameters are shown for stripmap mode for all systems (except PAMIR). In the following, we describe the four difference SAR sensors we have chosen for our study.

STAR-4

Intermap Technologies from Canada has developed several airborne interferometric SAR systems used commercially for topographic mapping. The radar is a X-band stripmap SAR with two horizontally displaced antennas (Lange and Gill, 2009; Wei and Coyne, 2008). Their TopoSAR system has two transmitters and two receivers and the transmitter can be alternated giving an effective doubling of the baseline (referred to as ping-pong) (Rosen et al., 2000). We have chosen the STAR-4 which provides a good compromise between high resolution and large area coverage rate. We consider this system a good representative of commercial airborne SAR.

PAMIR

The *Phased Array Multifunctional Imaging Radar* (PAMIR) is a high resolution phased array airborne SAR system with possibilities for cross-track interferometric mapping and along-track interferometry moving target indication (MTI) (Ender and Brenner, 2003; Brenner and Roessing, 2008). The interferometric array includes three vertically spaced antennas with baselines of 0.67 meters and 0.335 meters, giving a total baseline of one meter (Brenner and Roessing, 2008). This is a state-of-the-art airborne SAR, and with a recent upgrade to 3 GHz bandwidth the system has demonstrated SAR imagery with a resolution of 2×5 cm (Brenner, 2010). In repeat pass interferometry using a baseline of 17 m, PAMIR also demonstrated extreme resolution and accuracy in interferometry (Brenner et al., 2010). This radar is fully programmable and configurable. It is difficult to retrieve the actual performance from literature and we have therefore chosen parameters from a sliding spotlight interferometry test (Brenner and Roessing, 2008). Thus, the actual parameters for stripmap interferometry might be incorrect. The system is still interesting to consider, representing state-of-the-art in airborne SAR.

SRTM-X

In 2000, the *Shuttle Radar Topography Mission* (SRTM) successfully recorded single pass SAR interferometry data from the entire land mass of the earth between 60° N and 57° S, both in C-band and X-band (Rabus et al., 2003; Suchandt et al., 2006). This was the first single pass SAR interferometer in space. The maximum swath width of the C-band system was 225 km in scansar mode using 4 beams, and 45 km in X-band stripmap mode (Rabus et al., 2003). The main antenna in the shuttle was 12×0.5 meters large, and the outboard antenna was 6×0.4 meters for the X-band frequency. The secondary antenna was mounted on a 60 meters long boom to obtain a long baseline.

TanDEM-X

TerraSAR-X is a high resolution spaceborne SAR system for earth observations (Stangl et al., 2006). The radar antenna is a phased array with 384 elements, with the possibility of multiple modes of operation: Stripmap mode with 3 meter resolution and 30 km swath width; scansar mode with 16 meters resolution and 100 km swath width; and high resolution spotlight mode with one meter resolution and 10 km image size. In June 2010 the TanDEM-X was launched (Krieger et al., 2007). This is an ad-on satellite to TerraSAR-X for digital elevation mapping (DEM). Both satellites are considered one bistatic system with synchronous formation flying (Krieger et al., 2010) providing single pass high resolution interferometry with large baselines. Both satellites will fly in helix tracks, forming both horizontal and vertical baselines up to several kilometers (Krieger et al., 2010). The interferometric data collection with the two satellites can be performed in different ways since both platforms have the possibility to transmit and receive. TanDEM-X represents state-of-the-art in horizontal resolution and vertical accuracy in

Table 6.3: Imaging geometry for the airborne STAR-4 and PAMIR, the satellite-borne SRTM-X and TanDEM-X, and HISAS. The parameters are shown for stripmap mode for all systems (except PAMIR).

	STAR	PAMIR	SRTM	TanDEM	HISAS
Altitude [km]	8.5	3	233	515	0.02
R_{\max} [km]	17	7	410	728	0.2
$R_{\text{sw,gr}}$ [km]	11	2	45	30	0.17
$R_{\text{sw,sr}}$ [km]	7.7	1.7	36	21	0.17
Speed [m/s]	111	100	7500	7600	2
PRF_{\min} [Hz]	178	333	1250	3167	3.33
PRF_{\max} [Hz]	19430	86207	4167	7143	4.37
PRF [Hz]	1875	1000	1674	4500	3.75
SA length [m]	425	1111	1068	4740	80
SA time [s]	3.8	11	0.14	0.62	40
N_{λ} in SA	13600	35000	34167	151667	5333
N_p in SA	7180	11000	238	2800	150
Rate [Mpixels/s]	2.2	25	3.75	95	0.85

spaceborne single pass interferometry.

6.3.2 Synthetic aperture data collection

There are fundamental differences between InSAR and InSAS in environmental conditions, phase velocity, geometry, data collection, platform motion and navigation (see Section 2.5). In this section, we list some of the differences that affect synthetic aperture imaging and interferometry. Table 6.3 shows the key numbers for geometry and data collection for the different sensors. The last row shows the data-rate defined as the area coverage rate in square meters per second divided by the theoretical resolution in square meters. This gives a indication of the hardware complexity and the required processing power.

The imaging geometry of existing SAS systems are very similar, with a swath reaching from nadir to roughly ten times the altitude. This geometry is very different from spaceborne SAR systems, which have a much more vertical geometry. This reduces the effect of shadowing, but increases the effect of foreshortening and layover (Franceschetti and Lanari, 1999, page 37-41). An airborne SAR system usually has an imaging geometry somewhere between a SAS and an spaceborne SAR. Figure 6.5 shows the measurement geometry for the different sensors.




Figure 6.5: Comparison of the imaging geometries for the different sensors. The red area shows the actual swath. Note that all geometries have correct aspect except for HISAS, which has even more horizontal geometry than indicated.

Table 6.4: Interferometric parameters for the airborne STAR-4 and PAMIR, the satellite-borne SRTM-X and TanDEM-X and the HISAS. The parameters are shown for stripmap mode for all systems (except PAMIR).

	STAR-4	PAMIR	SRTM-X	TanDEM-X	HISAS
σ_z [m] (10dB)	1.47-8.94	1.96-2.61	1.84-2.10	0.52-0.54	0.01-0.13
$\frac{\partial\theta}{\partial z}$ [deg/m]	2.88-0.47	2.16-1.62	2.3-2.0	8.1-7.9	309-32
$z_{2\pi}$ [m]	125-760	167-222	157-178	44-45.5	1.2-11
f_θ [millicycles/res]	4.1-0.37	0.28-0.16	62.8-50.3	16.5-15.5	15-0.23
D_{crit} [m]	118-1312	1776-3146	478-596	15200-16100	9-622
k_B [no shift]	0.9919-	0.9994-	0.8745-	0.9670-	0.9696-
	0.9993	0.9997	0.8993	0.9689	0.9995

6.3.3 Interferometric processing

In the design of an interferometric system there are a number of important quantities to consider. In this section, we list parameters that indicate the sensitivity / accuracy that can be obtained, and the complexity / difficulty of the processing. We show the results in Table 6.4 and in Fig. 6.6.

Design parameters

As we remember from Section 3.3, the interferometric phase-difference, θ , can be converted to a depression angle relative to the interferometer, Φ

$$\Phi \approx \sin^{-1} \left(\frac{\lambda}{2\pi D} \theta \right), \quad (6.1)$$

where D is the interferometric baseline (i.e. the distance between the receiver arrays). We assume that the incoming waves are parallel. The depth relative to the sensor was presented in Equation 3.9

$$z = r \sin (\Phi + \Phi_0) \quad (6.2)$$

where r is the slant-range and Φ_0 the roll of the interferometric array.

A simple differentiation of Eqs. 6.1 and 6.2 gives the standard deviation of the depth estimate as a function of the standard deviation of the phase-difference estimates σ_θ

$$\sigma_z = \frac{\lambda}{2\pi} \frac{r \cos(\Phi + \Phi_0)}{D \cos \Phi} \sigma_\theta, \quad (6.3)$$

where r , B and Φ_0 are assumed known. Clearly, a large interferometric baseline with short wavelengths gives the best performance (lowest standard deviation in depth). However, as we will see later, an increase in the interferometric baseline also causes a reduction in coherence and robustness.

An important quantity is the height sensitivity (Bamler and Hartl, 1998)

$$\frac{\partial \theta}{\partial z} = \frac{2\pi}{\lambda} \frac{D}{r} \frac{\cos \Phi}{\cos(\Phi + \Phi_0)}, \quad (6.4)$$

which gives the phase-difference for a given geometry. Equivalently, this can be represented by the height ambiguity, or the height which causes one phase wrap at a given range

$$z_{2\pi} = 2\pi \left(\frac{\partial \theta}{\partial z} \right)^{-1} = \lambda \frac{r \cos(\Phi + \Phi_0)}{D \cos \Phi}. \quad (6.5)$$

The local *fringe frequency* in range is a key parameter of the interferometer

$$f_\theta = \frac{1}{2\pi} \frac{\partial \theta}{\partial r} = -\frac{z}{\lambda} \frac{D}{r^2} \frac{\cos \Phi}{\cos(\Phi + \Phi_0)}. \quad (6.6)$$

This expression gives the change in phase-difference as a function of range and is identical to the spectral shift of the signal. It can be used to calculate the maximum theoretical coherence between the receiver arrays. The *critical baseline*, D_{crit} is the baseline which gives a spectral shift equal to the bandwidth, B

$$D_{\text{crit}} = \lambda \frac{B}{c} \frac{r^2 \cos(\Phi + \Phi_0)}{z \cos \Phi}. \quad (6.7)$$

The baseline dependent coherence, k_B , can be approximated as (Bamler and Hartl, 1998)

$$k_D \approx 1 - \frac{D}{D_{\text{crit}}} \quad \text{for } D \ll D_{\text{crit}}, \quad (6.8)$$

with the total coherence written as

$$k = k_{\text{SNR}} k_D k_T. \quad (6.9)$$

Here k_{SNR} is the coherence due to a finite SNR and k_T is the temporal coherence.

In Table 6.4 and Fig. 6.6 we summarize the most important interferometric features for our example systems. There are rather large differences between the systems. Some of these are:

- HISAS has a large relative variation in range causing large variations in all the sensitivity parameters.
- TanDEM-X has large swath, high horizontal resolution and good vertical accuracy due to the large baseline.
- SRTM-X has relatively large spectral shift compared to the bandwidth.
- PAMIR has very low fringe frequency indicating that unwrapping is less difficult.
- STAR-4 has very large height ambiguity which indicates that unwrapping is not needed.
- PAMIR, HISAS and STAR-4 have large critical baselines compared to the actual baseline. This reduces the need to perform spectral shift filtering.
- TanDEM-X has a very large critical baseline, allowing for a high interferometric range resolution.
- HISAS has a very low height ambiguity at short range. This indicates the need for good unwrapping algorithms.
- STAR-4, PAMIR and SRTM-X have similar standard deviations of the height estimate for fixed number of pixels and SNR. This means that PAMIR has much higher vertical accuracy for fixed horizontal patch in square meters, due to the much higher resolution.

6.3.4 Summary

In this section, we have compared state-of-the-art interferometric SAR systems with the HISAS sonar. We have chosen the systems somewhat subjectively in order to shed light on what can be learned from InSAR. There are rather large differences between the different SAR systems. The differences can be grouped in different ways. Spaceborne sensors are different from airborne sensors. SAS is closer to airborne SAR. Complex and new sensors have better overall performance than simpler older sensors. This is in particular apparent when comparing SRTM-X with TanDEM-X.

Interferometric SAS and HISAS in particular, has very large variations in the interferometric performance over the entire swath due to the large range variation. At short range, the accuracy is extremely high, but the height ambiguity very low. At far range, the fringe frequency is very low and the critical baseline very high. At short range, phase unwrapping is probably the largest challenge, while at far range, lack of SNR is probably the largest challenge. In sum, these findings suggests that optimal interferometric processing should be range dependent.

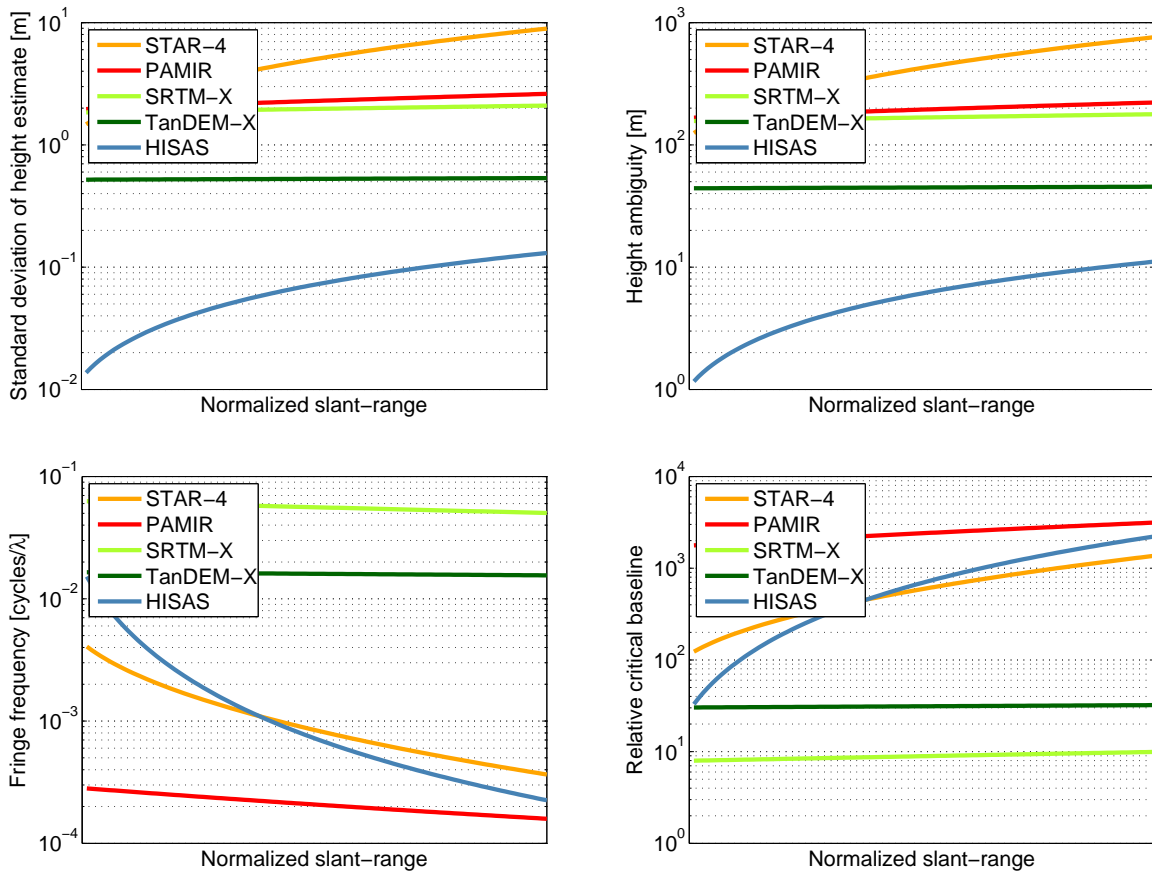


Figure 6.6: Interferometric performance of the airborne STAR-4 and PAMIR, the spaceborne SRTM-X and TanDEM-X and HISAS as a function of normalized slant-range. Upper left: theoretical height accuracy; Upper right: height ambiguity; Lower left: fringe frequency; Lower right: critical baseline.

HISAS has large relative bandwidth and large beamwidth compared to SAR systems. This indicates that the split-bandwidth interferometry technique (see Section 8.3) can be considered to simplify or avoid unwrapping problems.

A trend in interferometric SAR is the use of repeat pass interferometry for deformation mapping, temporal change detection and multi-baseline tomography. We believe that the future of InSAS brings these advanced techniques to the underwater environment too (Synnes et al., 2010).

Chapter 7

Resolution and precision assessment

In this section we discuss the theoretical, the practical and the achievable resolution and precision of the depth estimates. We focus on two quantities: Horizontal resolution, i.e. the resolution in x, y or in along-track and cross-track projected down onto the seafloor (ground-range), and the precision in the depth estimate (precision along the z -coordinate). We will also separate between sidescan and SAS bathymetry, although many of the arguments and equations will appear in both topics.

The horizontal resolutions of bathymetric estimates are complicated. We need to define a sensible and applicable definition of resolution in a bathymetric map. We choose two different approaches: When *calculating* the resolutions mathematically, we start by the underlying image resolution and investigate how the bathymetric processing may affect these measures. When we want to *measure* the achieved resolution in a bathymetric dataset we consider the flank of a distinct change in the bathymetry. We measure the horizontal distance between where the bathymetry is at one level to where it has reached the other level (assuming that the change in bathymetry is instantaneous and that the change is significantly larger than the precision of the estimates, so that the two levels can be easily separated). This estimate of resolution is *ad-hoc* and we do not claim that it can be deduced from the theoretical measures, but it is strongly influenced by our perception of resolution.

In this section and throughout the thesis, simulated data are generated using a point scatter simulator (with a few exceptions in Chapter 4). The simulator uses independent, non-directive scatterers which are randomly distributed in the horizontal plane and are located on a virtual surface in the vertical plane. The strength of each scatterer is calculated using a standard normal distribution. The simulator is implemented in time-domain and generates raw sonar data using LFM pulses. Sonar beampattern, transmission-loss and continuous motion through pings are simulated accurately. So is the sonar trajectory. However, the simulator does not simulate a realistic scene. The scattering mechanism is simplified and effects like shadow, multipath and specular reflections are omitted. A realistic speckle statistic is still achievable by choosing a large enough number of scatterers in each resolution cell. In this thesis we use approximately ten reflectors (or more) in each resolution cell. This means that a simulation for SAS interferometry has

to have a much higher density of scatterers than a simulation for sidescan bathymetry.

7.1 Sidescan bathymetry

As mentioned in Chapter 2, a sidescan image has a theoretical along-track resolution of (Van Trees, 2002, pages 46-51)

$$\delta_x^{\text{ssi}} \approx \frac{\lambda}{L}r = \frac{\lambda}{N_r d}r, \quad (7.1)$$

where λ is the wavelength, $L = N_r d$ is the length of a receiver array consisting of N_r elements of size d , and r is the range. Note that the along-track resolution is range-dependent and that this is the 3 dB resolution (i.e. a measure of the distance needed to separate two targets of equal strength by 3 dB) (Van Trees, 2002, pages 46-51). Other types of resolution can be defined (Van Trees, 2002, pages 46-51). In Section 2 we also defined the theoretical 3 dB range resolution of a sidescan image as (Franceschetti and Lanari, 1999, page 23)

$$\delta_r^{\text{ssi}} = \frac{c}{2B}, \quad (7.2)$$

where c is the sound velocity and B the bandwidth of the system.

In sidescan bathymetry we use a patch of data along a range-profile (i.e. a small segment of a time series) to estimate the depth (see Section 5.1.2). Therefore the theoretical sidescan bathymetry resolution should equal the sidescan image resolution along-track, and the number of independent samples in the patch times the sidescan image resolution cross-track. Along the array we do not apply any weighting, nor do we perform any other kind of processing which will deteriorate the resolution so the *practical* (and theoretical) along-track resolution in sidescan bathymetry equals the theoretical resolution of the sidescan images

$$\delta_x^{\text{ssb}} = \delta_x^{\text{ssi}} = \frac{\lambda}{N_r d}r. \quad (7.3)$$

Cross track the story is more complicated. Firstly, we have to consider the length of the patch of data. Secondly, we apply a weighting after the matched filtering which reduces the resolution. Thirdly, we render the sidescan image down onto a (non-flat) seafloor (which yields a resolution in y instead of in r). Fourthly, we apply weighting on the patch data before we correlate the interferometric arrays (which effectively reduces the size of the patch). The bathymetric y -resolution thus becomes

$$\delta_y^{\text{ssb}} = w_1 w_2 N_s T_{ry} \delta_y^{\text{ssi}} = w_1 w_2 N_s T_{ry} \frac{c}{2B}, \quad (7.4)$$

where N_s is the number of independent samples in the patch, w_1 a factor due to the window applied in the matched filter, w_2 a factor due to the window applied on the patch-data and T_{ry} a factor due to the resampling from range to y . While N_s can be

large, the three other terms are all close to one and can be written on an approximate form

$$\delta_y^{\text{ssb}} \approx \frac{4}{3} \frac{3}{4} N_s \frac{r}{\sqrt{r^2 - h^2}} \frac{c}{2B}, \quad (7.5)$$

where r is the range and h the depth of the seafloor at the patch location relative to the sonar depth. Here we have used $w_1 \approx 4/3$ and $w_2 \approx 3/4$ (Van Trees, 2002, chapter 3) and assumed a local flat seafloor (at the patch). In practice, we sample the data uniformly at theoretical range resolution along y , and select N'_s samples from the ground-range data

$$N'_s \approx \frac{4}{3} N_s \frac{r}{\sqrt{r^2 - h^2}}. \quad (7.6)$$

Consequently, the sidescan bathymetry resolution along y can be written

$$\delta_y^{\text{ssb}} \approx \frac{3}{4} N'_s \frac{c}{2B} = \frac{3}{4} P, \quad (7.7)$$

where N'_s is the number of samples in a ground-range profile sampled at theoretical range resolution and P the length of the ground-range patch in meters.

Figures 7.1 and 7.2 show results from sidescan bathymetry on a simulated dataset of a single ping in a noise-free environment. The simulated seafloor consists of a set of half meters delta-changes as can be seen in Panel 7.1(a). The depths are estimated on a large number of lines beamformed at different x -positions. Panel 7.1(b) shows the depth estimates. Notice the loss of correlation when a step change is inside the patch (see Panel 7.1(c)).

In Figure 7.2 we show how the method achieves close to the theoretical resolutions described in Equations 7.3 and 7.7. Panel 7.2(a) shows the range profile for two different lines. Notice that the half meter change in bathymetry is significantly larger than the precision of the estimates and that the transition between the different regions are almost linear. In Panel 7.2(b) we show the estimated length of the slope averaged over range for each line. In this simulation $\delta_y^{\text{ssb}} \approx 3.2$ meters and the estimated slope lengths match quite well with the theoretical resolution. Averaging over all lines gives an estimated y -resolution of 3.16 meters.

In Panel 7.2(c) we show the along-track results. The length of the slope now varies with range. The theoretical angular resolution is 0.72 degrees and is shown in solid black. The dashed lines show the estimated slope length which corresponds to an angular resolution of 0.61 degrees. This reduction of 15 percent is due to the fact that the slope-length not necessarily corresponds to the 3 dB resolution, but it should correspond to *some* image resolution. The important result is that the along-track sidescan bathymetry resolution also seems to be proportional to range, and have the same trend as the theoretical sidescan resolution.

Another important parameter is the *posting separation*, which is often confused with the resolutions mentioned above (when the posting separation equals the resolution, the data is critically sampled in such a way that each sample is independent of the others

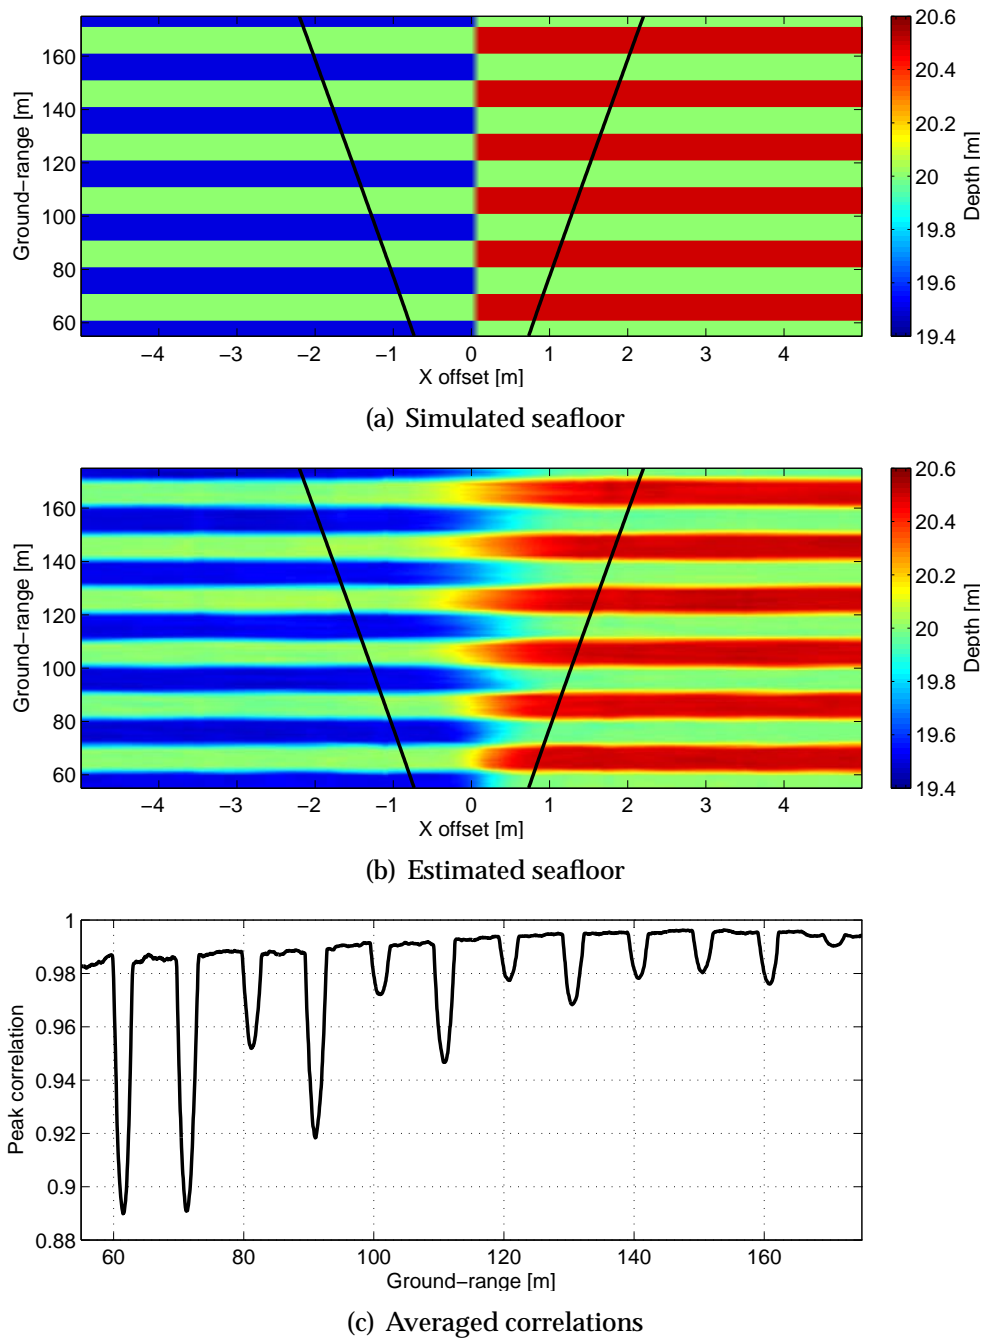





Figure 7.1: The figure shows results from sidescan bathymetry on a single simulated noise-free ping. The bathymetry is estimated on a number of offset beams. Panel (a) shows the simulated seafloor, while the estimated seafloor is presented in panel (b). Panel (c) shows the peak correlations as a function of ground-range averaged over offset beams. The beams close to zero offset are omitted, since they are affected by the along-track bathymetric change.



(a) Estimated range profiles



(b) Estimated resolution along ground-range



(c) Estimated resolution along-track

Figure 7.2: The figure shows results from sidescan bathymetry on a single simulated noise-free ping. The bathymetry is estimated on a number of offset beams. Panel (a) shows two of the estimated range profiles, offset 5 meters from broadside. Panel (b) shows the estimated ground-range resolution for each offset line. Panel (c) shows cross-sections along-track at different ranges. The colored lines represent different ranges. The theoretical and estimated resolutions are drawn on top. Notice that the apparent asymmetrical beam pattern is caused by the fact that the range changes with the bathymetry when the ground-range is kept constant.




Figure 7.3: STD of the sidescan bathymetry depth estimate errors as a function of ground-range for a set of different SNRs. The results are calculated from Equation 7.8 with $D=0.28$ meters, $f_0=100$ kHz, $c=1500$ m/s and $N = 128 \cos(\Phi)$, with $\Phi = \tan^{-1}(h/r_g)$ where r_g is the ground-range and $h = 20$ is the depth of the seafloor relative to the sonar.

and no information is lost due to undersampling). In sidescan bathymetry, the posting separation along x is given by the motion of the sensor between pings. Since the sensor can change its pointing direction from one ping to another, the posting separation may change as a function of ping and range (it may even become zero). For a SAS-system it usually equals slightly less than half the array-length. For HISAS it is typically 0.5 meter. Along ground-range y , the posting separation can be chosen arbitrarily, only affecting processing time. We use one meter as default. This gives an oversampling which ensures that the estimated surface is sufficiently continuous and smooth.

Next we consider the accuracy and precision of sidescan bathymetry. We define precision as a measure of the error relative to the true answer. In the case where there are no ambiguities and no bias, the precision equals the accuracy, which is a measure of noise in the estimates. In this thesis we mostly assume bias- and ambiguity-free estimates in the mathematical calculations, which means that we only calculate the accuracy. The precision is discussed by comparing results with the ground truth, where available, or by observing obvious flaws in the estimates.

As mentioned in Section 4.5 there are two different expressions which describe the accuracy of interferometric time delay estimation. Either one uses Equation 4.39 and assumes that the recorded data are deterministic, time delayed time series (which they are not, since the signal is scattered from a 3D geometry), or one uses Equation 4.40 and assumes that the recorded data are co-registered datasets with speckle statistics (which breaks down when there are objects on the seafloor). Since the two expressions are close to but not equal, we have chosen to consider sidescan bathymetry as time delay estimation and SAS interferometry as interferometric processing.

Combining Equations 3.24 and 4.39 gives an expression for the standard deviation

of the sidescan bathymetry estimates

$$\text{std}\{z_{\text{est}}\} \approx \frac{cr_2}{2D} \frac{1}{2\pi f_0} \frac{1}{\sqrt{N}} \sqrt{\frac{1}{\rho} + \frac{1}{2\rho^2}}, \quad (7.8)$$

where c is the sound speed, r_2 the range from receiver #2 to the seafloor, D the vertical baseline, N the number of independent samples and f_0 the center frequency. Remember from Equation 4.17 that the SNR, ρ , can be estimated from the coherence.

Figure 7.3 shows results from Equation 7.8 using HISAS settings, i.e. $D=0.28$ meters and $f_0=100$ kHz. The sound speed, c , is set to 1500 m/s and we have used a fixed patch length of 3.2 meters in ground-range. This is equivalent to $N = 128 \cos(\Phi)$ independent samples, where Φ is the depression angle. The increase at short range in Figure 7.3 is caused by the increasing depression angle, which reduces the number of independent samples in the patch. From the figure we see that a 10 dB SNR is sufficient to achieve a STD of the depth-error always less than 5 cm. This is around the desired performance, the problem is to achieve >10 dB at 200 meters range.

To evaluate if the theoretical accuracy is met, we have simulated a flat seafloor with 120000 reflectors scattered randomly out to 200 meters range. AWGN is added to the received time series and the estimated depths averaged over 1000 pings. The results can be found in Figure 7.4. The two upper panels show the coherence and corresponding SNR from the estimates. They also show the results from measuring the SNR directly on the simulated raw-data. A comparison reveals three distinct features: At short range the SNR from the peak correlation coefficient meets a threshold and falls of. This is caused by a combination of interpolation inaccuracies and baseline decorrelation. At medium range the two SNR estimates correspond quite well. At far range the SNR from the peak correlation coefficient falls of again. This discrepancy is yet to be explained.

The third panel in Figure 7.4 shows the STD and RMS error of the sidescan bathymetry depth estimates. These two quantities are almost identical, indicating that there is no significant bias in the estimates. In the same panel the predicted performance from the SNR estimates are included. They show a good correspondence with the estimated error. However, one can notice that the fall in peak correlation at far range increases the predicted STD, but it does not affect the actual performance of the depth estimates. This suggests that the peak correlation coefficient is unreliable when it falls below around 0.7 (which fits well with our correlation threshold value of 0.66).

The lower panel of Figure 7.4 summarizes the sidescan bathymetry performance. It shows that the ratio between the achieved STD of the depth-error and the predicted STD from the peak correlation coefficient is close to one. This means that Equation 7.8 can give a good estimate of the performance of sidescan bathymetry, when the SNR is calculated from the peak correlation coefficient. It also show that there is a significant loss in performance at short range mostly caused by geometrical effects. But since the performance is proportional to range, this effect is less significant for the overall performance.

The results from a large scale simulation of the sidescan bathymetry performance can be found in Figure 7.5. Part of a real AUV trajectory is used in the simulation, which




Figure 7.4: Results from averaging 1000 pings of sidescan bathymetry depth estimates on a flat simulated seafloor with AWGN. A HISAS sonar is used in the simulation. The upper panel shows coherence as a function of ground-range. The solid line shows the peak correlation coefficient from sidescan bathymetry while the dashed line show the coherence calculated from the inserted SNR. The second panel shows the corresponding SNR as a function of ground-range. The third panel shows the STD (solid line) and RMS error (dotted line) of the sidescan bathymetry depth estimates as a function of ground-range. We also show the predicted STD from the peak correlation coefficient (dashed line) and from the inserted SNR (dashed-dotted line). The lower panel shows the ratio between the estimated STD and the predicted STD (from the peak correlation coefficient) (solid line) as a function of ground-range. The dashed line shows the ratio between the estimated STD and the predicted STD (from the inserted SNR).

means that there are variations on all 6 degrees of freedom. The figure shows good performance through the full processing chain, including compensation for attitude and motion and mosaicking each ping into a common coordinate system.

7.2 SAS bathymetry

We start by revising the theoretical resolutions of SAS images. The range-resolution is the same as for sidescan images

$$\delta_r^{\text{sasi}} = \delta_r^{\text{ssi}} = \frac{c}{2B}. \quad (7.9)$$

The along-track resolution is given by the receiver-element size, d

$$\delta_x^{\text{sasi}} = \frac{d}{2}. \quad (7.10)$$

In contrast to sidescan bathymetry, SAS interferometry uses a 2D patch of image pixels for each depth estimate. Since the along-track resolution of the SAS images are range-independent, the bathymetric along-track resolution becomes

$$\delta_x^{\text{sasb}} = w_1 N_{sx} \frac{d}{2} = w_1 P_x, \quad (7.11)$$

where w_1 is a window-factor ≤ 1 and N_{sx} is the number of samples along x in the 2D patch. P_x is the patch size along x in meters. The ground-range resolution is found by the same procedure as in the previous chapter

$$\delta_y^{\text{sasb}} \approx w_2 N'_{sy} \frac{c}{2B} = w_2 P_y, \quad (7.12)$$

where w_2 is a window factor, N_{sy} the number of samples along y in a ground-range image sampled at theoretical range resolution, and P_y the patch size along y in meters. Although N_{sx} and N_{sy} are independent, we use $N_{sx} = N_{sy} = \sqrt{N_s}$. We also use flat weighting, which means that $w_1 = w_2 = 1$. All the SAS bathymetries have a posting separation equal to the input image resolution, which means that the images are over-sampled by a factor of $\sqrt{N_s}$ in each direction. This is done for visualization purposes. If data compression is important, the bathymetries can simply be downsampled by the same amount.

Figures 7.6 and 7.7 show the results from SAS interferometry on a simulated 20 by 20 meters seafloor. Figure 7.6 shows the estimated bathymetry along with the estimation error, while Figure 7.7 show the estimated resolutions. The conclusions are the same as in the previous section. The measured resolutions (using the flank length) is close to, but less than the theoretical resolutions.

In the upper panel of Figure 7.7, notice that the loss of correlation is less along y than along x . This is caused by the seafloor depth stepping down along positive y -direction,






Figure 7.5: The upper panel shows a simulated bathymetry for a point scatter simulation with 120000 reflectors. The simulated trajectory showed in solid black consists of 1000 pings from a true AUV trajectory. The center panel shows the estimated bathymetry using sidescan bathymetry with 200 patches one meter separated. The results are mosaicked into a half-meter grid and areas with no coverage or with a correlation value below 0.66 is colored white. The lower panel shows error in the estimates. Notice that the error increases with range.




(a) Estimated seafloor using SAS bathymetry




(b) Error in the estimated seafloor


Figure 7.6: The figure shows results from SAS interferometry on a 20 by 20 meters simulated seafloor using 13 million point reflectors. On average there are 9 reflectors for each resolution cell, which means that the SAS images have fully developed speckle. The simulated bathymetry consists of four quadrants with half a meter step changes. The upper panel shows the estimated SAS bathymetry. The lower panel shows the error in the estimated bathymetry.



(a) Averaged coherences



(b) Estimated resolution along x



(c) Estimated resolution along y

Figure 7.7: The figure shows detailed results from the estimated bathymetry in Figure 7.6. The upper panel shows the peak correlation coefficients averaged over each direction. The solid line shows the correlation as a function of x -position averaged along y -positions, while the dashed line shows the correlation as a function of y -position averaged along x -positions. The drop in correlation is caused by the step change in the bathymetry. The center panel shows the estimated x -resolutions using flank length. The lower panel shows the corresponding y -resolutions. The theoretical and average resolutions are drawn on top.

which causes an effect similar to layover; a small part of the slant-range data is without echo. Since the simulation is noise-free, the effect of blending two different depths into a single correlation window is reduced. This also causes a underestimation of the y -resolutions, which we see in the lower panel of Figure 7.7. If the seafloor had stepped up instead of down, a layover would have caused mixing of echo from different depths into single samples. The correlation loss would then have been increased instead of decreased, and the y -resolutions would have been overestimated instead of underestimated.

The accuracy of SAS depth estimation is in principle equal to the accuracy of the sidescan depth estimation. However, we use a slightly different expression in this case. Combining Equations 3.24 and 4.40 with $S \approx 0.5$ in Equation 4.40 we get

$$\text{std}\{z_{\text{est}}\} \approx \frac{cr_2}{2D} \frac{1}{2\pi f_0} \sqrt{\int (\theta - \theta_0)^2 p(\theta; k(\rho), N_s) d\theta}, \quad (7.13)$$

where $p(\theta; k(\rho), N_s)$ is given by Equation 4.36.

Figure 7.8 shows results from a simulation of a flat seafloor with ground-range out to 130 meters. (The range is limited due to time consumption of the simulation. A SAS simulation requires a much larger number of point scatterers than a sidescan simulation.) The figure illustrates much the same as Figure 7.4. The coherence in the upper panel falls off at short range due to baseline decorrelation. One can also see straddle loss due to the sampling, from the oscillating shape of the estimated coherence. At longer ranges, the measured and estimated SNRs agree well. The STD and the RMS values for the depth estimation error are virtually identical so we only show the STD in panel three of Figure 7.8. From the bottom panel we observe a small discrepancy between the measured STD and the one predicted by the coherence. Still, the fraction between them is always between 1 and 2 so the accuracy predicted from the measured coherence is also in this case a good indicator of the actual performance. The discrepancy is probably due to high sensitivity of coherence when the SNR is very high.

To verify the performance on a more realistic scenario we simulated a 20 by 20 meters seafloor with 13 million point reflectors and a bumpy seafloor. The trajectory is a section of a real AUV trajectory. The simulated seafloor, the estimated seafloor and the estimation error can be seen in Figure 7.9. The estimation error is largest where the gradient in the bathymetry is largest. This is partly due to the fact that the geometry changes within the estimation window, and to gridding inaccuracies in the simulator and in the calculation of the depth-error.

7.3 Summary

The horizontal resolutions of the estimated bathymetries can be easily deduced from the input image resolutions and the filter sizes. For sidescan bathymetry the along-track resolution is equal to the along-track sidescan image resolution (see Equation 7.3), and the







Figure 7.8: Results from averaging SAS interferometry depth estimates on a flat simulated sea-floor with AWGN. A HISAS sonar is used in the simulation. The upper panel shows coherence as a function of ground-range. The solid line shows the peak correlation coefficient from SAS bathymetry averaged over all x -positions, while the dashed line show the coherence calculated from the inserted SNR. The second panel shows the corresponding SNRs as a function of ground-range. The third panel shows the STD (solid line) of the depth estimation error as a function of ground-range. We also show the predicted STD from the peak correlation coefficient (dashed line) and from the inserted SNR (dashed-dotted line). The lower panel shows the ratio between the estimated STD and the predicted STD (from the coherence) (solid line) as a function of ground-range. The dashed line shows the ratio between the estimated STD and the predicted STD (from the inserted SNR).



(a) Simulated depth



(b) Estimated depth using SAS bathymetry



(c) Error in the estimated depth

Figure 7.9: The figure shows results from SAS depth estimation on a simulated 20 by 20 meters seafloor consisting of 13 million point reflectors on a bumpy seafloor. On average there are 9 reflectors for each resolution cell, which means that the SAS images have fully developed speckle. A HISAS sonar and part of a real AUV trajectory is used to simulate the data. The upper panel shows the simulated bathymetry, the center panel the estimated bathymetry and the lower panel the error in the estimates.

cross-track resolution is equal to the effective ground-range patch-length (see Equation 7.7). For SAS bathymetry, the resolutions are equal to the size of the estimation window in both dimensions (see Equations 7.11 and 7.12). These findings are in agreement with simulation results using flank length as a resolution measure.

The accuracies can be found either from time delay theory or from speckle statistics. For sidescan depth estimation we use time delay theory (see Equation 7.8) and find that the accuracy predicted from the measured peak correlation value corresponds well with the measured accuracy on a simulated dataset. For SAS depth estimation we use speckle statistics and found that the predicted accuracy (see Equation 7.13) is in agreement with the measured accuracy within a factor two, when the measured coherence is used in the predictions (i.e. the measured accuracy is between the theoretical accuracy and twice the theoretical accuracy).

Chapter 8

Results and studies

In this chapter we present analysis and results from selected studies. Some of the material is presented for the first time while other parts are reproductions from previously published material.

8.1 The phase-difference filter size and shape

The vertical accuracy and the horizontal resolutions of SAS bathymetry are dependent on the number of independent image-pixels used to estimate each depth (see Chapter 7). In this section we consider the effect of changing filter sizes in the phase-difference estimator (see Equation 4.29). We also examine the effect of different weights applied in the filter and of down-sampling the resulting depth maps. The results and presentation in this section is based on Sæbø et al. (2007c).

As we discussed in Chapter 7, both the vertical accuracy and the horizontal resolutions of SAS bathymetry are dependent on the number of independent image-pixels used to estimate each depth. Our implementation of the phase-difference estimator is such that the input gridding is not changed. For some applications, the most elevated feature in each resolution cell is more important than the average depth so any down-sampling must retain peaks if they exist in the bathymetry. A technique which keeps the shallowest depth values can then be applied to reduce the sample density to a suitable quantity for charting of large areas.

Depending on the application, different resolutions and different grid samplings may be desired. Resampling can be implemented in different ways at different levels in the processing

- **Pre-processing:** For a critically sampled SAS image one can reduce the resolution by either reducing the wavenumber coverage, or by incoherent averaging independent wavenumber regions or by simply selecting only some of the pixels. The third option is clearly suboptimal since it will produce an undersampled image. The second option can be used to reduce the amount of speckle (Jakowatz et al., 1996, pages 112-121), but incoherent averaging will destroy the phase

in the images. The first option is the best solution, but it should be noted that when considering bathymetric accuracy, it is always better to have as high as possible resolution in the input SAS images and rather perform down-sampling at a later stage. However, reducing wavenumber coverage reduces the computation time of both SAS image formation and of interferometric processing and that may outweigh the disadvantages in some cases.

- **Filtering:** The maximum likelihood phase filter reduces the resolution to the size of the applied filtering window, but retains the input grid sampling (see Chapter 7).
- **Post-processing:** At this step, the interferogram is oversampled with the size of the applied filter. Resampling the interferogram or the bathymetry is almost identical, since the geometrical conversions from phase-differences to depths are close to linear in small regions. Oversampled depth estimates give a smoother visualization of the bathymetry and is often the best choice for smaller regions. For larger regions, we want to reduce the data-amount and down-sample the grid sampling to the actual resolution. In some cases, even further down-sampling may be desired. Again there are several options. For optimum visualization, a simple average filter can be applied. For bathymetric mapping, picking the highest point in each cell may be a more robust solution. More advanced methods may use the SAS reflectivity in the down-sampling.

Theoretically, SAS bathymetry can provide depth maps with image resolution. The phase-difference can simply be estimated for each image pixel-pair without any filtering. Due to the speckly nature of the SAS images this works poorly in practice; the echo strength is random and, assuming a fixed noise level, will cause variations in the SNR. The precision of the depth estimates will therefore vary strongly, causing unacceptable low robustness. To illustrate the effect of different filtering techniques on real data, we have selected data from a simulated bumpy seafloor with random distributed point scatterers. Figure 8.1 shows the estimated bathymetry with and without the phase-filtering. Clearly, the depth maps agree to a large extent, but random pixels seem to fail for the full-resolution case. The simulation is noise-free and therefore also a best case scenario; introducing noise will only increase the problems for the full-resolution case. A noisy seafloor like the one in the lower panel of Figure 8.1 is not for post-processing or 3D-visualization.

To investigate the effect of different filtering techniques we selected an area outside Horten, Norway at around 200 meters water depth. The scene consists of a relatively flat seafloor with a 3 by 10 meters barge located in the middle. The imagery is generated using INS integrated DPCA motion estimates and wavenumber beamforming, and has a final resolution of 2.94 cm by 1.37 cm (along-track by range). Figure 8.2 shows the SAS image, the estimated bathymetry and the coherence using a standard filtering size of 9×9 pixels. Note the lack of coherence in the shadow regions.






Figure 8.1: The figure shows results for a noise-free point scatter simulation of a bumpy seafloor. The upper panel shows the simulated bathymetry as a function of ground-range (y) and along-track (x). The center panel shows the estimated bathymetry using SAS interferometry and a 9×9 pixels maximum likelihood phase filter. The bottom panel shows the same results without the phase filter (i.e. a 1×1 window).



(a) SAS image

(b) Estimated relative depth

(c) Estimated coherence



(d) 3D-rendering of relative depth

Figure 8.2: SAS image, bathymetry and coherence of a 15×20 meters area at around 200 meters water depth outside Horten, Norway. The SAS image is shown with 60 dB dynamic range, the colorcoded bathymetry ranges from 18.2 (red) to 19.6 (blue) meters depth and the coherence ranges from 0 (blue) to 1 (red). The interferogram was estimated using a 9×9 pixels maximum likelihood phase filter. The lower panel shows a 3D rendering of the object with colorcoded depth.

A 9×9 pixels maximum likelihood phase filter gives a reasonable trade-off between resolution, robustness and smoothing. The resulting bathymetry is shown in the upper left panel of Figure 8.3. The full resolution relative depth estimate (pixel-by-pixel phase-difference) is shown in the upper center panel. This map contains too much noise to be useful. Although noise can be filtered well by the human eye, it causes severe problems in the phase unwrapping and large errors in the depth maps. The relative depth based on filtering with a 5×5 pixel window is shown in the upper right panel of Figure 8.3. This allows better resolution than normal filtering, but also noisier estimates. For high SNR, this may prove to be a better solution than the default 9×9 window. An interesting point is that the two shallowest features are both 5 cm higher than for normal filtering. This could either be caused by larger variance in the estimate or it could be actual features which are excessively smoothed using the 9×9 window.

To cover the range of applicable filters we have also applied a 25×25 window. The detail-level seen in center left panel of Figure 8.3 is strongly reduced and small, but strong, scatterers are seen as square shapes with the size of the window. However, the smoothness of the bathymetry is increased due to noise suppression. This makes the bathymetry better suited for 3D-rendering. Note that the shallowest points are both 8 cm deeper than with default filtering, due to reduced noise and more smoothing.

The previous filters were all box-car or rectangular filters. The *shape* of the filter also affect the performance. We implemented filtering using a 9×9 filter with Gaussian shape truncated at 0.5. The result showed in center panel of Figure 8.3 is almost identical to the default filtering of the upper left panel, but has slightly lower SNR and better resolution. The edges in the bathymetry are also smoother, since the filtering window has a smoother cutoff. This is a good alternative to the default filtering. The depth of the shallowest points are unchanged from the default 9×9 filtering.

We also considered post-estimation smoothing; results of which are shown in the center right panel of Figure 8.3. The post-processed smoothing results were generated using default filtering and 19×19 pixels neighboring filtering on top (Gonzales and Woods, 1992, page 191). By comparing this result to the results using a 25×25 window, we see they have comparable resolution in the bathymetry, but the heavy filtered version has more square features caused by the window size. The smoothing reduces the effect of the window, since all points are weighted equally. Another important aspect is that the maximum likelihood filter assumes homogeneity inside the filter window. If the seafloor has sharp features a too large filter will cause decorrelation. If a smooth, oversampled bathymetry is desired, it could be better to apply smoothing instead of increasing the estimation window. The highest points of the bathymetry in the center right panel of Figure 8.3 are 7.5 cm lower than for the default filtering, which is comparable to the heavy filtered version.

Down-sampling the grid spacing of the bathymetry was performed using a 19×19 pixels neighboring filter and is shown in the lower left panel of Figure 8.3. This is the same as the previous case, but without an oversampled result. The visual representation is clearly poorer, but the approach is useful if data reduction is important. The highest points are about 12 cm lower than in the default filtering. An alternative down-




Figure 8.3: Relative depth of the barge using different filtering steps. The colorcoded bathymetry ranges from 18.2 (red) to 19.6 (blue) in all panels. Upper row, from left to right: Phase-filter with 9×9 , 1×1 and 5×5 window. Center row, from left to right: Phase-filter with 25×25 window, Gaussian phase filter with 9×9 window and post-smoothed bathymetry. Lower row, from left to right: Downsampled bathymetry, downsampled bathymetry by picking the most elevated target in each cell instead of averaging, and phase-filter with 9×9 on input images with 10 times poorer resolution in both dimensions.

sampling was tried by choosing the highest point in each cell instead of the average and is shown in the lower center panel of Figure 8.3. Naturally, this has the same highest points as in the default filtering, since the highest points are selected within each down-sampling cell. It can sometimes be important to ensure that none of the elevated targets are removed in the down-sampling.

The pre down-sampling method is shown in the lower right panel of Figure 8.3. This was performed by down-sampling the original imagery with a factor of 10 before processing with a 9×9 window. This means that there is less information and poorer SNR in the images. This effect is visible in the bathymetry, which has additional areas where the correlation is below the correlation threshold of 0.66. Using low-resolution images could be a quick-look option. The highest points are approximately 14 cm lower than for default filtering and can be caused by the fact that small elevated targets are not resolved in the images.

8.1.1 Summary

There is always a trade-off between accuracy and resolution: Larger filters results in more accurate estimates. We have also found that:

- Due to speckle, the size of the phase-difference filter should never be reduced to only a single pixels, even for extremely high SNR cases
- It is better to down-sample the depth-maps than the input images
- Oversampling the depth maps improves visualization
- The size of the phase-difference filter should be large enough to minimize variance, but small enough to ensure constant geometry within the filter

8.2 Performance of CUF in sidescan bathymetry

When using cross-correlation techniques in large baseline direction of arrival estimation, dilation (or scaling) of the signals can reduce the performance. This is a limiting factor in sidescan bathymetry. We introduced the cross-uncertainty function as a possible solution in Section 4.3, and studied its performance on idealized 1D-simulations in Section 4.4. In this section, we compare the performance of CCF and CUF on a simulated scenario and on real data.

The simulated data is a full-scale sonar simulation of 1000 pings with a range of 120 meter. The trajectory is taken from a real AUV navigation solution and the sea-floor is a random realization of a bumpy bathymetry with Gaussian distributed point-reflectors. The simulated sonar is similar to HISAS with 30 kHz LFM pulses centered around 100 kHz. Figures 8.4 and 8.5 show the results from the relative depth estimation using CCF and CUF. We see the following




Figure 8.4: Results from a full-scale sonar simulation. The trajectory consists of 1000 pings of a real AUV trajectory. The seafloor is a random seafloor with Gaussian distributed point reflectors. All panels show the result as a function of pings (abscissa) and range (mantissa). The range is 120 meters. The different panels show the following: (a) simulated relative bathymetry; (b) estimated dilation-factor using co-registered CUF.

- The dilation is in general insignificant at long range but increases strongly at near range.
- The dilation is bathymetry-dependent and reduces the error of a bumpy topography for all ranges.
- The correlation coefficients are higher for CUF than CCF. This holds for all ranges but are most dominant at close range.
- The bias and the standard deviation of the error of the depth estimate is less when using CUF relative to using CCF.

Our CUF estimation algorithm uses the estimated dilation-factor to regrid one of the signals in order to eliminate the relative dilation between them. This increases the correlation and thus the accuracy of the estimator. Figure 8.4(b) shows that the estimated dilation-factors also contains valuable information about the seafloor bathymetry. In Equation 3.14 we calculated the relationship between a relative seafloor depth and a relative dilation-factor. Assuming a flat seafloor, the dilation-factor gives us a direct estimate of the relative seafloor depth, independent of the time delay. But the dilation is also dependent of the slope on the seafloor. An extension to our method would therefore be to first run CUF and use the time delay to get a good estimate of the seafloor depth. Then a prediction of the dilation can be calculated from Equation 3.14 and subtracted from the actual dilation calculated from CUF. The remaining dilation can then be linked to the local slope of the seafloor.

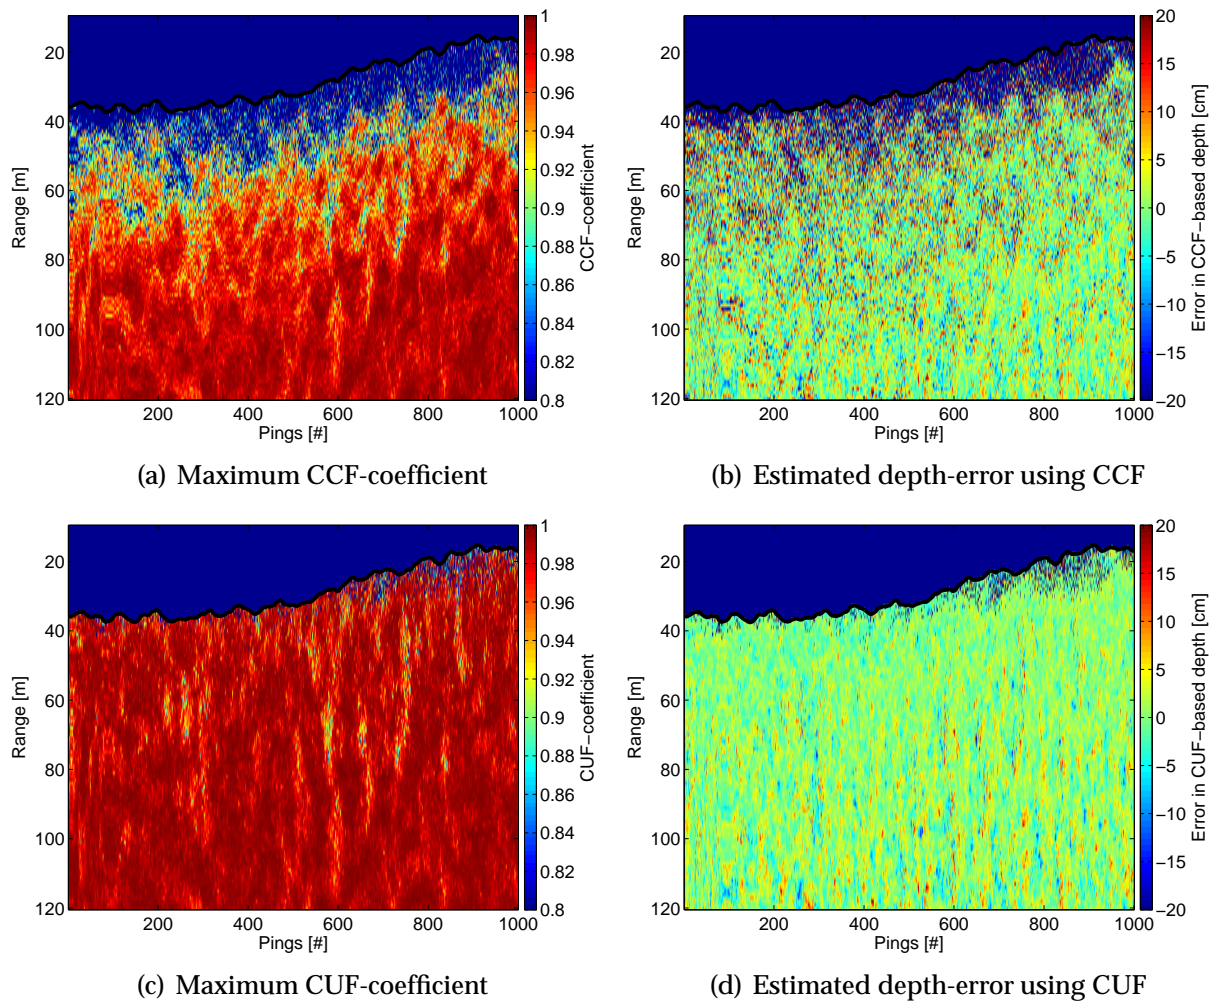


Figure 8.5: Results from a full-scale sonar simulation. The trajectory consists of 1000 pings of a real AUV trajectory. The seafloor is a random seafloor with Gaussian distributed point reflectors. All panels show the result as a function of pings (abscissa) and range (mantissa). The range is 120 meters. The different Panels show the following: (a) and (c) estimated correlation coefficients using CCF and CUF, respectively; (b) and (d) estimated depth-errors for the two methods.

In Figure 8.6 we compare CUF- and CCF-based relative depth estimation on real data collected by the HUGIN autonomous underwater vehicle carrying a HISAS prototype (Hansen et al., 2006). The data was collected in June, 2005 outside Horten, Norway, at around 200 meters water depth. During the collection of these data, the vehicle altitude was 16 meters. The center frequency of the transmitted LFM signal was 100 kHz and the bandwidth 30 kHz. The vertical baseline between the receiver arrays was 28 cm, equivalent to 18.7 wavelengths. We processed the sonar data as dynamic focused sidescan data in slant-range without roll and pitch compensation using a backprojection (or delay-and-sum) algorithm (see Chapter 2). Figure 8.6 shows a clear improvement




Figure 8.6: Experimental results from 688 pings using the HISAS prototype. The data were collected outside Horten, Norway, at approximately 200 meters water depth. All panels show the result as a function of pings (abscissa) and range (mantissa). The range is 90 meters. The different panels show the following (left to right): Upper panels: Estimated bathymetry using CCF and CUF, respectively. Lower panels: Corresponding correlation coefficients. Center panels: Difference in correlation coefficients and dilation-factors estimated using CUF.

in correlation for ranges less than 45 meters. All estimates with correlation value below 0.66 (equivalent to 3 dB SNR) have been discarded. The results indicate an effective range improvement of about 15 meters when using CUF relative to using CCF.

The rock-formation in the center of the scene is better represented by CUF estimates than by CCF estimates (when compared to other benchmark sensors). This is in agreement with our prediction that CUF is better suited for large topographical variations. The dilation factor contains information about the topography. The dilation-factor increases relative to the flat seafloor at the front of the rock-formation and decreases at the back indicating an estimate of local seafloor slope relative to the interferometer.

8.2.1 Summary

In this section we have compared the performance of CCF and CUF in seafloor depth estimation by means of a sidelooking interferometric sonar

Based on Monte Carlo simulations and field-collected data we conclude that using the wideband cross-uncertainty function for slant-range relative depth estimation gives a significant improvement in bathymetric accuracy. Notably, this improvement amounts to as much as an equivalent 20 dB SNR increase for typical sidelooking sonar geometries.

We also demonstrated that the dilation estimation that CUF employs, is justified by developing an improved mathematical model for slant-range interferometry. The simple delay-only model has large approximation errors at short-range and should be avoided. Note that the technique we have proposed in this section is an alternative to seafloor depth estimation using ground-range images (Sæbø et al., 2007b), where a-priori knowledge about the seafloor depth is required.

Although we have limited our work to sidescan sonar, it is equally applicable for synthetic aperture sonar and for synthetic aperture radar. For SAR the geometry is different so the effects of the dilation has to be investigated, but SAS has the same geometry as a sidescan sonar. The conclusions drawn in this section should therefore be directly transferable to SAS interferometry.

8.3 Split-bandwidth interferometry

Phase ambiguities are one of the major problems in interferometric processing and a common solution is to use a 2D phase unwrapper. We have chosen to run a standard Goldstein's Branch cut algorithm in our processing chain (see Section 4.6.8). However, there are still some cases where the unwrapper fails. In this section we present an alternative method for resolving phase ambiguities by dividing the signal bandwidth into sub-bands (Madsen et al., 1993; Bamler and Eineder, 2005). A good approach is to develop an estimator based on the probability density functions for each sub-band having the correct wrap interval. Budillon et al. (2005) suggest to either use a maximum likelihood method by exploiting deterministic contextual information consisting in the

approximation of the depth surface through local planes, or a Bayesian maximum a posteriori estimation technique. Another approach consists of fusing the interferograms from each sub-band with a coherence weighted Kalman filter (Lanari et al., 1996). These methods are outside the scope of this thesis. We describe a deterministic approach to the split-bandwidth phase estimation principle.

A sound implementation of a 2D phase unwrapper works well in cases where the *a priori* ground-range (from for example sidescan bathymetry) is sufficiently accurate, and when the seafloor can be represented by a continuous surface. In cases where there are disconnected regions in the bathymetry (or more accurate; in the interferogram) the unwrapper has problems connecting the regions. Figure 8.7 shows an example where the Goldstein's Branch cut algorithm fails. The figure shows the SAS image, the interferogram and the reconstructed bathymetry of a German WWII submarine at 200 meters water depth. It is clear from Figure 8.7(b) that the interferogram should be unwrapped from the sail and towards the tip of the submarine. But since the interferogram has the same value at the tip of the submarine as on the surround seafloor, the unwrapper also unwraps from the tip towards the sail. Clearly, these two paths will meet at a discontinuity. In 8.7(c) we show the estimated bathymetry after 2D unwrapping. The tip is apparently broken off the rest of the submarine and lies at the seafloor. From the shadow of the submarine in SAS imagery we know that this is incorrect – the submarine should slant upwards all the way to the tip. This is a classical example of the phase unwrapping problem.

The interferometric time delay is defined as (see Section 4.5)

$$\tau = \frac{\theta}{2\pi f_0}, \quad (8.1)$$

where θ is the measured phase-difference and f_0 the center frequency of the signal. Signals with different frequencies should provide the same seafloor depth estimate (ignoring any penetration into the seafloor), and given the same geometry, the same time delay, τ (Xu et al., 1994). This means that the interferometric phase-difference is a function of frequency only, $\theta(f_0) = 2\pi f_0 \tau$. Given two different frequencies, f_1 and f_2 , the following applies


$$\theta_1 = 2\pi f_1 \tau, \quad (8.2a)$$

$$\theta_2 = 2\pi f_2 \tau. \quad (8.2b)$$

This means that $\theta_1/f_1 = \theta_2/f_2$. The problem is that both θ_1 and θ_2 are known only modulo 2π , so the equation we have to solve is


$$\frac{(\theta_1 + m2\pi)}{f_1} = \frac{(\theta_2 + n2\pi)}{f_2}, \quad (8.3)$$

where m and n are unknown integers. We are then left with two unknowns and one equation. Assuming k frequencies would have provided $k - 1$ equations and k unknowns, so the problem is always under-determined.




(a) SAS image

(b) Interferogram

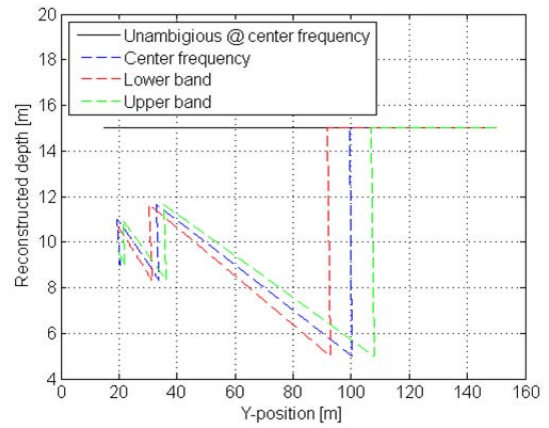


(c) Estimated bathymetry


Figure 8.7: The figure illustrates the problem of robust 2D phase unwrapping. The left panel shows a SAS image of the sail and forward part of the German WWII submarine U-735. The length of the visible part of the submarine is around 40 meters. The data were collected using HISAS on the HUGIN 1000-MR, at around 200 meters water depth. The upper right panel shows the corresponding interferogram (notice that the cut-out is slightly different). The lower panel shows a 3D representation of the estimated seafloor depth after unwrapping the interferogram with the Goldstein's Branch cut algorithm. Notice that the front part of the submarine seems to be broken off the rest.




(a) Phase-difference



(b) Reconstructed depth



(c) Ambiguous depths from lower band



(d) Ambiguous depths from upper band

Figure 8.8: The figure shows the results from a numerical study on the effect of phase ambiguities. The upper left panel shows the theoretical phase-differences for a HISAS sonar observing a flat seafloor 15 meters below the sonar. The center frequency is 100 kHz and the sound speed 1500 m/s. The unambiguous phase-differences are shown in solid black. The ambiguous (or wrapped) phase-differences for the full frequency band are shown in blue. The ambiguous phase-differences for the lower and upper half of the frequency band are shown in red and green, respectively. The upper right panel shows the corresponding reconstructed depths. The lower left panel shows the reconstructed depths for the lower half of the frequency band after adding different 2π cycles. The lower right panel shows the same for the upper half of the frequency band.

The solution is to observe that a 2π wrap will give a different contribution on each side of the equation. In other words, if Equation 8.3 apply for \hat{m} and \hat{n} , then it does not have a solution for $m = \hat{m} + 1$, unless $n = \hat{n} + f_2/f_1$. For $m = \hat{m} + 2$, n must be equal to $\hat{n} + 2f_2/f_1$, and so on. By observing that the first incorrect solution occurs when lf_2/f_1 is integer, where l is an integer different from zero, it should be sufficient to only allow a few 2π to be added to each interferogram. Clearly, without this restriction there is an infinite number of solutions. For HISAS, $f_1 = 92.5$ kHz and $f_2 = 107.5$ kHz, which means that lf_2/f_1 is close to integer for $l = 6$. In our implementation, n and m are restricted to be within ± 3 . An unambiguous phase can then be reconstructed as long as the interferograms are no more than 3 times 2π cycles wrong to start with. In practice m and n are found by minimizing the function

$$c(n, m) = \frac{(\theta_1 + m2\pi)}{f_1} - \frac{(\theta_2 + n2\pi)}{f_2}, \quad |m| \leq 3 \quad |n| \leq 3 \quad (8.4)$$

In Figure 8.8 we show the results from a simple numerical study. The upper left panel shows how the wrapped phase-differences differ for the two frequency bands. The upper right panel shows that this also causes a difference in the depth estimates. In the lower panels of Figure 8.8 we show all depth estimates when adding 2π cycles. The correct depth (and the correct value for n and m) is found for each y -position by finding the minimum difference between all depth estimates in the bottom left panel and bottom right panel of Figure 8.8(d).

Our method of split-bandwidth phase estimation is very simple, but have some drawbacks in practice:

- The sub-band images have reduced resolution compared to the original images. This also reduces the final bathymetric resolution. A solution is to use the split-bandwidth method to find the correct 2π interval for the full bandwidth interferogram. Problems can appear in the up-sampling of the split-bandwidth estimates when m or n differ between neighboring pixels.
- The frequencies of each sub-band have to be accurately known. In our experience, the band-pass filtering caused unsymmetrical sub-bands, due to a taper on the original frequency band. The effective frequencies of the sub-bands therefore had to be estimated by a center-of-mass method, a method which itself has limited accuracy.
- Phase-difference estimates with low SNR and high variance may be corrected with a random number of 2π , independent of the neighboring pixels.

In Figure 8.9 we have tested the split-bandwidth method on a simulated bumpy seafloor. Figure 8.9(a) shows the simulated seafloor. In order to provoke a failure in the phase unwrapping algorithm we used a poor *a priori* ground-range (which means that the SAS images were poorly co-registered). Figure 8.9(b) shows the estimated seafloor depth after phase unwrapping the interferogram. The depth estimates are consistent,

but the whole area has been unwrapped into the wrong 2π interval. Figure 8.9(c) shows the corresponding result after split-bandwidth phase estimation. The original bandwidth of the SAS images was 30 kHz around 100 kHz. After bandstop filtering, two sub-bands of 15 kHz were used in the absolute phase estimation. Similar results were achieved by filtering the raw-data instead of the images. Figure 8.9(c) shows a clear improvement over the standard phase unwrapper (notice that the offset of the estimated bathymetry is caused by the poor *a priori* ground-range). However, as predicted, random estimates seem to fail severely. In this case, these random errors could be removed by running the Goldstein's Branch cut algorithm on top of split-bandwidth phase estimation (see Figure 8.9(d)).

A typical scenario where standard unwrappers fail, are on large man-made objects like a wreck. Figure 8.10 shows a SAS image of a 40 meters long wreck. Figure 8.11(a) shows the depth estimates of the same area. Notice the probable failure in the depth estimate in the upper right part of the image. Figure 8.11(b) shows the corresponding result using the split-bandwidth method. The wrap problem is eliminated, but instead random errors can be seen across the scene. These failures are due to the pixel-based implementation of split-bandwidth phase estimation. Treating each pixel as deterministic and independent variables is an invalid assumption and will, as shown, fail randomly.

8.3.1 Summary

Although standard phase unwrappers perform very well in most scenarios, they still have problems with disconnected regions, like large man-made objects. One approach is to use bandwidth to resolve the phase ambiguities. In this section we have described and tested a very simple approach which filters the SAS images into sub-bands and compares the sub-band interferograms pixel for pixel in order to find the correct 2π . Although the method and the implementation are simple and assume deterministic and independent estimates, it still showed promising results on regions where the Goldstein's Branch cut algorithm failed. However, the estimated depth maps suffered from random errors. A more thorough study should try to expand to for example one of the methods described in Budillon et al. (2005).

8.4 Filtering of the depth maps

Interferometric synthetic aperture methods benefit from the coherent properties of speckle. However, low intensity pixels inherent in speckle suffer from low signal-to-noise ratio, rendering full-resolution phase-differencing unpractical. A common solution is to use a maximum likelihood phase-difference (MLP) estimator on a few neighboring pixels (Hanssen, 2001, page 96), compromising between phase accuracy and horizontal resolution (see Section 8.1). In this section we discuss filtering of depth maps after applying the MLP estimator. The results and presentation in this section is based on Sæbø

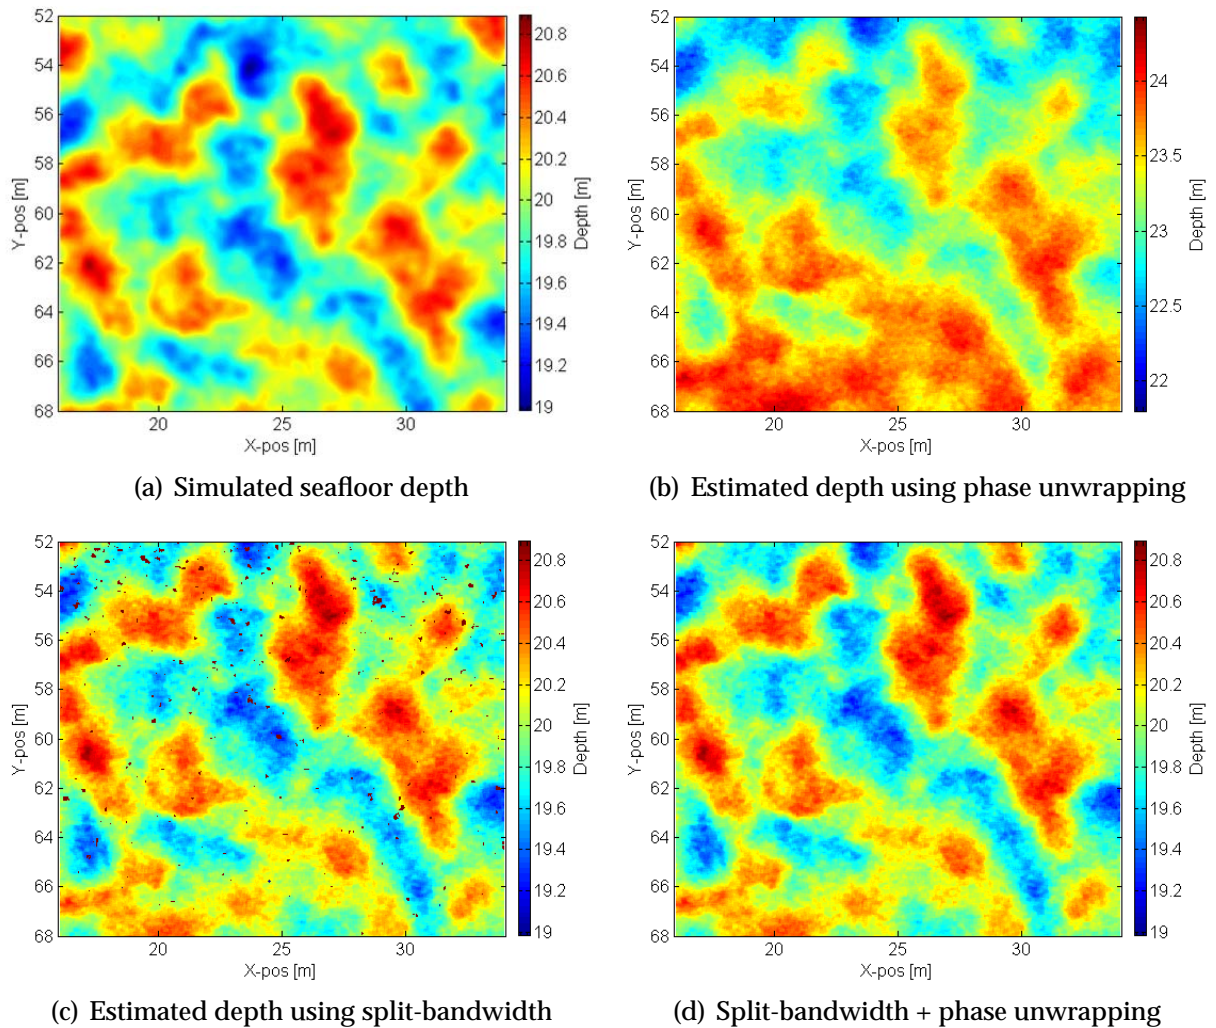


Figure 8.9: The figure shows the results from a point scatter simulation of a bumpy seafloor. The depth of the simulated seafloor is shown in the upper left panel. The reconstructed seafloor depth using phase unwrapping (and a poor *a priori* ground-range) is shown in the upper right panel. In the lower left panel we show the reconstructed seafloor depth using the split-bandwidth method. In the lower right panel we show the results using split-bandwidth followed by phase unwrapping.

et al. (2009).

Within the MLP-filter, the phase-differences are assumed to be homogeneous, which sets an upper limit to the size of the filter. Exceeding this size causes a drop in coherence, which again decreases the accuracy of the phase-difference estimates. This means that the MLP-filter is unable to decrease the variance of the depth estimates to a desired level.

In this section we describe and demonstrate how the interferometric depth estimates can be filtered in order to decrease the variances. Instead of applying a filter on the


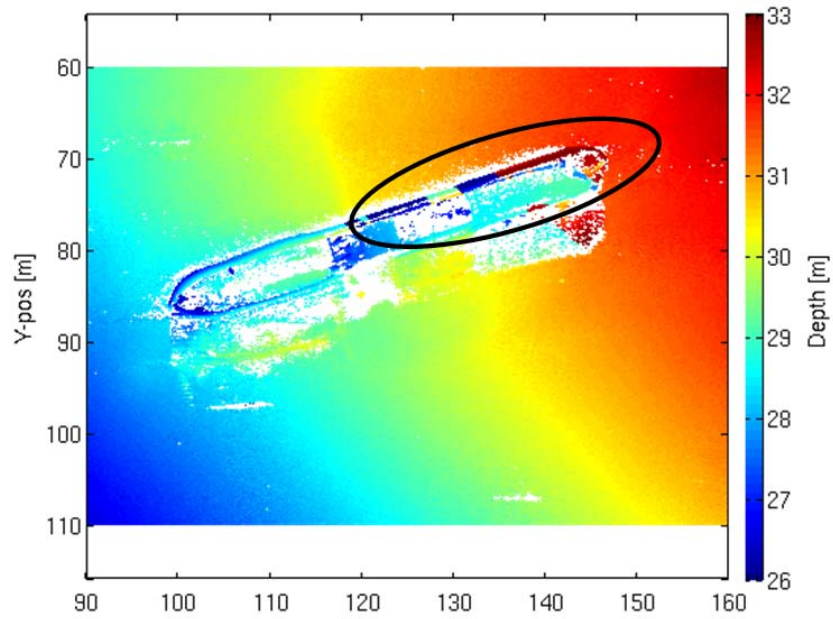


Figure 8.10: The figure shows a 60 dB SAS image of a 40 meters long wreck. The data were collected using HISAS on the HUGIN 1000-MR. The wreck is elevated 2-3 meters above the seafloor at around 30 meters water depth.


complex interferometric images we filter the depth estimates in post-processing. This removes the need for homogeneous phase-differences within the filter window.

We have compared a weighted smoothing (WS) filter (Brownrigg, 1984), a weighted median (WM) filter (Brownrigg, 1984) and a weighted version of the bilateral (WB) filter (Tomasi and Manduchi, 1998), and found that the WB-filter performs best. The WS-filter degrades the sharp transitions to an unacceptable degree while the WM-filter fails to smooth flat areas sufficiently.

A very useful quantity of the interferometric processing is the coherence of the estimated bathymetry (Hanssen, 2001, pages 96-98) (see Section 4.7). The coherence can be converted to an equivalent signal-to-noise ratio (see Section 4.1.3). Both the coherence and the SNR can be employed as input weights to the filters we present. We have found that SNR is best suited since it better separates weights at high SNRs. In addition, we remove low SNR estimates completely by thresholding the correlation coefficient at $2/3$.



(a) Estimated depth using phase unwrapping



(b) Estimated depth using split-bandwidth

Figure 8.11: The figure shows depth estimation results for the same area as in Figure 8.10. The upper panel shows the estimated seafloor depth using standard phase unwrapping and the lower panel shows the corresponding result using split-bandwidth. In the upper panel a phase unwrapping failure is highlighted by a black circle.

We define a set of filtering weights, w

$$w = \begin{cases} \rho & : k \geq 2/3 \\ 0 & : k < 2/3 \end{cases} \quad (8.5)$$

In addition we have tested the square root of SNR and the logarithm of SNR, but they both performed poorer than SNR itself.

8.4.1 Maximum likelihood phase-difference filter

The variances in the phase-difference estimates can be reduced by increasing the size of the Maximum likelihood phase-difference filter. The phase-differences are estimated by (see Section 4.5)

$$\theta = \angle \left\{ \frac{\sum_{i=1}^N a[i]b^*[i]}{\sqrt{\sum_{i=1}^N |a[i]|^2 |b^*[i]|^2}} \right\}, \quad (8.6)$$

where $a[i]$ is the i 'th pixel of the master image and $b[i]$ is the corresponding pixel in the slave image. The relative depths, z_1 are found by converting the phase-difference to a ground-range translation and applying Equations 3.20. This filter is always used in the interferometric processing, independent of subsequent filtering. The advantage with this filter is that it is a maximum likelihood filter, but it assumes homogeneity within the filter size.

8.4.2 Weighted smoothing filter

The weighted smoothing filter is a weighted neighboring filter applied on the relative depth estimates,

$$z_2 = \frac{\sum_i w_i z_i}{\sum_i w_i}, \quad (8.7)$$

where w are the weights defined in Equation 8.5 and z is the original depth estimates. The weighted smoothing filter performs well on slowly varying seafloors but it degrades edges and small objects.

8.4.3 Weighted median filter

The weighted median filter is a standard edge-preserving filter (Gonzales and Woods, 1992, pages 191-195). Contrary to traditional weighted median filters, our version is implemented with adaptive, non-integer weights. First the window depth values are sorted in ascending order. Using this sample order, the smallest index is found whose

cumulative weight is equal or greater than half the total weight sum. The corresponding depth is the weighted median value.

$$z_3 = z_k, \quad \text{where} \quad \min_k \left\{ \sum_{i=1}^k \geq \frac{1}{2} \sum_{i=1}^n \right\}, \quad (8.8)$$

where n is the number of window pixels, z are the original depths and w are weights sorted in ascending depth. The advantage of the weighted median-filter is that it removes wild points while preserving edges. The disadvantage is that it has an unpredictable statistical behavior and it does not smooth slowly varying regions sufficiently.

8.4.4 Weighted bilateral filter

The weighted bilateral filter is the most complex filter. It consists of a Gaussian smoothing in both horizontal and vertical direction. Thus each sample is weighted according to three criteria: The horizontal distance from the center sample, the difference in value (vertical distance) from the center sample, and the weight

$$z_4 = \frac{\sum_i G(d_i)G(v_i)w_i z_i}{\sum_i G(d_i)G(v_i)w_i}, \quad (8.9)$$

where d are the horizontal distances, v the differences in value and G the Gauss-function with standard deviation σ

$$G(x) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2}{2\sigma^2}\right). \quad (8.10)$$

The weighted bilateral filter smooths slowly varying regions while preserving edges and therefore provides high resolution on sharp objects and less resolution on a flat seafloor. These properties are both desirable for visualization of 3D surfaces.

8.4.5 Experimental results

We have tested the different filters on a selection of experimental data collected with HISSAS. The original bathymetry often shows an unnatural large variability on flat regions of the seafloor, caused by a marginal SNR in the interferometric processing. On rocks or man-made objects, the reflectivity and thus the SNR are usually higher. Therefore the accuracy of the estimated shape of objects can be much higher.

An ideal bathymetric filter should smooth slowly varying parts of the seafloor and preserve edges, while using the coherence (or SNR) as weights to limit noise contributions. One approach is to change the size of the MLP filter adaptively (see Section 8.1).

Figure 8.12 shows an example of the estimated bathymetry on a small part of a 30 meters long wreck. Notice two different sections: The flat seafloor at the left and the

small bridge at the upper right. In the original bathymetry the bridge is quite sharp, but the seafloor is contaminated by noise. After applying the SM-filter, the flat seafloor is smoother, but the bridge is rounded at the edges. The WM-filter does a better job of preserving the edges of the bridge, but it does not smooth the flat seafloor sufficiently. By applying the WB-filter, we clearly see that the seafloor is smoothed while the edges of the bridge are as distinct as in the original bathymetry.

In Figure 8.13 we show a patch of seafloor with small rocks. The WS-filter (center panel) actually removes some of the small objects. These objects are quite easy to detect in the bathymetry after WB-filtering. WM-filtering on this scene does not contribute to a smoother surface and is omitted from the figure. On all other scenes we have tested we have found similar results: The WB-filter gives the best trade-off between smoothness and edge preservation.

8.4.6 Summary

Due to the speckle in the SAS images, interferometric SAS estimates will have large variability. Increasing the size of the maximum likelihood filter used in the interferometric phase-difference estimates will reduce the variability at the cost of reduced horizontal resolution. This assumes that the interferometric data are homogeneous inside the filter size. We have presented three different filters which can be applied in post-processing directly on the bathymetry. The filters do not require homogeneity within the filter window. We have applied the different filters on experimental data collected by HUGIN carrying HISAS. We found that the weighted bilateral filter achieves the best performance in smoothing slowly varying areas while preserving edges of distinct objects.

8.5 Comparison of HISAS and EM 3000

High resolution seabed mapping is traditionally performed using multibeam echo sounders (Lurton, 2002, pages 268-282). Since individual beams are formed and steered in different directions, layover effects are rare. Also, the relatively steep observation geometry, makes the MBEs less susceptible to shadow, surface multipaths and errors in the sound speed profile. However, the area coverage rate is limited since the swath is only around four times the altitude. The coverage rate can be increased by running at larger altitude, but the horizontal resolutions of the estimated depth maps are also proportional to range.

Sidescan bathymetry or seafloor depth estimation using interferometric side looking sonar is an alternative technique to MBEs. Lurton (2000) showed that the acoustical measurement precision of sidescan bathymetry is sufficient for seafloor depth estimation, as long as the system is reasonably designed and the footprint shift is accounted for. Other problems, like the phase ambiguities, can be solved by unwrapping techniques (Ghiglia and Pritt, 1998) or by using more than two vertical receivers (Bird and Mullins, 2005). Sidescan bathymetry on AUVs can provide high area coverage rate,





Figure 8.12: Upper left: SAS image of a 30 meters long wreck at 340 meters water depth. Upper right: Corresponding coherence. The four lower panels show a cut-out of the estimated bathymetry. Upper left: No filter; Upper right: Smoothing filter; Lower left: Median filter; Lower right: Bilateral filter.




Figure 8.13: Estimated bathymetry of a 20 by 20 meters rocky seafloor at approximately 30 meters water depth. The depth varies from 25 meters (red) to 35 meters (blue). Upper panel: Unfiltered bathymetry. Center panel: Smoothed bathymetry. Lower panel: WB filtered bathymetry. Note that the original bathymetry has an unnatural large variability caused by marginal SNR. A WS filter reduces this effect, but also smooths real objects like the two rocks in the lower part of the scene. The bilateral filter gives a favorable trade-off between a smooth seafloor and distinct objects.



Figure 8.14: Picture of a HUGIN 1000-MR with a starboard HISAS (left part of picture) and an EM 3000.

with a resolution comparable to MBEs (Hegrenæs et al., 2010). The area coverage rate is, however, diminished due to the blind zone at nadir (Hagen and Hansen, 2007).

The HUGIN 1000-MR is equipped with an EM 3000 multibeam echo sounder (EM3, 2003) (see Figure 8.14). In this section we compare the seafloor mapping ability of the EM 3000 with HISAS. The results and presentation in this section is based on Sæbø and Langli (2010).

The HISAS and EM 3000 sensors are operated simultaneously. The accuracies and resolutions of the estimated depth maps are dependent on the sensor designs and the signal-to-noise ratio of the received signals, but there are some external factors which limit the absolute depth measurements for both sensors

- An error in the AUV position translates to a *bias* in the depth measurement.
- Errors in the AUV attitudes cause a rotation of the depth measurement which gives a *linear range-dependent error*.
- A timing-error between the navigation and the sonar causes an error which is dependent on the high-frequency AUV motion. The HUGIN AUV has for example a high-frequency roll variation which induces artificial *ripples* in the bathymetric maps for incorrect timings.
- An error in the average sound speed scales the range and effects both position and depth of the measurement. This causes a *linear range-dependent error*, which is separable from attitude errors by comparing the different sides.
- An incorrect sound speed profile causes a *non-linear range-dependent error*.




Figure 8.15: Sketch of the sensor geometry on board the HUGIN AUV. The orange circle illustrates a cross-section of the AUV, the gray area the HISAS swath and the green and red areas the EM 3000 swath. While HISAS uses time-difference of arrival (TDOA) to locate the seafloor depth, the EM switches between time of arrival (TOA) (red area) and TDOA (green area).

- An error in the internal lever-arms (e.g. the distance from the navigation unit to the sensor) will cause an *attitude dependent bias*.

It is fairly easy to calibrate range-dependent errors by inspecting the estimated bathymetries of overlapping regions in a lawn-mower pattern. A timing error can be detected by correlating along-track seafloor variations with short-time AUV attitudes. The most challenging errors are errors in the AUV position and internal lever-arm errors.

8.5.1 The EM 3000

The EM 3000 (EM3, 2003) is a 300 kHz high-resolution multibeam echo sounder for seabed mapping and inspection (Lurton, 2002, pages 268-282). Its depth rating and physical dimensions makes it well suited for use on HUGIN AUVs. The swath-width of 130 degrees provides ground-range coverage up to four times the sensor altitude. The angular resolution is obtained using the Mills cross principle (Urlick, 1983, page 62), and the beamwidths of the 127 beams ranges from 1.5 degrees at the vertical to 3.5 degrees at ± 65 degrees. For beams having close to normal incidence, the range is typically determined by a center of gravity amplitude detection principle (or time of arrival, see Figure 8.15). For more horizontal geometries, the echoes from the seafloor are less well-defined and amplitude detection can only provide coarse range estimates. Range at the beam pointing angles are then found using the split-beam technique (Burdic, 1984, pages 327-331): the receiver array is divided into two subarrays, for each subarray a beam is computed, and range is found comparing the phase of the two half-beams and identifying the instant of zero phase-difference. The last step is almost identical to the interferometric processing on HISAS. With good external sensors, the system delivers bathymetry with precision of 5 cm RMS in shallow water.

When HISAS is operated in SAS mode the sampling criterion limits the speed of the AUV. A typical velocity is around 2 m/s, which at 20 meters altitude gives a theoretical area coverage rate of 0.6 km²/h for the EM. However, when SAS processing is not en-

abled, the speed of the vehicle can be increased. For the HUGIN 1000-MR the maximum operating speed is around 3 m/s, which at 20 meters altitude results in a mapping rate of $0.9 \text{ km}^2/\text{h}$. Additional increase in the rate can be made by increasing the altitude. However, at some point the resolution becomes too poor, the range too large or the ping spacing too separated. A reasonable altitude is 40 meters which gives the EM an area coverage rate of $1.7 \text{ km}^2/\text{h}$. In comparison, a typical area coverage rate for HISAS is $1.8 \text{ km}^2/\text{h}$ (Hagen and Hansen, 2007).

8.5.2 Large scale comparison between EM 3000 and HISAS

In Figure 8.16 we present results from a shallow water survey in January 2009. The data are collected outside Horten of Norway, where the water depth varies between 10 and 20 meters. The AUV has run in a lawn-mower pattern seven meters above the seafloor. In the figure we show the aligned bathymetry for the EM 3000 and for HISAS. We also show the differences between the two depth maps.

The difference in area coverage rate is clearly seen by comparing the upper panels of Figure 8.16. In addition to the increase in the surveyed area, HISAS has been run with heavy overlap and each area of the seafloor is observed four times. This is done to ensure full coverage with the EM.

The bathymetric maps in Figure 8.16 seem to agree very well, but the lower panel of the figure shows that there are differences. There is both a small bias of 15.7 cm and a small systematic discrepancy. We have seen similar bias at deep waters with flat seafloor and a constant sound speed profile and believe that the bias is caused by an erroneous lever-arm on one of the systems. This theory is strengthened from observations where the apparent bias changes when the vehicle pitches.

The systematic error (seen as a cross-track pattern in the lower panel of Figure 8.16) is less dominant than the bias and can either be caused by a lever-arm error or by an incorrect sound speed profile. The estimated sound speed profile varies with approximately 0.1 percent from the position of the AUV down to the seafloor. A linear error in the profile of 0.01 percent will cause a depth-error of 5 cm at 85 meters range. The standard deviation of the differences between the bathymetric maps is only 1.7 cm, so ignoring the bias, the two sensors produce almost identical depth maps. The difference in frequency gives a different penetration level, so for some types of seafloor the sensors should disagree. Also, the resolution difference will contribute to a small standard deviation in the differences.

In Figure 8.17 we investigate the two bathymetries in Figure 8.16 in more detail. The upper right panel shows a horizontal slice through the two depth maps and the lower panel shows a vertical slice. The location of the slices are shown in the upper left panel. Both slices illustrate that the agreement between EM 3000 bathymetry and HISAS bathymetry is excellent, after accounting for the bias of approximately 15 cm.

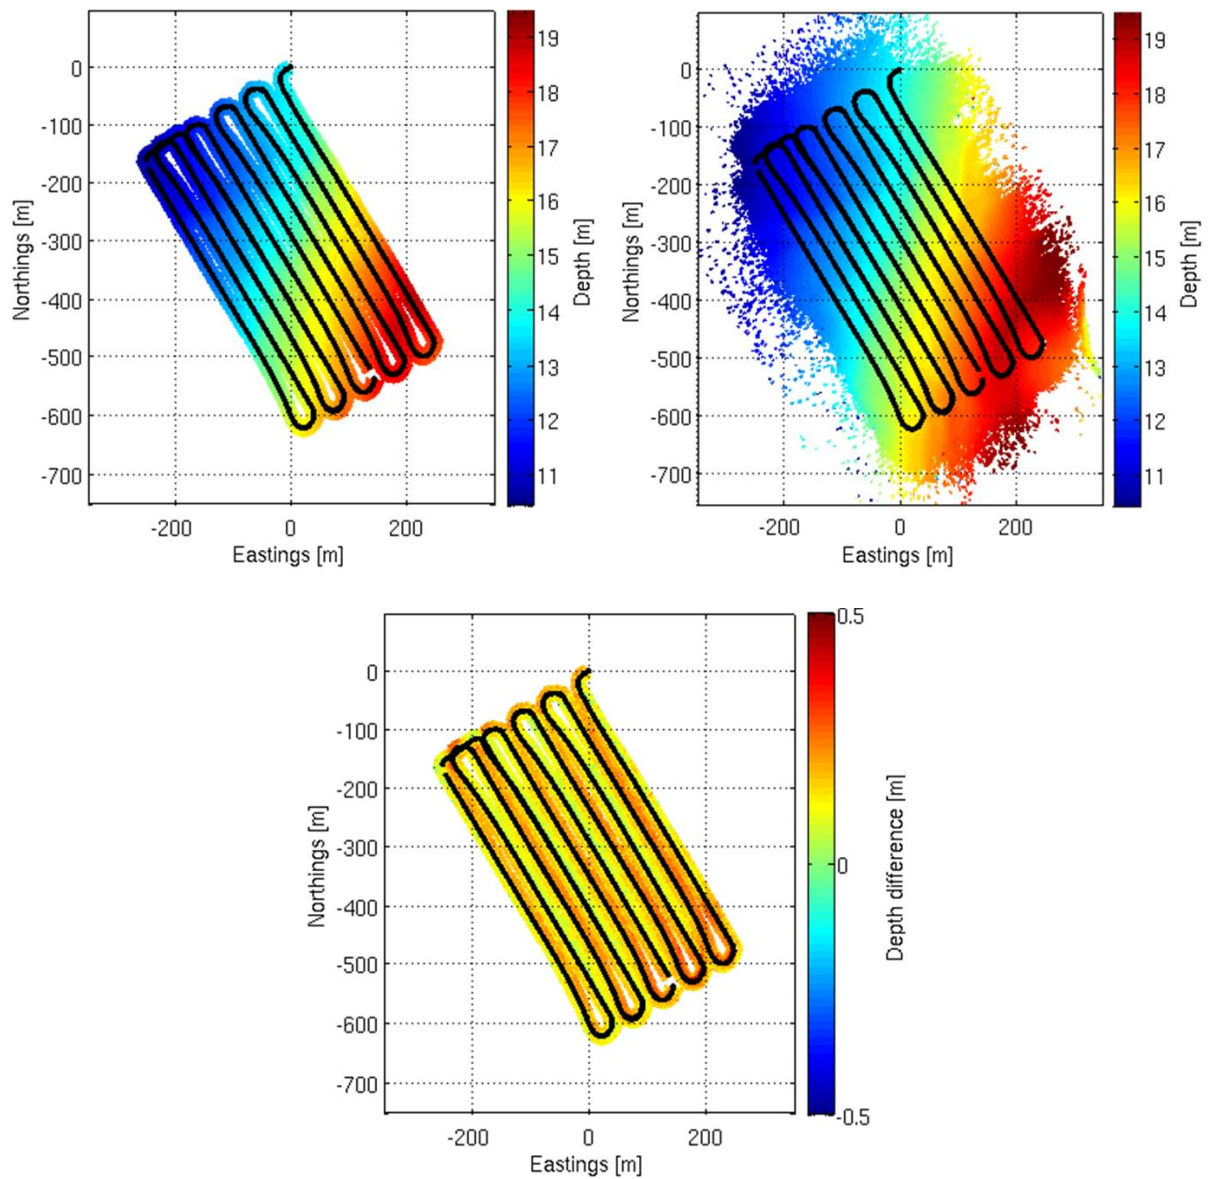


Figure 8.16: Comparison of EM 3000 depth estimates and HISAS depth estimates for a shallow water survey outside Horten in January 2009. The upper left panel shows the geo-referenced depth measurements using the EM 3000. The black line shows the AUV trajectory. Notice that there are small gaps in the depth map due to the limited swath-width of the EM 3000. The upper right panel shows the estimates using HISAS, collected simultaneously. Notice that the coverage is significantly increased. The lower panel shows the difference between the two depth maps. The average difference is 15.7 cm while the standard deviation of the difference is 1.7 cm. Notice that there is a small but systematic cross-track discrepancy.

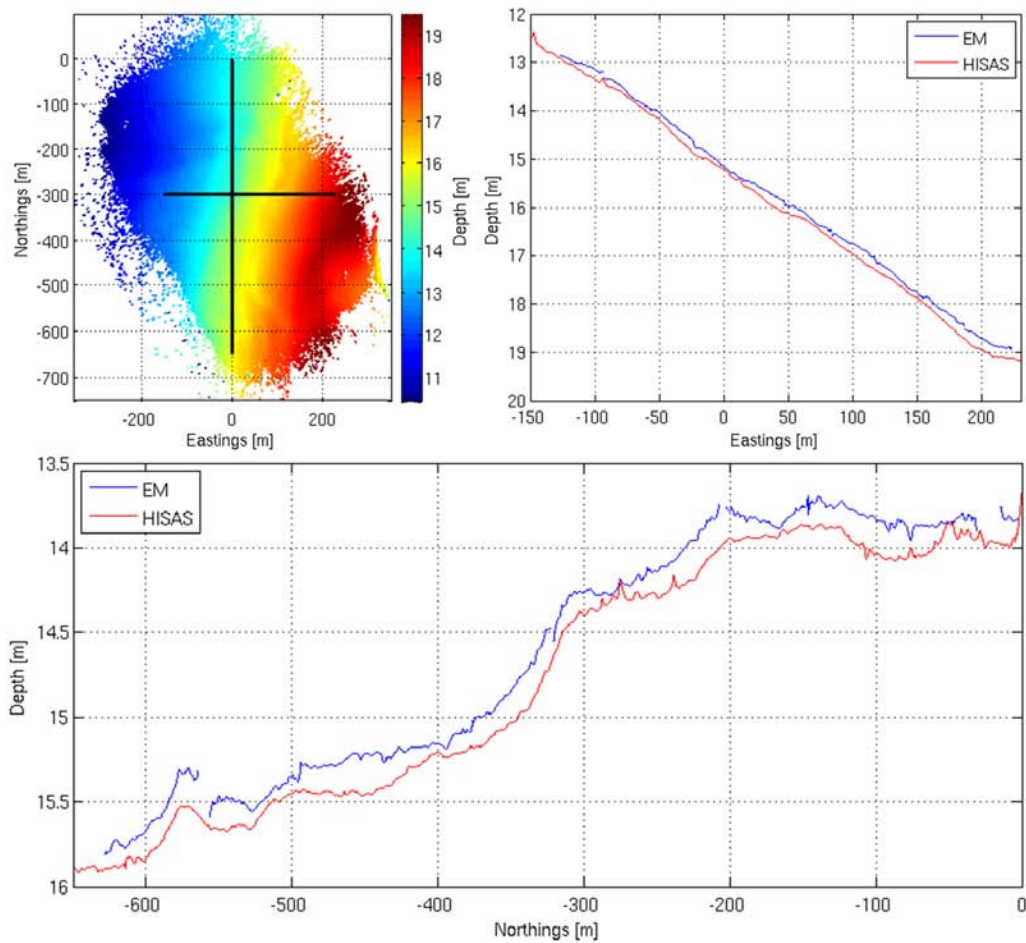


Figure 8.17: Details of the comparison presented in Figure 8.16. In the upper right panel we show the estimated geo-referenced depths at the positions illustrated by the horizontal black line in the upper left panel. In the lower panel we show the same comparison at the positions of the vertical black line in the upper left panel. There is a small, but clear bias between the EM and HISAS, which gives the difference between the blue and red curves.

8.5.3 Small scale comparison between EM 3000 and HISAS

In this section we focus on mapping of small objects. We have chosen a deep water dataset collected in August 2008 outside the western coast of Norway. Geo-referenced depth maps of a 30 by 30 meters area have been calculated using EM 3000, HISAS sidescan bathymetry and HISAS SAS bathymetry. The area contains three torpedo nets from World War II, which are around 5 by 5 meters in extension and elevated a meter over the seafloor. Figure 8.18 shows an overview of the bathymetry, made using HISAS in sidescan mode. At around position (175, 250), a white square indicates the 30 by 30 meters area we focus on. While the EM 3000 only covers the area once, HISAS observes the same area from the four nearest survey lines. Therefore, HISAS can provide four

individual depth maps of the same area.

EM and HISAS sidescan bathymetry The upper panels of Figure 8.19 shows the depth maps from the EM and from sidescan bathymetry, for the region of interest. There is only a small bias of 3.6 cm between the maps and the standard deviation of the difference is only 1.0 cm. There are some noticeable differences:

- The resolution of the EM bathymetry is slightly better than HISAS sidescan bathymetry. This difference is partly due to the fact that sidescan bathymetry is based on all passings and combines short-range and long-range data. If the geo-coding algorithm for HISAS had been set to select the measurement with shortest range instead of using all available measurements, the resolution would have been more similar. Any navigation error between the lines will also degrade the resolution of the depth map. Finally, a reduction of the length of the correlation window will increase the resolution in HISAS bathymetry.
- HISAS bathymetry is less noisy and less detailed. This is related to the above effect; the more averaging, the less noise and the smoother the surface.




Figure 8.18: Sidescan bathymetry measurements using HISAS in deep waters outside the western coast of Norway in August 2008. The black lines show the AUV path, while the horizontal and vertical white line illustrate positions analyzed further in Figure 8.21. In this section we focus on the 30 by 30 meters area enclosed by the white square (at the intersection between the two white lines).




Figure 8.19: Details of the area highlighted in Figure 8.18. The upper left panel shows a depth map using EM 3000 measurements and the upper right panel the same using HISAS in sidescan mode. The depths are shown in a common grid. The mean difference between the depth maps is 3.6 cm while the standard deviation of the difference is 1.0 cm. The center left panel shows the SAS image for pass #1 and the lower left panel the SAS image for pass #4. The corresponding right panels show the depth maps estimated using SAS interferometry. The bias between the EM depth and the depth from pass #1 is -32 cm, while the bias between EM and pass #4 is 52 cm. All the biases are subtracted in the figure.

HISAS SAS bathymetry The lower left panels in Figure 8.19 show SAS images from two sides of the area, both with a range of above 100 meters. The lower right panels show the corresponding SAS bathymetries. From the SAS images it is clear that the acoustic echo looks quite different from each side. This is due to the small depression angle, which causes shadowing and layover effects. This is also reflected in the SAS bathymetries. Different parts of the objects are mapped for each side and to avoid holes in the bathymetry we have to combine the bathymetries from the two sides.

The SAS bathymetries in Figure 8.19 have a bias relative to the EM bathymetry of -32 cm and 52 cm, respectively. To investigate if these biases are related to the SAS interferometry method, we have made depth maps from the same passings using sidescan bathymetry. In Figure 8.20 we show the individual depth maps from the four passings. In the figure, all biases are subtracted, but relative to the EM the biases are -34 cm, -4 cm, 7 cm and 47 cm for passings #1-4. These observations are presented more clearly in Figure 8.21 where two cross-sections of each depth map are shown. Apparently, there is a drift in the estimated depths using HISAS, relative to the depth map from EM. We have considered some of the possible causes for this drift:

- A roll error is unlikely since this would have been visible as linear errors in the right panel of Figure 8.21.
- An error in the average sound speed speed is unlikely since this also would have been visible in the right panel of Figure 8.21.
- An incorrect sound speed profile is unlikely since the sound speed is more or less constant at 330 meters water depth.
- It is not a timing-error or an internal lever-arm error since these errors do not cause a drifting depth-error.

There is one plausible explanation for the biases: An erroneous drift in the estimated AUV-position. In the region of interest at the time of the data-collection the tide can explain some of the drift, but not all. An error in the depth of the AUV of the relevant scale is, however, within the reported accuracy of the navigation system. A thorough investigation of the depth accuracy of the AUV and other causes for the drift in the depth are outside the scope of this thesis.

Ignoring the discussed biases, Figure 8.20 shows an excellent agreement between the individual depth maps. There are visible differences, but this should be expected considering that the views are from two different sides and two different ranges. The positions of the objects seem to be quite similar, indicating that there is no apparent horizontal drift in the AUV navigation.

8.5.4 Summary

This section is a summary of preliminary studies into the performance of seafloor mapping using HISAS SAS. As a benchmark sensor, we have chosen an EM 3000, which

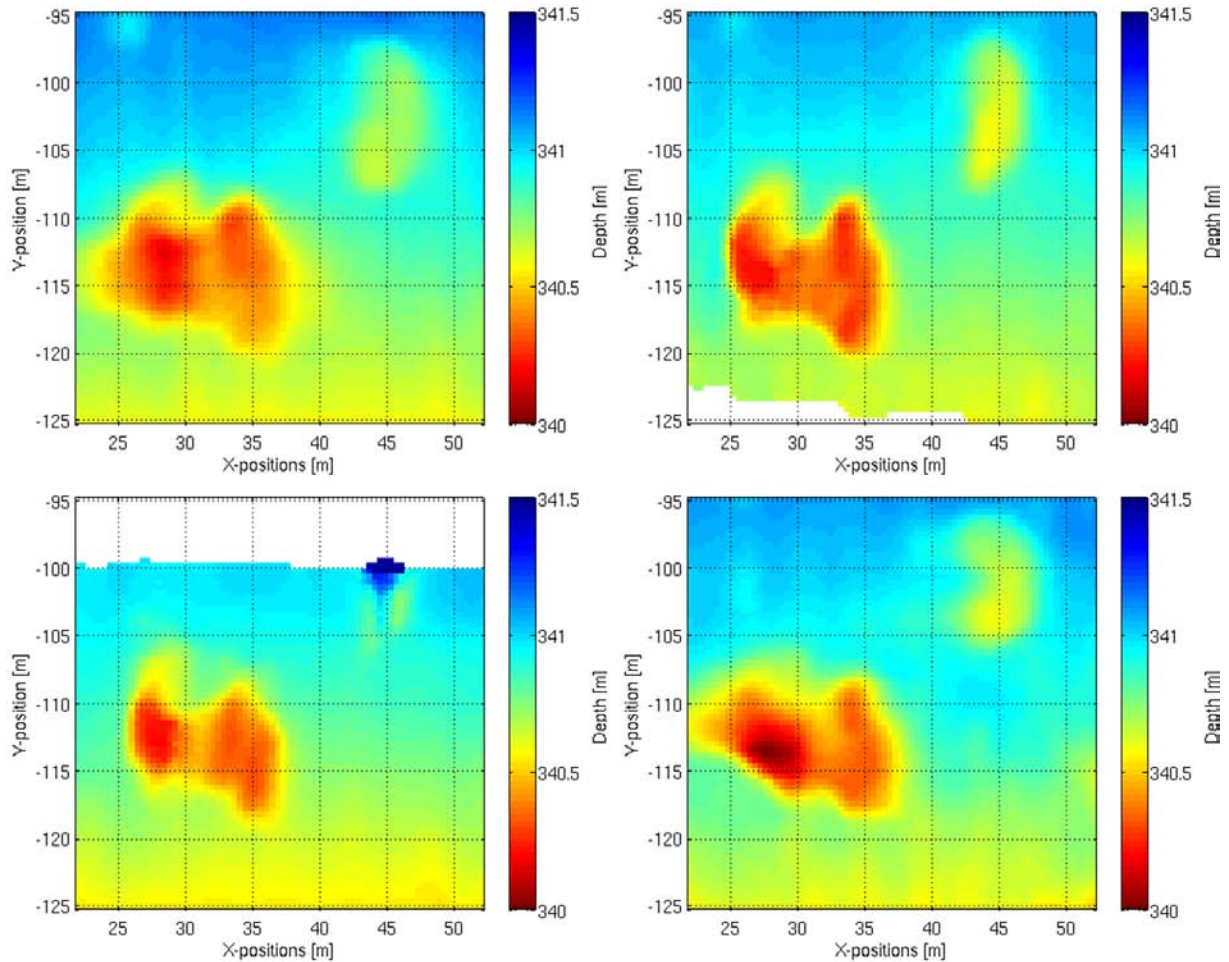


Figure 8.20: Details of the area highlighted in Figure 8.18. The panels show depth maps using HISAS sidescan bathymetry. The upper left panel shows the map from pass #1, the upper right from pass #2, the lower left from pass #3 and the lower right from pass #4. In the upper panels the area is observed from below (high y -position), in the lower panels, the area is observed from above. Pass #1 and #4 are observed from long range (around 100 m) while pass #2 and #3 are observed from shorter range. The blank areas in pass #2 and #3 are too close to nadir (less than 45 degrees). Notice a phase wrap error at close range in pass #3. The biases between each pass are subtracted in the figure. The altitude of the AUV was 25 meters in all four passings.

through experience has proven to be a reliable and accurate mapping sensor. A study of the theoretical performance shows that the depth accuracy of the two systems is similar. When it comes to horizontal resolution, the EM has slightly better, but comparable performance to sidescan bathymetry processed HISAS data. On the other hand, HISAS has a considerable larger area coverage rate. Since the swaths of the two sensors are almost non-overlapping, combining these two sensors together makes a very efficient mapping system.




Figure 8.21: Details of the positions highlighted in Figure 8.18. The left panel shows the estimated geo-referenced depths using sidescan bathymetry at the positions illustrated by the vertical white line in Figure 8.18. The right panel shows the same comparison at the positions of the horizontal white line. The average depths are shown in black, while the colored curves show the depth for each individual pass. Clearly, there is a bias on pass #1 and #4.

The SAS bathymetry processed HISAS data has approximately the same area coverage, and a much higher horizontal resolution. The computational burden in full-swath SAS bathymetry is large, but SAS bathymetry on selected regions is an excellent classification aid.

We have investigated the actual performance on two different sets of data. The first survey took place at shallow water where we compared geo-referenced depth maps of a one km² area using HISAS sidescan bathymetry and EM bathymetry. The depth maps showed very good agreement, but we found a small bias and a small cross-track error. These are most likely caused by an external lever-arm error on one of the systems.

The second survey was a deep water survey where we looked in more detail at the differences between HISAS and EM bathymetry. In this case we found no significant bias, but we noticed that the depth of the selected region were drifting as a function of time for HISAS. After eliminating a number of different error-sources, we found it likely that this drift is caused by an error in the AUV-position. Disregarding the drift, the performance of the two sensors was as expected from the specifications. Also on smaller objects, the EM and HISAS can be used as one combined system.

We also compared the EM and HISAS sidescan bathymetry depth maps to HISAS SAS bathymetry depth maps and found that the resolutions of the SAS bathymetry depth maps are superior. Due to the combination of high resolution and small depression angle, SAS bathymetry is more susceptible to shadowing and layover. Combining SAS bathymetry from two sides will efficiently combat this effect at the expense of reduced area coverage rate.

8.6 Using HISAS for pipeline surveying

In this section, we investigate how to use interferometric SAS in pipeline inspection. The results and presentation in this section is based on Sæbø et al. (2010) and Hansen et al. (2010b). Pipeline inspection is usually performed by using a short-range sensor close to the pipeline, which gives limited information about the surroundings. By using an interferometric SAS, a large area including the pipeline can be mapped and imaged in one swath. In June 2009 the Royal Norwegian Navy performed a survey outside the western coast of Norway using HISAS on an HUGIN 1000-MR. The main purpose of the mission was to survey an area with a pipeline, exposed, partly exposed and buried. The particular pipeline has an inner diameter of $12'' \approx 30$ cm, and known position prior to the mission.

A pipeline inspection survey with imaging sonar may contain the following tasks (Gauer et al., 1999)

- Detection of burial for exposed pipelines
- Detection of exposure for buried pipelines
- Detection of pipeline free span and estimation of span
- Detection of damages of pipelines
- Detection of buckling of pipelines
- Detection of debris near the pipeline

For long automated AUV surveys and short range sensors, tracking of the pipeline (Pettillot et al., 2002) might be mandatory. With the swath width of HISAS and the navigation accuracy of the HUGIN AUV, tracking actually only becomes important for multi-hour missions (given that the pipeline position is known). In this section, we focus on the sensor, interferometric SAS, and the particular benefits and challenges in inspection of the pipeline and the area around.

Figure 8.22 shows an overview of the HUGIN mission. The color-coded bathymetry is estimated using HISAS sidescan bathymetry. The water depth varies from 200 to 400 meters in the test area. The tracks are shown in black and the pipeline in brown. We focus on two different areas from this mission: Area A and Area B as indicated by the two squares. Figure 8.23 shows a sidescan sonar image from one mission line from one of the sides. Note that the pipeline is covered by gravel in several locations.

In Area B the AUV runs parallel to the pipeline at around 60 meters range to both sides. We have co-registered (Jakowatz et al., 1996, 288-302) data from the two sides and merged SAS images and SAS bathymetries. The results are presented in Section 8.6.1. In Area A the AUV has run perpendicular to the pipeline and passed directly above it. Processed data from Area A are presented in Section 8.6.2. In Section 8.6.3 we present results from shape estimation of the pipeline from both datasets.




Figure 8.22: Overview of the pipeline mission. The data are collected by the Royal Norwegian Navy, outside the western coast of Norway




Figure 8.23: Sidescan sonar image of the pipeline as a function of pings and range. The range is 0 meters (bottom) to 180 meters (top). The approximate distance traveled is 2000 m.

Figure 8.24(a) shows an example SAS image from the pipeline mission. The image shows a 50 by 100 meters scene, and the zoomed panel shows 10 by 10 meters cutout. Note the very high resolution. The pipeline is exposed, partly buried, and buried. The pipeline's surroundings are particularly important – external factors are responsible for many incidents of pipeline damage. Simple examples may be misplaced fishing trawls or anchors which have been entangled with the pipeline or it can be a change in the landmass or underwater avalanches, which causes the pipeline to slide down a slope. Figure 8.24(b) shows an example where an anchor has been released nearby a pipeline and where there may be cause for a closer inspection.

Figure 8.25 shows some of the mechanisms involved in scattering of acoustic waves from an exposed cylinder (a pipeline) on the seafloor. A few simplified assumptions can be made:

- The specular reflection is the first arrival (closest range in the sonar image)
- If the pipeline is sufficiently smooth, the specular reflection will dominate the backscattered signal
- If the pipeline material is non-smooth (seen with the sonar), diffuse scattering (non-specular) will occur
- The top of the pipeline may not be observed since it requires diffuse scattering
- The acoustic signal may travel via the pipeline and then the seafloor, then into the receiver (multipath). In this case, the direction of arrival will be incorrect (i.e. incorrect bathymetry). The backscattered echo will also have a small error in range.
- The acoustic signal may travel via the seafloor and then the pipeline, then into the receiver (multipath). In this case, the direction of arrival from the pipeline will be correct (i.e. correct bathymetric depth).




Figure 8.24: The left panel shows an example SAS image of a 50 by 100 meters area from the pipeline mission. The cutout shows a small object, possibly a barrel. The right panel shows an example SAS image of a 130 by 100 meters area. The large cutout has range from 40 meters (left) to 140 meters (right). The pipeline is located in the left part of the image. The zoom shows an anchor on the seafloor around 50 meters away from the pipeline.

- Behind the pipeline, there is a shadow zone. The length of this zone is dependent on the exposure of the pipeline and the topography in the zone.

A more thorough treatment of the problem requires a study of scattering of acoustic waves from geometrical shapes with varying smoothness (Brekhovskikh and Lysanov, 1982; Ogilvy, 1991). In this section, we concentrate on a more phenomenological description from a sensor observation point of view.

A fundamental question to be answered is whether this particular pipeline produces non-specular reflection when observed with a HISAS type sonar. To answer this, we consider a single ping. Figure 8.26 shows a sectorscan image (or a 2D phased array image of one ping). The pipeline is at approximately 163 meters range. Note that sonar broadside is not perpendicular to the pipeline. The image is constructed using standard delay-and-sum (Johnson and Dudgeon, 1993, pages 112-119) with a Kaiser tapering to suppress sidelobes (Harris, 1978). This to ensure that sidelobes from the specular reflection is not confused with the scattering from other directions. Figure 8.27 shows the relative backscattering strength from the pipeline, based on the same data as in Figure




Figure 8.25: Scattering geometry operating HISAS parallel to a pipeline

8.26. The small oscillations in the backscattering strength is due to the sidelobes in the radiation pattern. The theoretical angular resolution of the phased array is

$$\frac{\lambda}{L} = \frac{1}{80} \approx 0.7^\circ. \quad (8.11)$$

We see that the backscattering strength from the pipeline is maximum at around 3.5 degrees off broadside, at the direction for specular reflection. The directivity pattern for the transmitter and the receiver is not compensated for. This effect is, however, weak, compared with the observed results. The backscattering strength falls off in both directions as expected, but not very dramatically. This indicates that there is a substantial amount of backscatter in the non-specular horizontal direction. This can be caused by two different effects:

- The pipeline has small scale roughness causing diffuse scattering
- The signal is specularly reflected via the pipeline through a bottom bounce.

8.6.1 Tracks parallel to the pipeline

In this section we consider two mission lines (or tracks) in Area B in Figure 8.22: One track west of the pipeline with an approximate range of 59 meters to the pipeline and one track east of the pipeline at a range of approximately 61 meters. Both tracks are close to parallel to the pipeline with an altitude of approximately 25 meters. Figure 8.23 shows one of the sidescan images from a parallel track to the pipeline.

For tracks parallel with the pipeline *specular reflection*, *shadowing* and *multipath* becomes important. Specular reflection gives a strong reflection exactly perpendicular to




Figure 8.26: Sectorscan image (phased array 2D image of one ping). The pipeline is at approximately 163 meters range.




Figure 8.27: Relative backscatter from the pipeline as function of angle.

the surface of the pipe. This reflection is the first echo from the pipeline in the image and gives a very distinct indication of the position of the pipeline. Note that if the pipeline is non-parallel with the vehicle track, as is the case in Figure 8.26, then specular reflection is not normal to the SAS-track.


Figure 8.28 shows example SAS images from a section of the pipeline. In the upper left panel, the pipeline is observed from east and in the upper right panel the pipeline is observed from west. Note the changed location of the shadow behind the pipeline. The lower (large) panel shows the fusion of the images from each side (see section 8.6.1).

Co-registration details


The specular reflections from the pipeline occur in different locations when observed from the two different sides (see Figure 8.28). Thus, co-registering the two images (Jakowatz et al., 1996, 288-302) has to be based on other features. We have used five locations on the rope seen below the pipeline to manually co-register the images. In addition to an alignment, we had to apply a 0.4 degrees rotation of one of the images. In our study, we did not apply scaling, since the effect in this case is insignificant. The lower panel of Figure 8.28 shows the intensity-averaged combined image of the pipeline. By examining the rope we see that the images seem successfully co-registered, almost everywhere. There are two visible errors in the fusion, circled out in Figure 8.29(c). These errors are probably due to a projection error due to an uncompensated error in the elevation of the rope.

Characterization of the pipeline


By measuring the range to the pipe, R , the height of the sensor relative to the seafloor (at the pipeline location), H , and the length of the shadow, r , we can estimate the height



(a) SAS image from the east track



(b) SAS image from the west track



(c) SAS image merged from east and west track

Figure 8.28: All panels shows SAS images from the same 60 by 40 meters on the seafloor. The AUV has run along the x -axis and range is along the y -axis. The two upper panels show the image from the east and west tracks respectively. The lower panel shows the result of merging the east and west tracks on an intensity basis.




Figure 8.29: SAS image cutout containing the pipeline and the rope. Note the incorrect fusion of the rope on two places indicated in the right panel. This is due to an uncompensated elevation of the rope.

of the pipeline relative to the seafloor, h

$$h = \frac{Hr}{R + r}. \quad (8.12)$$

This is a common method of estimating height of small objects with sidescan sonar (Reed et al., 2004).

We have estimated the height of the pipeline relative to the seafloor in a number of representative locations and found that it is typically elevated 29 cm. The actual pipeline has a diameter of 30 cm. The method is, however, not entirely conclusive. The bathymetry in the shadow causes uncertainty, and the burial of the pipeline is unknown.

In the combined SAS image in Figure 8.28(c), we see that the pipeline has a certain width. By assuming that the first arrival is from the specular reflection in the vertical plane in both sides, the diameter d can be approximated from the width w of the pipeline

$$d \approx \frac{2w}{\cos \Phi_1 + \cos \Phi_2}, \quad (8.13)$$

where θ_1 and θ_2 are the incident angles to the pipeline relative to the sonar (see Figure 8.30). By averaging the estimate over 30 meters we find that $d \approx 46$ cm. The actual pipeline diameter was 30 cm. The reason for this discrepancy may be related to the error in the combined interferometry (see the next subsection).

Combined Interferometry

We have applied the same co-registering technique on the interferometric estimates. Figure 8.31 shows the estimated bathymetry of the same region of the pipeline combined from two sides. The two sides are combined by selecting the region of maximum coherence (Hanssen, 2001, pages 96-98) in the two single-look bathymetries, i.e. there is




Figure 8.30: Geometry of a two-sided pipeline survey. The orange circle at the upper left section is the AUV traveling out of the paper and the one at the upper right illustrates the AUV traveling into the paper. The measured width between the specular reflections, w , is equal to 46 cm.




Figure 8.31: Interferometric bathymetry from the same region as imaged in Figure 8.28. The bathymetry is combined from two tracks by selecting the side with the highest coherence for each pixel. Note the small elevated feature on the left side of the pipeline. This is part of the rope.

no averaging of the estimates. We removed a 27 cm vertical bias between the two bathymetries to get a smooth surface. This is most likely caused by inaccurate navigation or tidal changes causing an apparent depth change. Note that the combined SAS interferometry of the pipeline does not have the form of a cylinder – it appears flat at the center. This indicates that there is little backscatter from the top of the pipeline, or in the non-specular direction.




Figure 8.32: A 100 meter section of the pipeline seen from a track perpendicular to the pipeline. The upper panel shows the interferometric bathymetry and the lower panel shows the SAS image rendered onto the bathymetry.

8.6.2 Track perpendicular to the pipeline

In this section, we study data from Area A in Figure 8.22, where the AUV track is perpendicular to the pipeline. The sonar look-direction is then parallel to the pipeline. Specular reflection, shadowing and multipath is not dominant (as it was in the previous section, where the vehicle track was parallel to the pipeline). In this case, diffuse scattering (or backscattering in the non-specular direction) becomes important.

Figure 8.32 shows the estimated bathymetry from the starboard sonar (upper panel) along with the SAS image rendered onto the bathymetry (lower panel). The pipeline can be clearly seen in both panels. Note that the pipeline shape is more cylindrical compared to what we got from the parallel tracks. In the SAS image we see the pipeline joints as dark regions every 12.5 meters. These rather substantial differences in interferometric response are discussed further in the following sections.




Figure 8.33: A 8 by 100 meter zoom of a 100 meter section of the pipeline bathymetry in Figure 8.33. Note that the aspect ratio is not equal in the horizontal plane. The relative height of the pipeline is approximately 30 cm.

8.6.3 Estimating the shape of the pipeline

In the two previous sections we showed that the data collected from parallel and perpendicular tracks to the pipeline have very different characteristics. In this section we investigate these differences further.

Figure 8.33 shows a 8 by 100 meter cut-out from the data in Figure 8.32. By selecting the local maximum we can detect and track the peak of the pipeline, which gives the 3D position of the pipeline. This can be used to measure the local curvature of the pipeline.

The peak location can also be used to shift the bathymetry data in such a way that the pipeline is aligned along one axis. After linear detrending the local slope of the seafloor we can average the bathymetry over the 100 meters, to get an estimate over the pipeline profile. The left panel of Figure 8.34 shows the estimated pipeline profile after shifting, detrending and averaging. The observed shape can intuitively seem to be a slight mismatch to a cylindrical pipeline, but this is actually as theory would predict:

- The interferometric estimate is not true 3D, but rather 2.5D - there is only one depth at each horizontal position. This is due to the sensor geometry – the vertical receiver antenna is a sparse array containing only two elements, thereby only one

measurement can be done.

- We apply a window of 9 by 9 image-samples to produce one interferometric estimate (see Section 8.1). This is done to reduce variance in the depth estimate. In effect, the bathymetry is lowpass filtered with an 18 cm window.

The measured height of the pipeline relative to the surrounding in the left panel of Figure 8.34 is around 30 cm, which fits with the diameter of the pipeline. The burial of the pipeline is unknown. We also tried to estimate the diameter based on the curvature of the profile, but due to the lowpass filtering, the profile is not cylindrical and the estimate varied from 33 to 40 cm depending on where we measured.

We have also performed similar analysis on the SAS image and averaged the intensity data to get a profile of the pipeline. The estimated width (and thereby diameter) varied from 30 to 57 cm depending on whether we measured -3 dB or -10 dB from the peak. This result is somewhat expected as the main lobe width of the imaging system is specified as 2.5 cm at -3 dB but wider at -10 dB.


In the right panel of Figure 8.34 we see an estimate of the pipeline profile from the merged bathymetry data using the parallel tracks. We have applied the same shifting and averaging method as described above to produce the profile. The profile show several important features: Firstly, it is clear that the profile is dominated by the two specular reflections. Since the AUV height is different for the two tracks (see Figure 8.30) the specular reflections are at different depths. They are also smeared in range due to the interferometric window. Secondly, it seems that there is little echo from the top of the pipeline. There are some energy at the center of the pipeline profile, but the height is not above the specular reflections. Thirdly, there are energy below the seafloor surface, which indicates that we receive a multipath bounced onto the seafloor before it is echoed back into the sonar.

8.6.4 Summary

In this section we have shown preliminary results from a pipeline survey using HISAS as primary sensor. We have focused on two different sections of the pipeline, one where the AUV tracks were parallel to the pipeline on each side and one where the track was perpendicular to the pipeline.

In the parallel tracks the pipeline is easily identifiable. The backscattered field is affected by specular reflection, shadowing and multipath. These features are also present in the interferometric data, but the estimated bathymetry is still usable to separate the pipeline from the rest of the seafloor. The shadowing can be eliminated by merging the data from two sides. The bathymetric pipeline profile does not fit a cylindrical shape, as expected, since the interferometric bathymetry estimate is dominated by the specular reflection.

In the track perpendicular to the pipeline, the pipeline is easily distinguished from the rest of the seafloor, both in the SAS images and in the interferometric data. The bathymetric profile of the pipeline looked more like a lowpass filtered cylindrical object.



(a) Estimated pipeline profile from the perpendicular track (b) Estimated pipeline profile from the parallel tracks

Figure 8.34: Pipeline profile estimated by tracking the pipeline peak in the bathymetry, shifting and aligning the pipeline, detrending the local slope of the seafloor and averaging the bathymetry. The left panel show the results from the perpendicular track. The data is from a 100 meter section of the bathymetry showed in Figures 8.32 and 8.33. The right panel shows the result from the merged bathymetry of the parallel tracks. The data is from a 60 meter section of the bathymetry showed in Figure 8.31. The black line is the mean profile, while the colors indicates the distribution of points at each location. Red means a large number of samples at that location and blue means a low number of samples.

From the results we have presented it is clear that HISAS can be used to get useful information from a pipeline and its surroundings. The high resolution imagery and bathymetry of the seafloor is ideal for locating small objects in the proximity of the pipeline. At the pipeline the data are less intuitive – due to several acoustic effects which complicate the interpretation.

We have found that it may be beneficial to observe the pipeline not only from parallel tracks, which is most efficient, but also from perpendicular tracks. This can produce valuable information not easily retrievable otherwise. Whether this is worth the cost of increased survey time is left to further investigation.

Chapter 9

Summary and conclusions

In this thesis we have developed methods for robust and accurate interferometric SAS. We have emphasized that *a priori* depth information is an important prerequisite for SAS imaging and developed a solution through sidescan depth estimation. We have described the integration of depth information into the navigation and the image formation. Finally, we have developed a method for interferometric SAS, and performed detailed studies into possible advances and applications.

One of two main topics in this thesis is a thorough description of the geometry in interferometric sidelooking sonar. An important conclusion is that the interferometric time series are not only time delayed relative to each other, but also dilated. We have developed a new mathematical description of the geometry in sidelooking sonar and shown how an approximation-free solution for the depth can be obtained.

The other main topic is time delay estimation. We have argued that a standard time delay estimator like the cross-correlation function is sub-optimum on interferometric time series due to the relative dilation. We presented two solutions: The cross-uncertainty function and a nonlinear co-registration by interpolating the data onto an *a priori* ground-range. We showed that the cross-uncertainty function has better performance than the cross-correlation function on time delayed time series and approximately equal performance on non-dilated time series. We also introduced an improved cross-correlation based estimator on data with asymmetrical spectra.

We have developed two depth estimation routines: Sidescan bathymetry and SAS interferometry. We introduced a well performing unwrapper for cross-correlation based sidescan bathymetry estimates, and briefly discussed the concept of iterative SAS interferometry. A short summary on the difference between SAS interferometry and SAR interferometry tried to shed light on where typical SAS-systems are similar to SAR-systems.

There is a lack of a standard on how to measure resolution in bathymetric maps. We suggested *flank length* as a measure of the horizontal resolution in depth maps. We estimated the flank length on simulated datasets to verify the theoretical values. The conclusion was the implemented sidescan bathymetry and SAS interferometry performed very close to the theoretical limits when it comes to horizontal resolution. When

it comes to the vertical precision of the depth estimation, it followed the theoretical performance within a factor of two. An important observation is that it is the SNR estimated from the coherence which limits the vertical precision, not the actual SNR. This is because the estimated SNR also includes model errors like baseline decorrelation.

We presented results from a number of studies using experimental HISAS data. The topics are to a large extent unrelated. The conclusions are as follow:

- It seems unrealistic to achieve depth estimates with full SAS image resolution. Due to the speckly nature of SAS images, some kind of filtering has to be applied. The phase-difference filter is a maximum likelihood estimate of the interferometric phase-difference. The size of the filter is a trade-off between vertical precision and horizontal resolution. At some level the precision will not be improved by increasing the size of the filter, since it assumes homogeneity. A better solution is then to post-smooth the depth estimates.
- In slant-range sidescan bathymetry (or other cases where the time delay estimate is limited by dilation of the time series) the cross-uncertainty function outperforms a standard cross-correlation based technique with as much as an equivalent 20 dB SNR increase. A cross-uncertainty based technique is therefore a good alternative to the non-linear co-registration in ground-range (which we have used in the rest of this thesis). The dilation-factor may also provide useful information which cannot be obtained from the time delay estimate itself.
- A split-bandwidth technique which divides the bandwidth into sub-bands, can be used to resolve the 2π ambiguities in SAS interferometry. We implemented a simple technique, which outperformed a standard 2D unwrapping technique (the Goldstein Branch cut algorithm) on large scale. However, due to the simplicity of the method it failed on a small number of random pixels.
- The weighted bilateral filter outperforms a large phase-difference filter, a weighted smoothing filter and a weighted median filter when it comes to smoothing depth maps, while retaining sharp edges.
- Sidescan bathymetry using HISAS has very similar performance to depth estimation using an EM 3000 multibeam echo sounder. For simultaneously operation using an HUGIN 1000-MR AUV in typical geometry, the EM 3000 has slightly better resolution while HISAS has larger area coverage rate. HISAS SAS bathymetry has large area coverage rate *and* superior resolution. We also found that there is a small discrepancy remaining between the two sensors.
- A HISAS interferometric SAS can be used to get useful information from a pipeline and its surroundings. We have found that due to acoustic effects which complicates the interpretation, it can be beneficial to observe the pipeline not only from parallel tracks, but also from perpendicular tracks.

9.1 Suggested future work

Some of the topics we investigated in this thesis have not been addressed to their full extent, due to a limited time span. A few of them which deserves further attention, are listed below:

- An expanded cross-uncertainty estimator which estimates the dilation both from the geometry and from the data itself. A goal should be to find a continuous model-based estimate of the delay and the dilation instead of the patch-based method presented in this thesis.
- A better and more robust split-bandwidth implementation based on the probability density functions for each sub-band having the correct wrap interval. Either by using a maximum likelihood method by exploiting deterministic contextual information consisting in the approximation of the depth surface through local planes (Budillon et al., 2005) or by fusing the interferograms from each sub-band with a coherence weighted Kalman filter (Lanari et al., 1996).
- A study of the discrepancy between HISAS and the EM 3000 when operated simultaneously on the HUGIN 1000-MR AUV.
- An even more thorough study of the acoustic effects when using sidelooking sonar to survey elevated man-made objects like a pipeline.

Bibliography

EM 3000, Multibeam echo sounder, rev. C edition, October 2003.

Triton Imaging, Inc. eXtended Triton Format (XTF), rev. 26 edition, December 2008.

Kongsberg EM Series Multibeam echo sounder: EM datagram formats, 850th edition, January 2010.

MATLAB 7 MAT-File Format, version 7.11 edition, September 2010.

R. A. Altes. *Some invariance properties of the wide-band ambiguity function*. *J. Acoust. Soc. Am.*, **53**(4): 1154–1160, 1973.

J. T. Ambrose and A. Geneva. *Wide swath bathymetric sonar error modeling for a hydrographic survey ship*. In *OCEANS '95 MTS/IEEE*, pp. 339–346. San Diego, CA, USA, October 1995.

R. Bamler and M. Eineder. *Accuracy of differential shift estimation by correlation and split-bandwidth interferometry for wideband and delta-k sar systems*. *IEEE Trans. Geosci. Remote Sensing.*, **2**(2): 151–155, 2005.

R. Bamler and P. Hartl. *Synthetic aperture radar interferometry*. *Inverse problems*, **14**(4): 1–54, 1998.

K. L. Bell, Y. Steinberg, Y. Ephraim and H. L. Van Trees. *Extended Ziv-Zakai lower bound for vector parameter estimation*. *IEEE Trans. Inform. Theory*, **43**(2): 624–637, 1997.

A. Bellettini and M. A. Pinto. *Theoretical accuracy of synthetic aperture sonar micronavigation using a displaced phase-center antenna*. *IEEE J. Oceanic Eng.*, **27**(4): 780–789, 2002.

S. Benedetto and E. Biglieri. *Principles of digital transmission: With wireless applications*. Kluwer Academic/Plenum Publishers, 233 Spring Street, New York, NY 10013, USA, 1999.

J. S. Bird and P. Kraeutner. *Cramer-Rao bound investigation of swath bathymetry accuracy*. In *Proceedings of Oceans 2001 MTS/IEEE*, pp. 1640–1647. Honolulu, HI, USA, November 2001.

- J. S. Bird and G. K. Mullins. *Analysis of swath bathymetry sonar accuracy*. *IEEE J. Oceanic Eng.*, **30** (2): 372–390, 2005.
- W. W. Bonifant Jr, M. A. Richards and J. H. McClellan. *Interferometric height estimation of the seafloor via synthetic aperture sonar in the presence of motion errors*. *IEE Proc. Radar, Sonar Navig.*, **147** (6): 322–330, 2000.
- L. Brekhovskikh and Y. Lysanov. *Fundamentals of ocean acoustics*, volume 8 of *Springer Series in electrophysics*. Springer-Verlag, Berlin, Germany, 1982.
- A. R. Brenner. *Proof of concept for airborne SAR imaging with 5 cm resolution in X-band*. In *Proceedings of EUSAR 2010*. Aachen, Germany, June 2010.
- A. R. Brenner and L. Roessing. *Radar imaging of urban areas by means of very high-resolution SAR and interferometric SAR*. *IEEE Trans. Geosci. Remote Sensing.*, **46**(10): 2971–2982, October 2008.
- A. R. Brenner, L. Roessing and P. Berens. *Potential of very high resolution SAR interferometry for urban building analysis*. In *Proceedings of EUSAR 2010*. Aachen, Germany, June 2010.
- D. R. K. Brownrigg. *The weighted median filter*. In *In Communications of the ACM*, volume 27 (8), pp. 807–818. 1984.
- M. P. Bruce. *A processing requirement and resolution capability comparison of side-scan and synthetic-aperture sonars*. *IEEE J. Oceanic Eng.*, **17**(1): 106–117, 1992.
- A. Budillon, G. Ferraiuolo, V. Pascazio and G. Schirinzi. *Multichannel SAR interferometry via classical and Bayesian estimation techniques*. *Journal on Applied Signal Processing*, **20**: 3180–3193, 2005.
- W. S. Burdick. *Underwater acoustic system analysis*. Prentice-Hall, 113 Sylvan Avenue, Englewood Cliffs, NJ 07632, USA, 1984.
- H. J. Callow. *Signal processing for synthetic aperture sonar image enhancement*. Ph.D. thesis, University of Canterbury, Christchurch, New Zealand, 2003.
- H. J. Callow, M. P. Hayes and P. T. Gough. *Motion-compensation improvement for widebeam, multiple-receiver SAS systems*. *IEEE J. Oceanic Eng.*, **34** (3): 262–268, 2009.
- W. G. Carrara, R. S. Goodman and R. M. Majewski. *Spotlight synthetic aperture radar: Signal processing algorithms*. Artech House, 685 Canton Street, Norwood, MA 02062, USA, 1995.
- Y. T. Chan and K. C. Ho. *Joint time-scaling and TDOA estimation: Analysis and fast approximation*. *IEEE Trans. Signal Processing*, **53** (8): 2625–2634, 2005.

- J. C. Curlander and R. N. McDonough. *Synthetic aperture radar: Systems and signal processing*. Wiley Series in remote sensing. John Wiley & Sons, 605 Third Avenue, New York, NY 10158, USA, 1991.
- L. J. Cutrona. *Comparison of sonar system performance achievable using synthetic-aperture techniques with the performance achievable by more conventional means*. *J. Acoust. Soc. Am.*, **58** (2): 336–348, August 1975.
- P. N. Denbigh. *Swath bathymetry: Principles of operation and an analysis of errors*. *IEEE J. Oceanic Eng.*, **14** (4): 289–298, 1989.
- P. N. Denbigh. *Signal processing strategies for a bathymetric sidescan sonar*. *IEEE J. Oceanic Eng.*, **19** (3): 382–390, 1994.
- J. V. DiFranco and W. L. Rubin. *Radar detection*. Artech House, 685 Canton Street, Norwood, MA 02062, USA, 1980.
- L. B. Dozier, J. S. Hanna and C. R. Pearson. *Treatments of incoherent scattering for the parabolic equation and astral propagation models*. In *Ocean Variability & Acoustic Propagation*, edited by J. Potter and A. Warn-Varnas, pp. 265–281. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1991.
- J. H. G. Ender and A. R. Brenner. *PAMIR a wideband phased array SAR/MTI system*. *IEE Proc. Radar, Sonar Navig.*, **150**(3): 165–172, 2003.
- C. Falsi, D. Dardari, L. Mucchi and M. Z. Win. *Time of arrival estimation for uwb localizers in realistic environments*. *Journal on Applied Signal Processing*, **2006**: 152–164, 2006.
- J. R. Fienup. *Phase error correction by shear averaging*. In *Signal Recovery and synthesis III*, pp. 134–137. Optical Society of America, 1989.
- S. M. Flatte, R. Dashen, W. H. Munk, K. M. Watson and F. Zachariasen. *Sound transmission through a fluctuating ocean*. Cambridge monographs on mechanics and applied mathematics. Cambridge University Press, Trumpington Street, Cambridge CB2 1RP, UK, 1979.
- G. Fornaro and G. Franceschetti. *Image registration in interferometric SAR processing*. *IEE Proc. Radar, Sonar Navig.*, **142**(6): 313–320, December 1995.
- T. G. Fossum, T. O. Sæbø, B. Langli, H. J. Callow and R. E. Hansen. *HISAS 1030 - high resolution interferometric synthetic aperture sonar*. In *Proceedings of the Canadian Hydrographic Conference and National Surveyors Conference*. Victoria, BC, Canada, May 2008.
- G. Franceschetti and R. Lanari. *Synthetic aperture radar processing*. CRC Press, 2000 Corporate Blvd. Boca Raton, FL 33431, USA, 1999.

- K. Gade. *NavLab, a generic simulation and post-processing tool for navigation*. *European Journal of Navigation*, **2**(4): 51–59, November 2004.
- F. Gatelli, A. M. Guarnieri, F. Parizzi, P. Pasquali, C. Prati and F. Rocca. *The wavenumber shift in SAR interferometry*. *IEEE Trans. Geosci. Remote Sensing.*, **32**(4): 855–865, July 1994.
- R. Gauer, A. McFadzean and C. Reid. *An automated sidescan sonar pipeline inspection system*. In *OCEANS '99 MTS/IEEE*, pp. 811–816. Seattle, WA, USA, September 1999.
- N. Gebert, G. Krieger and A. Moreira. *Multichannel azimuth processing in ScanSAR and TOPS mode operation*. *IEEE Trans. Geosci. Remote Sensing.*, **48**(7): 2994–3308, July 2010.
- D. C. Ghiglia and M. D. Pritt. *Two-dimensional phase unwrapping: Theory, algorithms, and software*. John Wiley & Sons, 605 Third Avenue, New York, NY 10158, USA, 1998.
- I. A. Glover and P. M. Grant. *Digital communications*. Prentice-Hall Europe, Hemel Hempstead, Hertfordshire HP2 7EZ, UK, 1998.
- R. D. Gonzales and R. E. Woods. *Digital image processing*. Addison-Wesley Publishing Company, Reading, MA, USA, 1992.
- P. T. Gough and D. W. Hawkins. *Imaging algorithms for a strip-map synthetic aperture sonar: Minimizing the effects of aperture errors and aperture undersampling*. *IEEE J. Oceanic Eng.*, **22** (1): 27–39, 1997.
- H. D. Griffiths, T. A. Rafik, Z. Meng, C. F. N. Cowan, H. Shafeeu and D. K. Anthony. *Interferometric synthetic aperture sonar for high-resolution 3-D mapping of the seabed*. *IEE Proc. Radar, Sonar Navig.*, **144** (2): 96–103, 1997.
- J. Groen, R. E. Hansen, H. J. Callow, J. C. Sabel and T. O. Sæbø. *Shadow enhancement in synthetic aperture sonar using fixed focusing*. *IEEE J. Oceanic Eng.*, **34** (3): 269–284, 2009.
- P. E. Hagen and R. E. Hansen. *Area coverage rate of synthetic aperture sonars*. In *Proceedings of Oceans 2007 Europe*. Aberdeen, UK, June 2007.
- P. E. Hagen, N. J. Størkersen and K. Vestgård. *Hugin – use of uuv technology in marine applications*. In *OCEANS '99 MTS/IEEE*, pp. 967–972. Seattle, WA, USA, September 1999.
- P. E. Hagen, N. Størkersen and K. Vestgård. *The HUGIN 1000 military AUV system*. In *UDT Europa 2004 Conference Proceedings*. Nice, France, June 2004.
- P. E. Hagen, N. Størkersen and K. Vestgård. *HUGIN AUV's in marine research – today and tomorrow*. In *Proceedings from International Workshop on Underwater Robotics (IWUR)*. Genoa, Italy, November 2005.

- P. E. Hagen, R. E. Hansen and B. Langli. *Interferometric synthetic aperture sonar for the HUGIN 1000-MR AUV*. In *UDT Pacific 2006 Conference Proceedings*. San Diego, CA, USA, November 2006.
- R. E. Hansen, T. O. Sæbø, H. J. Callow, P. E. Hagen and B. Langli. *Synthetic aperture sonar processing for the HUGIN AUV*. In *Proceedings of Oceans 2005 Europe*. Brest, France, June 2005.
- R. E. Hansen, T. O. Sæbø, H. J. Callow and B. Langli. *The SENSOTEK interferometric synthetic aperture sonar: Results from HUGIN AUV trials*. In *Proceedings of the 8th European Conference on Underwater Acoustics*. Carvoeiro, Portugal, June 2006.
- R. E. Hansen, H. J. Callow and T. O. Sæbø. *The effect of sound velocity variations on synthetic aperture sonar*. In *Proceedings of Underwater Acoustic Measurements 2007*. Crete, Greece, June 2007.
- R. E. Hansen, H. J. Callow, T. O. Sæbø, S. A. Synnes, P. E. Hagen, T. G. Fossum and B. Langli. *Synthetic aperture sonar in challenging environments: Results from the HISAS 1030*. In *Proceedings of Underwater Acoustic Measurements 2009*. Nafplion, Greece, June 2009. CDROM (ISBN 978-960-98883-4-9).
- R. E. Hansen, H. J. Callow, T. O. Sæbø and S. A. V. Synnes. *Challenges in seafloor imaging and mapping with synthetic aperture sonar*. In *Proceedings of the Eighth European conference on synthetic aperture radar*. Aachen, Germany, June 2010a. CDROM (ISBN 978-3-8007-3272-2).
- R. E. Hansen, T. O. Sæbø, H. J. Callow and P. E. Hagen. *Interferometric synthetic aperture sonar in pipeline inspection*. In *Proceedings of OCEANS 2010 MTS/IEEE*. Sydney, Australia, May 2010b. CDROM (ISBN 978-1-4244-5222-4).
- A. Hanssen, J. Kongsli, R. E. Hansen and S. Chapman. *Statistics of synthetic aperture sonar images*. In *Proceedings of Oceans 2003 MTS/IEEE*, pp. 2635–2640. San Diego, CA, USA, September 2003.
- R. F. Hanssen. *Radar interferometry: Data interpretation and error analysis*. Kluwer Academic Publishers, P. O. Box 17, 3300 AA Dordrecht, The Netherlands, 2001.
- F. J. Harris. *On the use of windows for harmonic analysis with the discrete fourier transform*. *Proceedings of the IEEE*, **66** (1): 51–83, 1978.
- D. W. Hawkins. *Synthetic aperture imaging algorithms: With application to wideband sonar*. Ph.D. thesis, University of Canterbury, Christchurch, New Zealand, 1996.
- M. P. Hayes and P. T. Gough. *Synthetic aperture sonar: A review of current status*. *IEEE J. Oceanic Eng.*, **34** (3): 207–224, 2009.

- Ø. Hegrenæs, T. O. Sæbø, P. E. Hagen and B. Jalving. *Horizontal mapping accuracy in hydrographic AUV surveys*. In *Proceedings of the IEEE Autonomous Underwater Vehicles 2010*. Monterey, CA, USA, September 2010.
- S. Holm. *Optimum FFT-based frequency acquisition with application to COSPAS-SARSAT*. *IEEE Trans. Aerospace and Electronic Systems.*, **29** (2): 464–475, 1993.
- S.-H. Hong, H.-S. Jung and J.-S. Won. *Extraction of ground control points (GCPs) from synthetic aperture radar images and SRTM DEM*. *International Journal of Remote Sensing*, **27**(18): 3813–3829, 2006.
- D. R. Jackson and M. D. Richardson. *High-frequency seafloor acoustics*. Springer Monograph series in underwater acoustics. Springer Science+Business Media, 233 Spring Street, New York, NY 10013, USA, 2007.
- J. 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, P. O. Box 17, 3300 AA Dordrecht, The Netherlands, 1996.
- B. Jalving, K. Gade, O. K. Hagen and K. Vestgård. *A toolbox of aiding techniques for the HUGIN AUV integrated inertial navigation system*. In *Proceedings of Oceans 2003 MTS/IEEE*, pp. 1146–1153. San Diego, CA, USA, September 2003.
- D. H. Johnson and D. E. Dudgeon. *Array signal processing: Concepts and techniques*. Signal processing series. Prentice Hall, 113 Sylvan Avenue, Englewood Cliffs, NJ 07632, USA, 1993.
- C. D. de Jong, G. Lachapelle, A. Skone and I. A. Elema. *Hydrography*. Series on Mathematical geodesy and positioning. VSSD, Leeghwaterstraat 42, 2628 CA Delft, The Netherlands, 2002.
- L. R. Joughin and D. P. Winebrenner. *Effective number of looks for a multilook interferometric phase distribution*. In *Proceedings of Geoscience and Remote Sensing Symposium (IGARSS) 1994*, pp. 2276–2278. Pasadena, CA, USA, August 1994.
- D. Just and R. Bamler. *Phase statistics of interferograms with applications to synthetic aperture radar*. *Appl. Opt.*, **33**(20): 4361–4368, 1994.
- S. M. Kay. *Fundamentals of statistical signal processing. Estimation theory*. Prentice Hall, 113 Sylvan Avenue, Englewood Cliffs, NJ 07632, USA, 1993.
- D. Kincaid and W. Cheney. *Numerical analysis*. Brooks/Cole Publishing Company, 511 Forest Lodge Road, Pacific Grove, CA 93950, USA, second edition, 1991.
- B. Kjellesvig. *Improved height estimation algorithms for synthetic aperture sonar interferometry*. Master's thesis, University of Oslo, Oslo, Norway, 2006.

- C. K. Knapp and G. C. Carter. *The generalized correlation method for estimation of time delay*. *IEEE Trans. Acoust., Speech, Signal Processing*, **ASSP-24**(4): 320–327, 1976.
- G. Krieger, A. Moreira, H. Fiedler, I. Hajnsek, M. Werner, M. Younis and M. Zink. *TanDEM-X: A satellite formation for high-resolution SAR interferometry*. *IEEE Trans. Geosci. Remote Sensing.*, **45**(11): 3317–3341, November 2007.
- G. Krieger, I. Hajnsek, K. P. Papathanassiou, M. Younis and A. Moreira. *Interferometric synthetic aperture radar (SAR) missions employing formation flying*. *Proceedings of the IEEE*, **98**(5): 816–843, May 2010.
- R. Lanari, G. Fornaro, D. Riccio, M. Migliaccio, K. P. Papathanassiou, J. R. Moreira, M. Schwabisch, L. Dutra, G. Puglisi, G. Franceschetti and M. Coltelli. *Generation of digital elevation models by using SIR-C/X-SAR multifrequency two-pass interferometry: The Etna case study*. *IEEE Trans. Geosci. Remote Sensing.*, **34**(5): 1097–1114, 1996.
- M. Lange and P. Gill. *An airborne interferometric SAR system for high performance 3D mapping*. In *Proc. of SPIE Radar Sensor Technology XIII*, volume 7308. Orlando, FL, USA, April 2009.
- N. Levanon. *Radar principles*. John Wiley & Sons, 605 Third Avenue, New York, NY 10158, USA, 1988.
- G. Llorc-Pujol, C. Sintès and X. Lurton. *A new approach for fast and high-resolution interferometric bathymetry*. In *Proceedings of Oceans 2006 Asia Pacific*. Singapore, May 2006.
- X. Lurton. *Swath bathymetry using phase difference: Theoretical analysis of acoustical measurement precision*. *IEEE J. Oceanic Eng.*, **25** (3): 351–363, 2000.
- X. Lurton. *An introduction to underwater acoustics: Principles and applications*. Springer Praxis Publishing, Chichester, UK, 2002.
- A. P. Lyons and D. A. Abraham. *Statistical characterization of high-frequency shallow-water seafloor backscatter*. *J. Acoust. Soc. Am.*, **106** (3): 1307–1315, 1999.
- S. N. Madsen, H. A. Zebker and J. Martin. *Topographic mapping using radar interferometry: Processing techniques*. *IEEE Trans. Geosci. Remote Sensing.*, **31**(1): 246–256, 1993.
- J. H. McClellan, R. W. Schafer and M. A. Yoder. *Signal processing first*. Prentice Hall, Upper Saddle River, NJ 07458, USA, 2003.
- Ø. Midtgaard, T. O. Sæbø and H. J. Callow. *Automated recognition of short-tethered objects with interferometric synthetic aperture sonar*. In *Proceedings of Oceans 2007 MTS/IEEE*. Vancouver, Canada, October 2007.
- Ø. Midtgaard, T. O. Sæbø and R. E. Hansen. *Estimation of target detection/classification performance using interferometric sonar coherence*. In *Proceedings of Underwater Acoustic Measurements 2009*. Nafplion, Greece, June 2009. CDROM (ISBN 978-960-98883-4-9).

- S. K. Mitra. *Digital signal processing: A computer based approach*. McGraw-Hill, 1221 Avenue of the Americas, New York, NY 10020, USA, 2002.
- R. Moddemeijer. *On the determination of the position of extrema of sampled correlators*. *IEEE Trans. Signal Processing*, **39** (1): 216–219, 1991.
- C. de Moustier, P. F. Lonsdale and A. N. Shor. *Simultaneous operation of the Sea Beam multibeam echo-sounder and the SeaMARC II bathymetric sidescan sonar system*. *IEEE J. Oceanic Eng.*, **15**(2): 84–94, 1990.
- J. A. Ogilvy. *Theory of wave scattering from random rough surfaces*. Institute of Physics Publishing, Redcliffe Way, Bristol BS1 6NX, UK, 1991.
- C. Oliver and S. Quegan. *Understanding synthetic aperture radar images*. SciTech Publishing, Raleigh, NC 27613, USA, 2004.
- P. Z. Peebles. *Probability, random variables and random signal principles*. McGraw-Hill series in electrical and computer engineering. McGraw-Hill, 1221 Avenue of the Americas, New York, NY 10020, USA, fourth edition, 2001.
- Y. Petillot, S. Reed and J. Bell. *Real time AUV pipeline detection and tracking using side scan sonar and multi-beam echosounder*. In *OCEANS '02 MTS/IEEE*, pp. 217–222. Biloxi, MS, USA, October 2002.
- M. Pinto. 2003. Private communication.
- M. A. Pinto. *High resolution seafloor imaging with synthetic aperture sonar*. *IEEE Oceanic Eng. Newsletter*, pp. 15–20, Summer 2002.
- A. H. Quazi. *An overview on the time delay estimate in active and passive systems for target localization*. *IEEE Trans. Acoust., Speech, Signal Processing*, **ASSP-29**(3): 527–533, 1981.
- B. Rabus, M. Eineder, A. Roth and R. Bamler. *The shuttle radar topography mission – a new class of digital elevation models acquired by spaceborne radar*. *Photogramm. Rem. Sens.*, **57**(4): 241–262, 2003.
- S. Reed, Y. Petillot and J. Bell. *Automated approach to classification of mine-like objects in sidescan sonar using highlight and shadow information*. *IEE Proc. Radar, Sonar Navig.*, **151**: 48–56, 2004.
- M. A. Richards. *Fundamentals of radar signal processing*. The McGraw Hill Companies, Two Penn Plaza, New York, NY 10121, USA, 2005.
- D. W. Ricker. *Echo signal processing*. Kluwer Academic Publishers, P. O. Box 17, 3300 AA Dordrecht, The Netherlands, 2003.
- E. Rodriguez and J. M. Martin. *Theory and design of interferometric synthetic aperture radars*. *Radar and Signal Processing, IEE Proceedings F*, **139** (2): 147–159, 1992.

- K. D. Rolt and H. Schmidt. *Effects of refraction on synthetic aperture sonar imaging*. *J. Acoust. Soc. Am.*, **95**(6): 3424–3429, 1994.
- P. A. Rosen, S. Hensley, I. R. Joughin, F. K. Li, S. N. Madsen, E. Rodriguez and R. M. Goldstein. *Synthetic aperture radar interferometry*. *Proceedings of the IEEE*, **88**(3): 333–382, 2000.
- K. Rottmann. *Matematisk formelsamling*. Bracan Forlag, Nørhaven, Denmark, third edition, 1995. In Norwegian.
- T. O. Sæbø and R. E. Hansen. *A comparison between interferometric SAS and interferometric SAR*. In *Proceedings of the Institute of Acoustics*. Lerici, Italy, September 2010.
- T. O. Sæbø and B. Langli. *Comparison of EM 3000 multibeam echo sounder and HISAS 1030 interferometric synthetic aperture sonar for seafloor mapping*. In *Proceedings of ECUA 2004*, pp. 451–461. Istanbul, Turkey, July 2010.
- T. O. Sæbø, R. E. Hansen, H. J. Callow and B. Kjellesvig. *Using the cross-ambiguity function for improving sidelooking sonar height estimation*. In *Proceedings of Underwater Acoustic Measurements 2007*, pp. 309–316. Crete, Greece, June 2007a.
- T. O. Sæbø, R. E. Hansen and A. Hanssen. *Relative height estimation by cross-correlating ground-range synthetic aperture sonar images*. *IEEE J. Oceanic Eng.*, **32**(4): 971–982, October 2007b.
- T. O. Sæbø, B. Langli, H. J. Callow, E. O. Hammerstad and R. E. Hansen. *Bathymetric capabilities of the HISAS interferometric synthetic aperture sonar*. In *Proceedings of Oceans 2007 MTS/IEEE*. Vancouver, Canada, October 2007c.
- T. O. Sæbø, R. E. Hansen and Ø. Midtgaard. *Filtering of high resolution interferometric synthetic aperture sonar*. In *Proceedings of Underwater Acoustic Measurements 2009*. Nafplion, Greece, June 2009. CDROM (ISBN 978-960-98883-4-9).
- T. O. Sæbø, H. J. Callow and P. E. Hagen. *Pipeline inspection with synthetic aperture sonar*. In *Proceedings of the 33th Scandinavian Symposium on Physical Acoustics*. Geilo, Norway, February 2010. Online (<http://www.iet.ntnu.no/en/groups/akustikk/meetings/SSPA/2010>).
- R. Shiavi. *Introduction to Applied Statistical Signal Analysis*. Academic Press, 525 B. St. Suite 1900, San Diego, CA 92101, USA, 1999.
- G. Shippey, P. Ulriksen and Q. Liu. *Wideband swath bathymetry, SAS autofocus and underwater navigation fixes: Three related problems in echo/image correlation*. In *IAC/IEE Sonar Signal Processing Conference*, pp. 483–491. Weymouth, UK, December 1998.
- M. Soumekh. *Fourier array imaging*. Prentice Hall, 113 Sylvan Avenue, Englewood Cliffs, NJ 07632, USA, 1994.

- M. Soumekh. *Synthetic aperture radar signal processing*. John Wiley & Sons, 605 Third Avenue, New York, NY 10158, USA, 1999.
- M. Stangl, R. Werninghaus, B. Schweizer, C. Fischer, M. Brandfass, J. Mittermayer and H. Breit. *TerraSAR-X technologies and first results*. *IEE Proc. Radar, Sonar Navig.*, **153**(2): 86–95, April 2006.
- S. Suchandt, M. Eineder, H. Breit and H. Runge. *Analysis of ground moving objects using SRTM/X-SAR data*. *Photogramm. Rem. Sens.*, **61**(3-4): 209–224, 2006.
- S. A. Synnes. *SAS multipath analysis*. FFI/RAPPORT To be published, Norwegian Defence Research Establishment, 2008.
- S. A. Synnes, R. E. Hansen and T. O. Sæbø. *Assessment of shallow water performance using interferometric sonar coherence*. In *Proceedings of Underwater Acoustic Measurements 2009*. Nafplion, Greece, June 2009. CDROM (ISBN 978-960-98883-4-9).
- S. A. V. Synnes, H. J. Callow, T. O. Sæbø and R. E. Hansen. *Multipass coherent processing on synthetic aperture sonar data*. In *Proceedings of European Conference on Underwater Acoustic (ECUA) 2010*. Istanbul, Turkey, July 2010.
- J. D. Taylor. *Introduction to ultra-wideband radar systems*. CRC Press, 2000 Corporate Blvd. Boca Raton, FL 33431, USA, 1994.
- C. Tomasi and R. Manduchi. *Bilateral filtering for gray and color images*. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 839–846. 1998.
- R. Touzi and A. Lopes. *Statistics of the stokes parameters and of the complex coherence parameters in one-look and multilook speckle fields*. *IEEE Trans. Geosci. Remote Sensing.*, **34**(2): 519–531, 1996.
- R. Touzi, A. Lopes, J. Bruniquel and P. W. Vachon. *Coherence estimation for SAR imagery*. *IEEE Trans. Geosci. Remote Sensing.*, **37**(1): 135–149, 1999.
- R. J. Urick. *Principles of underwater sound*. Mcgraw-Hill Book Company, Los Altos, CA, USA, third edition, 1983.
- H. L. Van Trees. *Optimum array processing. Part IV of detection, estimation, and modulation theory*. John Wiley & Sons, 605 Third Avenue, New York, NY 10158, USA, 2002.
- L. Wang, A. Bellettini, R. Hollett, A. Tesei, M. Pinto, S. Chapman and K. Gade. *InSAS'00: Interferometric SAS and INS aided SAS imaging*. In *Proceedings of Oceans 2001 MTS/IEEE*, pp. 179–187. Honolulu, HI, USA, November 2001.
- M. Wei and T. Coyne. *Integrated airborne IFSAR mapping system*. In *Proceedings of XXth ISPRS*. Beijing, China, July 2008.

- W. Xu, C. Chang, L. K. Kwok, H. Lim and W. C. A. Heng. *Phase-unwrapping of SAR interferogram with multi-frequency or multi-baseline*. In *Proceedings of the IEEE 1994 International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 730–732. Istanbul, Turkey, July 1994.
- H. A. Zebker and J. Villasenor. *Decorrelation in interferometric radar echoes*. *IEEE Trans. Geosci. Remote Sensing.*, **30**(5): 950–959, 1992.
- L. J. Ziomek. *Underwater acoustics: A linear systems theory approach*. Academic Press, 24-28 Oval Road, London NW1 7DX, UK, 1985.
- J. Ziv and M. Zakai. *Some lower bounds on signal parameter estimation*. *IEEE Trans. Inform. Theory*, **IT-15**(3): 386–391, 1969.

About the author — *Torstein Olsmo Sæbø* was born in 1977 in Bergen, Norway where he grew up. He received the cand.mag. (B.E.) degree in physics in 2000 and the cand. scient. (M.Sc.) degree in astrophysics in 2002, from the University of Oslo, Oslo, Norway. Since 2002, he has been working as a Scientist at the Maritime Systems Division at the Norwegian Defence Research Establishment (FFI), Kjeller, Norway, specializing in the field of synthetic aperture sonar. In 2004 he started to pursue a Ph.D. degree in the field of synthetic aperture sonar interferometry at the University of Tromsø, Department of Physics and Technology. In November 2010, he defended the current Ph.D. thesis, entitled “Seafloor Depth Estimation by means of Interferometric Synthetic Aperture Sonar”. The topic reflects his research interest, which are in the areas of signal processing, synthetic aperture beamforming and interferometry.
