

Estimating 3D motion from radar data

Exploiting an omnidirectional radar array for motion estimation in the context of SAR imaging on agile platforms

Keith T.J. Klein

Estimating 3D motion from radar data

Exploiting an omnidirectional radar array for
motion estimation in the context of SAR
imaging on agile platforms

by

Keith T.J. Klein

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Friday, November 20, 2020 at 9:30 AM.

Student number:	4298292	
Project duration:	March 9, 2020 – november 20, 2020	
Thesis committee:	Prof. Dr. A. Yarovoy,	TU Delft, chair
	Dr. F. Uysal,	TU Delft, supervisor
	Prof. Dr. ir. G.J.T. Leus,	TU Delft
	Dr. ir. J.J.M. de Wit,	TNO

This thesis is confidential and cannot be made public until November 20, 2021.

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

The work presented in this thesis was performed as part of a graduation internship at TNO, the Netherlands.

Company supervisors: Dr. M. Caro Cuenca,
ir. M.P.G. Otten,
Dr. ir. J.J.M. de Wit



Preface

In the midst of a pandemic, working on my thesis was not always easy. I would like to thank everyone who encouraged me when I had doubts, and provided input to my work. Most importantly, I would like to thank my supervisors who, even with minimal offline interaction, closely guided me through the entire process.

I would like to thank my TU Delft supervisor Faruk Uysal, for his valuable input and feedback, and guiding me through the process of submitting a conference paper for my work. I would also like to thank Alexander Yarovoy for providing useful feedback and input at critical points during the course of my project.

During the course of my internship I learned a lot about radar systems, and was able to draw from many of the things I had come across during my study program, allowing me to gain a better fundamental understanding of each individual aspect. Performing my internship at TNO also allowed me to verify my work using actual data obtained from one of their radar systems. I would like to thank my supervisors at TNO, Miguel Caro Cuenca, Matern Otten, and Jacco de Wit, for providing useful feedback and asking critical questions during our meetings, as well as providing the means for verifying my work on experimental data.

Keith T.J. Klein
Delft, November 2020

Abstract

Acquiring an accurate estimate of position is a challenging problem in coherent radar processing techniques such as Synthetic Aperture Radar (SAR). Even more so, for light and agile platforms such as multi-copters. Due to their unpredictable flight path, their motion must be accurately measured during data acquisition, and compensated for during processing. To obtain a focused SAR image, each pulse location must be accurately known to ensure that each pulse is added coherently to the imaging grid. Traditionally, this is achieved with an Inertial Navigation System (INS). While an INS can provide reasonable performance, its weight and size are often a constraint for agile platforms, limiting the available options and attainable accuracy.

In this study, we perform an analysis on the applicability of an omnidirectional radar array for explicitly estimating the motion of a multi-copter platform, and improving on the positioning accuracy achieved by the on board INS. Building on existing 1D SAR motion compensation techniques, we develop new methods for estimating the 3D motion of the radar platform by estimating its height and velocity. In addition, we also present a novel 3D autofocus technique termed multi-beam autofocus. This technique allows for the correction of 3D trajectory errors from pulse to pulse by exploiting the beamforming capabilities of the array, and focusing multiple regions as the image is created.

Using an Extended Kalman Filter (EKF), we obtain position estimates from the radar velocity measurements based on the last known INS position. We experimentally verify that using our velocity estimation method alone, the positioning performance is already improved compared to that of a state-of-the-art INS, allowing for INS-free imaging using arbitrary flight paths. This enables imaging in GNSS-denied environments, and has the potential to further reduce the weight of the platform. We also show that fusing the estimates obtained from our method with the existing INS output yields an additional performance increase in terms of SAR image focus, improving the resolvability and detectability of weak targets. The presented results open further avenues of research, not only in agile SAR imaging but also in autonomous GNSS-denied navigation.

Contents

List of Figures	ix
List of Tables	xi
List Of Abbreviations	xiii
1 Introduction	1
I Background	3
2 FMCW Imaging Radar	5
2.1 Data Acquisition	5
2.1.1 Range Compression	6
2.1.2 Doppler processing	6
2.2 Digital Beamforming	7
2.3 Synthetic Aperture Radar	9
2.3.1 SAR Geometry	9
2.3.2 Data Acquisition	9
2.3.3 Azimuth Resolution	10
2.3.4 Doppler Centroid	10
2.3.5 Backprojection Imaging	10
3 SAR Motion Compensation	13
3.1 Doppler Centroid Estimation	14
3.2 Autofocus	15
3.3 Kalman Filtering	16
3.3.1 Standard Kalman Filter	16
3.3.2 Extended Kalman Filter	17
II Problem Analysis and System Design	19
4 Observable States	21
4.1 Velocity	21
4.2 Attitude	23
4.3 Relative Horizontal Position	24
4.4 Altitude	25
4.5 Summary	27
5 State Estimation Using Radar Data	29
5.1 Height Estimation	29
5.1.1 Change-point Detection	29
5.1.2 Considerations For Experimental Data	32
5.1.3 Overview	34
5.2 Velocity Estimation	35
5.2.1 Velocity Magnitude	35
5.2.2 Doppler Centroid	39
5.2.3 Overview Of Velocity Estimation	46
5.3 Multi-Beam Autofocus	47
5.3.1 The Autofocus Problem	47
5.3.2 Focusing A Single Target	49
5.3.3 Ego-motion Estimation	49
5.3.4 Proof Of Concept	52
5.3.5 Conclusion	54

5.4	Summary	56
6	Sensor Fusion	57
6.1	Fusion Framework	57
6.1.1	System Model	58
6.1.2	State Observations	58
6.1.3	Noise Model	59
6.1.4	Initial Results	60
6.1.5	Tightly Coupled Kalman Filter	62
6.2	Incorporating Multi-beam autofocus	63
6.3	Conclusion	63
III	Evaluation and discussion	65
7	Quantifying Performance	67
7.1	SAR Quality Metrics	67
7.1.1	Peak To Background Ratio	68
7.1.2	Spatial Resolution	68
7.1.3	Image Registration	68
7.2	Short-term Performance	69
7.3	Long-term Performance	73
8	Conclusion and Recommendations	75
8.1	Conclusion	76
8.2	Recommendations	77
	Appendix	79
A	SAR motion Compensation	81
A.1	Doppler centroid estimation	81
A.1.1	Frequency domain Estimation	81
A.1.2	Time domain estimation	82
A.1.3	Image processing techniques	82
A.2	Autofocus Using Row-By-Row coordinate Descent	83
B	Raw Data simulator	85
B.1	Signal model and Simulation geometry	85
B.2	Target Distribution	85
B.2.1	Simulated Trajectories	86
C	Experimental Setup	89
C.1	Radar System Description	89
C.2	Measurement Setup	91
C.2.1	Recording area	91
C.2.2	Radar setup	91
C.2.3	Flight trajectories	91
D	Geometry And Pre-processing	93
D.1	Global Reference Frame and Euler Angles	93
D.1.1	Computing Antenna Positions	94
D.2	Data Processing	95
D.2.1	Downsampling	95
D.2.2	Range Compression And Loss Compensation	96
D.2.3	Doppler Processing	96
	Bibliography	97

List of Figures

2.1	FMCW frequency sweep	5
2.2	Simulated point target before and after applying the keystone transform.	7
2.3	Beamforming on a circular antenna array.	8
2.4	Antenna patterns for different beam configurations.	8
2.5	Geometry for a side-looking stripmap SAR system in both the zero, and non-zero squint case.	9
2.6	SAR imaging illustration	11
3.1	Doppler velocity estimation in automotive and airborne radar systems.	14
4.1	Drone-radar antenna footprint geometry.	22
4.2	Geometry of the radar beam when steering in different directions.	22
4.3	Simulated range-Doppler image for different steering directions.	23
4.4	Simulated range-Doppler image when pitch is non-zero.	23
4.5	Beam footprint when pitch is non-zero.	24
4.6	Computed offset caused by pitch in range-Doppler image.	25
4.7	Principle of radar altimetry in conventional systems and the response from the drone-radar.	25
4.8	Drone-radar range profile on a log-scale.	26
5.1	A comparison between the range profile of the sum of antenna elements for both a coherent, and non-coherent summation.	30
5.2	Estimated change-point in the range profile.	30
5.3	Comparison between gradient and change point detection for altitude estimation.	31
5.4	Digital elevation model of the scene of interest.	33
5.5	Height estimate compensated by Digital Elevation Model.	33
5.6	Radar altimeter flowchart.	34
5.7	Highlighted clutter edge and Doppler centroid.	35
5.8	Range-Doppler spectrum interpolation.	36
5.9	1D spectrum after interpolation.	36
5.10	Comparison between estimated velocity and INS velocity for 45 second segment.	37
5.11	Comparison between estimated velocity and INS velocity for 45 second segment.	37
5.12	Comparison between estimated velocity and INS velocity for 30 second segment.	37
5.13	Velocity magnitude estimation flowchart.	38
5.14	Spectrum overlap illustration.	39
5.15	Range-Doppler image before and after removing strong targets.	40
5.16	Comparison between the 1D spectrum before and after CFAR removal.	40
5.17	Doppler spectrum before and after interpolating axis to equivalent centroid angle.	41
5.18	Doppler spectrum correlated with the nominal antenna pattern.	41
5.19	Estimated centroid angles for a linear simulated data-set.	41
5.20	illustration of measured centroid angles in local reference frame.	42
5.21	Illustration of a Janus configuration using 4 beams on the drone radar.	44
5.22	centroid angle offset caused by non-zero attitude.	44
5.23	Comparison between estimated course and INS course for 45 s segment.	45
5.24	Comparison between estimated course and INS course for 45 s segment.	45
5.25	Comparison between estimated course and INS course for 30 s segment	45
5.26	Flowchart of the entire velocity estimation procedure.	46
5.27	Illustration of pixel matrix generation for multi-beam autofocus.	47
5.28	Induced acceleration error in simulated trajectory.	49
5.29	Point target before and after autofocusing.	50
5.30	Corrected trajectory when only using a single point target.	50

5.31	Illustration of using more focusing regions for trajectory correction.	51
5.32	Illustration of beam steering to different focussing regions.	52
5.33	Geometry used for autofocus simulation.	53
5.34	Estimated trajectory using autofocus with no added noise.	53
5.35	The estimated trajectory error in meters.	54
5.36	The estimated trajectory with an SNR of -27 dB.	54
5.37	Estimation error for an SNR of -27 dB.	55
5.38	A top-level flowchart for the proposed multi-beam autofocus algorithm. The inputs are highlighted in green. The selection of focusing regions is highlighted in orange, as this is currently done manually but could be automated using a detection approach. Note that the projection to the imaging grid is done twice per pulse. Once to interpolate the range bins, and once to update the phase corrections based on the corrected trajectory.	55
6.1	Kalman filter output for simulated trajectories.	60
6.2	Kalman filter outputs for experimental data trajectories.	61
6.3	This figure shows the horizontal error in meters between the radar-only output of the Kalman filter and the Xsens position estimate as shown in Figure 6.2	61
6.4	This figure shows the horizontal error in meters between the height output of the Kalman filter and the Xsens height estimate.	61
6.5	Tightly coupled Kalman filter output for experimental data.	62
6.6	Kalman filter output for experimental data using only inertial sensors.	63
6.7	bias estimation in accelerometers and gyroscopes.	63
7.1	Corner reflector resolution calculation.	68
7.2	Illustration of effects of positioning errors on image registration.	69
7.3	Visual comparison between different sensor configurations for a linear trajectory.	70
7.4	Visual comparison between different sensor configurations for a dynamic trajectory.	71
7.5	trajectories used to compute image focus.	72
7.6	Average PTBR for corner reflectors in the linear trajectory.	72
7.7	Average spatial resolution for linear trajectory	73
7.8	Average PTBR for dynamic trajectory.	73
7.9	Average spatial resolution for dynamic trajectory.	74
7.10	Course error for each trajectory.	74
7.11	The range error for each trajectory	74
B.1	Simulated antenna geometry	86
B.2	Range Doppler image for simulated clutter.	86
C.1	Circular radar designed by TNO	89
C.2	The radar system mounted to a multi-copter platform hovering above the ground.	90
C.3	Photos of the measurement area.	91
C.4	Photos of the drone-radar system.	92
C.5	Images of the experimental trajectories.	92
D.1	The order of rotations as used throughout this thesis. the accents above each axis represent how many times that axis has already been rotated.	94
D.2	The global and local reference frame visualized for the drone hovering above the ground.	94
D.3	The antenna setup used for the tracks in Figure C.5	95
D.4	The average range profile for an experimental data segment 1 second in length. It can be seen that the range loss is approximately quadratic.	96

List of Tables

4.1	Observable kinematic states.	27
5.1	Error statistics of the edge velocity estimator for simulated data.	37
5.2	Error statistics of the centroid estimator using simulated data.	45
6.1	Observations obtained from radar data.	58
C.1	Overview of the system specifications for the radar system developed by TNO.	90
C.2	Overview of measurement accuracy for the Xsens MTi-G-710[1].	90
C.3	Signal parameters used for all recordings.	91

List Of Abbreviations

AHRS	Attitude and Heading Reference System
CA-CFAR	Cell Averaging Constant False Alarm Rate
CDE	Correlation Doppler Estimation
CMQP	Constant Modulus Quadratic Program
DCE	Doppler Centroid Estimation
DEM	Digital Elevation Model
EKF	Extended Kalman Filter
FMCW	Frequency Modulated Continuous Wave
GDE	Geometry-based Doppler Estimation
GNSS	Global Navigation Satellite System
HPBW	Half-Power Beam Width
IMU	Inertial Measurement Unit
INS	Inertial Navigation System
LFM	Linear Frequency Modulated
MCA	Multi-Channel Autofocus
MDA	Map Drift Autofocus
MIMO	Multiple Input Multiple Output
NP	Non Polynomial
PGA	Phase Gradient Autofocus
PRF	Pulse Repetition Frequency
PTBR	Peak-To-Background Ratio
SAR	Synthetic Aperture Radar
SCR	Signal-to-Clutter Ratio
SDE	Sign Doppler Estimation
SDP	Semidefinite Program
SDR	Semidefinite Relaxation
SNR	Signal-to-Noise Ratio
TFMRA	Threshold First Maximum Retracking

1

Introduction

Radar systems are extremely useful and versatile sensors. The information contained within the reflected pulse can be used to derive many parameters about the immediate environment and the targets within it. Since radar is an active sensor using radio frequencies, it can be used in all lighting and most weather conditions. As a result, radar systems are applied in a wide variety of fields ranging from military applications to remote sensing and the automotive industry.

The application of digital signal processing to radar data enables more complex radar modes. A notable example is Synthetic Aperture Radar (SAR). In SAR imaging, a large antenna aperture is synthesized by taking advantage of the radar platform's motion. A fine azimuth resolution can be obtained by coherently combining the contribution of each pulse along the radar trajectory. One of the many challenges in SAR imaging is acquiring an accurate estimate of the relative position of the radar platform along the synthetic aperture. This is less of a problem in satellite and aircraft-mounted SAR systems, where the platform is relatively stable, and its motion can be accurately predicted. In these systems, small phase errors can be compensated using an onboard Inertial Measurement Unit (IMU).

SAR systems based on multi-copter or 'drone' platforms have become the topic of increasing research in recent literature[13, 44, 45]. They are a popular choice for surveillance and object detection in remote areas, as they are easily deployable and highly maneuverable. This makes it possible to quickly create high-resolution images of areas of interest. However, as it is still an emerging field in SAR, several problems need to be solved. One challenge, in particular, is that due to their agility, drone platforms are much more prone to errors in positioning. The positioning system must be able to track quick movements caused by sudden changes in direction and/or perturbations caused by environmental conditions. If left uncompensated, motion errors can cause severe degradation in image quality [38].

For accurate positioning, many IMUs combine measurements from a Global Navigation Satellite System (GNSS) receiver to get an accurate estimate of the position. This is also known as a GNSS-aided Inertial Navigation System (INS). For the sake of brevity, when using the term "INS" throughout this thesis, it will imply a navigation system using both inertial sensors and a GNSS module. The inertial sensors can accurately track short-term movements while the GNSS can compensate for drift. While highly accurate INSes exist, the size and weight of the INS is a constraint for drone platforms. This somewhat limits their accuracy, as shown in the discussion section in [15] by M. Duersch. The radar data itself also contains information about the platform motion and is also used for motion compensation in SAR imaging, providing one-dimensional compensation along the radar line of sight. Motion compensation algorithms can be roughly separated into two groups. Namely, autofocus, which compensates phase errors from pulse to pulse, and Doppler centroid estimation, which compensates a constant velocity component. Notable examples of both methods are phase gradient autofocus by Eichel et al. [16], and sign Doppler estimation by S. Madsen [34].

In principle, the short term accuracy of the IMU, combined with the long term stability of the GNSS module, should provide a fairly accurate estimate of the radar position. However, as previously mentioned, due to weight constraints, the available positioning systems are somewhat limited in accuracy. In practice, this leads to distortions caused by errors in motion estimation. In addition, in case of GNSS outages, the position estimate from an INS drifts very quickly due to bias in the inertial sensors.

To a certain extent, radar-aided positioning is achieved with traditional SAR motion compensation techniques such as autofocus and Doppler centroid estimation. However, for conventional radar systems, these

techniques provide only one-dimensional compensation along the antenna line of sight and are specifically designed for fixed-wing platforms. Without the constraint of a single beam direction, traditional motion compensation techniques can be extended to take advantage of an omnidirectional antenna array to create a 3D estimate of platform motion.

This thesis will investigate whether the accuracy of the on-board positioning system can be improved by incorporating kinematic estimates acquired from radar data using an omnidirectional radar array. We will specifically investigate this in the context of an agile multi-copter platform, or "drone-radar". Improved positioning accuracy will provide better image focus, and increased detectability of weak targets. Accurate positioning also creates the potential for more complex imaging modes, such as interferometric SAR and 3D imaging. Moreover, radar-aided positioning can be useful in GNSS-denied scenarios, where IMU drift can be compensated with accurate estimates from radar data.

The combination of the agile multi-copter platform combined with an omnidirectional radar creates the opportunity for novel signal processing algorithms to be implemented, which take advantage of this unique architecture. In summary, this thesis will answer the following question:

- To what extent can onboard INS-based positioning accuracy be improved by incorporating data from an omnidirectional radar array?

This question can be divided into the following sub-questions:

- What are the current algorithms being employed for positioning and motion compensation in high resolution radar modes?
- Which relevant parameters can be estimated from radar data, and by what means?
- How can the circular radar array be used to improve the accuracy of the parameter estimation techniques?
- How can the estimates from the circular radar array be used to improve the positioning accuracy of the radar platform?

The rest of this thesis will be organized as follows. In the next two chapters, we will present the background information required to answer the research questions. Chapter 2 will give an overview of the signal processing principles behind FMCW radar, and also introduces the concepts of digital beamforming and SAR in this context. In Chapter 3, we look to the literature and analyze the current techniques being employed for SAR motion compensation.

Using the background information as a starting point, Chapter 4 begins to answer the posed research questions by determining how the radar data can be used to obtain information about its motion and position. In Chapter 5, we adapt existing motion estimation techniques, and introduce a processing chain to estimate the height of the radar above ground, and the velocity of the radar in the horizontal plane. As an addition for future work, we also propose a novel 3D coherent positioning algorithm termed multi-beam autofocus which is verified using simulated data. In Chapter 6, we use an Extended Kalman filter (EKF) to obtain position estimates from the velocity observations, and fuse data from different sensors.

The accuracy of the obtained position estimates is evaluated in Chapter 7 by analyzing the focus and registration of subsequent SAR images created with the estimated trajectories. The thesis concludes by reflecting on the obtained results, and providing definitive answers to the posed research questions. We also present a short discussion, analyzing the prospects for future work and providing recommendations for possible additions and improvements.

I

Background

2

FMCW Imaging Radar

FMCW radars have many advantages over pulsed radars and are becoming increasingly popular in modern radar systems. They are even referred to as "radars for the future" in a book on FMCW design by M. Jankiraman [26]. Some of the mentioned advantages are low power usage, lower sampling requirements for short-range intervals, and high range resolution. Since a drone-radar is also a "radar for the future", it makes sense that it too is based on the FMCW principle.

As an introduction to the signal processing techniques presented in later chapters, this chapter gives a general overview of the FMCW principle, and its processing steps. The basic principles behind digital beam-forming and SAR are also explained to provide a foundation for the techniques presented in subsequent chapters.

2.1. Data Acquisition

FMCW radars use continuous frequency-modulated signals instead of pulses for data acquisition. The experimental drone-radar used to verify our method (Appendix C) employs a linear FMCW waveform. A Linear Frequency Modulated (LFM) waveform consists of a series of linear frequency sweeps between a specified bandwidth B , at a fixed sweep rate (Figure 2.1).




Figure 2.1: The principle behind linear FMCW signals. Apart from the up-sweep, there is also a short down-sweep. This is not included in the beat signal, but influences the sweep repetition frequency.

One of the advantages of FMCW radar over pulsed radar, is that it requires less frequent sampling. As the signal is only sampled after mixing, the sampling frequency is only determined by the maximum frequency of the beat signal. To show this, consider the transmitted signal

$$s_t(\tau) = w_r(\tau) \cos(2\pi f_0\tau + \pi K_r\tau^2), \quad (2.1)$$

where f_0 is the sweep start frequency, τ is the fast time index, K_r is the sweep rate, and w_r is the pulse envelope. The received signal is an attenuated, time delayed version of the transmit signal. For a co-located

receiver-transmitter pair, the received signal is given by

$$s_r(\tau, \eta) = A_0 w_r \left(\tau - \frac{2R(\eta)}{c} \right) w_a(\eta - \eta_c) \cos \left(2\pi f_0 \left(\tau - \frac{2R(\eta)}{c} \right) + \pi K_r \left(\tau - \frac{2R(\eta)}{c} \right)^2 \right), \quad (2.2)$$

where η is the slow-time index, and η_c is the slow-time index at the beam center crossing time. The term A_0 is a value which describes the attenuation in amplitude. Upon reception, the receive signal is mixed with the transmit signal

$$s_b(\tau, \eta) = s_t(\tau) s_r(\tau, \eta).$$

By noting that

$$\cos(x) \cos(y) = \frac{1}{2} (\cos(x-y) + \cos(x+y)),$$

it can be seen that the signal after mixing has two components. One of which is a sum of phases, and one of which is a difference of phases. The sum component is removed by filtering the signal and sampling only the difference component. Considering only the phase difference term $\theta(\tau, \eta)$, and simplifying, the phase term can be expressed as

$$\theta(\tau, \eta) = \frac{4\pi f_0 R(\eta)}{c} + K_r \frac{4\pi R(\eta)}{c} \tau - \pi K_r \frac{R(\eta)^2}{c^2}, \quad (2.3)$$

where the first term is a constant phase term determined by the intrinsic phase of the scattering cell, and the range to the scattering cell. The second term is the range dependent beat frequency in rad/s. The third term is an unwanted residual phase term which is a by-product of processing. The residual phase is very small at short range but can be compensated if needed. Since the signal is mixed before sampling, the requirements for the sampling frequency depend linearly on the required unambiguous range interval.

2.1.1. Range Compression

The expression in (2.3) already hints at a method for determining the range to targets. As each range cell has a specific beat frequency, range compression can be performed by analyzing the frequency components in the beat signal using a Fourier transform. The beat frequency has a linear relationship to the range of the scattering cell through

$$f_b(\eta) = \frac{2K_r R(\eta)}{c},$$

where each frequency bin then corresponds to a certain range bin,

$$R(\eta) = \frac{f_b(\eta) c}{2K_r}. \quad (2.4)$$

The range resolution after range compression depends on the bandwidth, B of the transmitted signal, and is given by

$$\Delta_r = \frac{c}{2B}.$$

2.1.2. Doppler processing

In some cases, it is useful to transform the signal to the Doppler domain. This allows for the detection of moving targets, or the estimation of their velocity. In the case relevant to this thesis, the Doppler domain can provide information on the motion of the platform itself through Doppler centroid estimation. Chapter 3 will explain this in further detail. An FMCW signal can be transformed into the range-Doppler domain using a 2D Fourier transform. The fast-time frequency components represent the range compressed signal, while the slow-time frequency components represent the rate of change of the carrier dependent phase.

From (2.3), the carrier-dependent phase of a point target is given by

$$\phi(\eta) = \frac{4\pi f_0 R(\eta)}{c} = \frac{4\pi}{\lambda} R(\eta), \quad (2.5)$$

where λ is the center wavelength of the carrier signal. The Doppler frequency of the target is the first-order derivative of this phase change, and is related to the radial velocity V_r by

$$f_D(\eta) = \frac{1}{2\pi} \frac{d\phi(\eta)}{d\eta} = \frac{2}{\lambda} V_r(\eta). \quad (2.6)$$

A problem that can arise in Doppler processing, is the smearing of targets in the range and Doppler dimensions. This is caused by the coupling which is inherent to the signal. Consider the phase history of the FMCW signal given in (2.3). If the range is approximated by an initial value, velocity and acceleration, the phase history can be written as

$$\theta(\tau, \eta) \approx \frac{4\pi}{c} (f_0 + K_r \tau) (r_0 + \dot{r} \eta + \frac{\ddot{r}}{2} \eta^2). \quad (2.7)$$

This shows that there is a coupling between the slow-time and fast-time dimension. In particular for the linear velocity component. A simple solution is achieved by using the Keystone transform [40]. The Keystone transform compensates for the velocity component by appropriately rescaling the slow time axis for each fast time index using an interpolation filter described by

$$\tilde{\eta} = \frac{f_0}{f_0 + K_r \tau} \eta. \quad (2.8)$$

An example of the Keystone transform applied to a constant velocity point target can be seen in Figure 2.2.

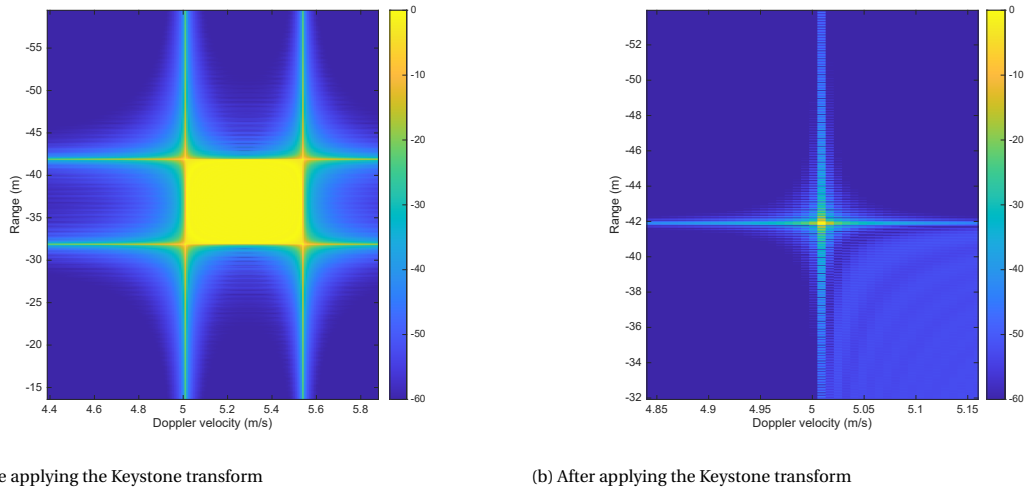


Figure 2.2: These images show the range-Doppler image of a simulated point target at a constant velocity before and after applying the Keystone transform. The Keystone transform focuses the target in range and Doppler at the expense of some interpolation artifacts.

2.2. Digital Beamforming

To obtain angular resolution with any antenna array, some kind of beamforming needs to be applied. Beamforming can be applied either implicitly through an imaging algorithm, or explicitly, by directly combining the signal from multiple antenna elements using a weighted sum. Digital beamforming can be described as the multiplication of the signal from each antenna element with a certain weight vector, as described by

$$\mathbf{y} = \mathbf{u}^T \mathbf{S}.$$

If there are a total of N antenna elements, and M signal samples per element, \mathbf{u} is the $N \times 1$ vector of weights, and \mathbf{y} is the $1 \times M$ output signal. Each row in the $N \times M$ matrix \mathbf{S} corresponds to the signal from a certain antenna element.

The simplest method of applying digital beamforming to an arbitrary antenna array, is using the steering vector. The steering vector is a complex-valued vector, for which the entries specify the phase shift experienced by each antenna element from a plane wave with a specific incidence angle. The steering vector can be computed using the generalized array factor

$$AF = \sum_{n=0}^{N-1} J_n \exp jk(\hat{r} \cdot \mathbf{r}_n), \quad (2.9)$$

where J_n is the complex coefficient associated with the n th element, k is the wavenumber, \hat{r} is the unit vector in the direction of interest, and \mathbf{r}_n is the position vector to the n th element with respect to a fixed reference

point on the antenna array. By taking the center element as the reference point, the components of steering vector are given by the generalized array factor sum components

$$a(\hat{r}) = \begin{bmatrix} \exp jk(\hat{r} \cdot \mathbf{r}_1) \\ \vdots \\ \exp jk(\hat{r} \cdot \mathbf{r}_N) \end{bmatrix}. \quad (2.10)$$

A simple beamformer uses the conjugate of the steering vector as its weights. The amplitudes of the weights can be used to select the active antenna elements. This technique is known as phase-shift beamforming, or delay-and-sum beamforming, and can be intuitively understood as a projection of the antenna elements to a plane with normal vector in the direction of interest. This is illustrated in Figure 2.3 for a circular receiver configuration. More advanced beamformers adjust the amplitude and phase of the weight coefficients to maximize the Signal to Noise Ratio (SNR) or reduce interference from certain directions. Since we are interested in maximizing the response from a given steering direction regardless of signal properties, the phase shift beamformer provides an optimal solution.

The generalized array factor can be used to determine the beam pattern of different receiver configurations. This is shown in Figure 2.4. Here, the value of J_n is determined by the antenna pattern of a single element. From the graph, the best obtainable angular resolution in azimuth is achieved with 9 receiver elements with a Half-Power Beamwidth (HPBW) of approximately 15 degrees.




Figure 2.3: Illustration of beamforming for a circular antenna array using 9 receiver elements. The exact phase shift depends on both the azimuth and elevation angle




Figure 2.4: The simulated antenna pattern using different numbers of receiver elements. The image on the left shows the entire antenna pattern, while the image on the right shows the crossings of the Half Power line for each configuration.


2.3. Synthetic Aperture Radar

Similarly to how angular resolution is improved when using multiple antenna elements, the position of the radar at different time instances can be used as "virtual" antenna elements. This is also known as Synthetic Aperture Radar (SAR). SAR is typically used to create high-resolution, geometrically accurate images of an area of interest.

2.3.1. SAR Geometry

Typical imaging SAR systems are side-looking, and follow a straight flight path. In Figure 2.5a, an illustration of strip map SAR data acquisition is shown for a single point target. The synthetic aperture is defined as the portion of the sensor path where the target is illuminated by the radar beam.

In some cases, the radar beam is not orthogonal to the sensor path. In this case, the beam is said to be squinted. This is illustrated in Figure 2.5b. Non-zero squint causes the zero Doppler time to be offset from the beam center crossing time.



(a) SAR data acquisition geometry for a squint angle equal to zero

(b) SAR data acquisition geometry in the non-zero squint case

Figure 2.5: Geometry for a side-looking stripmap SAR system in both the zero, and non-zero squint case.

The flight path can also deviate from the ideal straight path. For typical airborne SAR systems, squint and non-ideal trajectories are not desired as it increases the complexity of the processing. However, these effects are unavoidable in most cases and are dealt with using motion estimation/compensation techniques. An overview of different SAR motion compensation techniques is given in Chapter 3.

2.3.2. Data Acquisition

As the radar travels along the sensor path, periodic pulses are transmitted with a certain pulse repetition frequency (PRF). For FMCW radar, this is equal to the number of sweeps per second. For each position along the sensor path, the distance to a specific point scatterer on the ground is different. This causes an extra phase modulation of the return signal from pulse to pulse (slow-time). This is also known as Doppler shift.

Since, the carrier frequency of an FMCW signal is removed through mixing, the sampled signal is already demodulated. For a point scatterer at a distance $R(\eta)$ from the radar, where η is the slow-time index, the beat signal is given by

$$s_b(\tau, \eta) = A_0 w_a(\eta - \eta_c) \cdot \cos\{\theta(\tau, \eta)\}. \quad (2.11)$$

In the equation, A_0 is the amplitude of the received signal, and $w_a(\eta - \eta_c)$ is the received signal strength as a function of the antenna beam pattern. The antenna beam pattern can be approximated by a sinc-squared function

$$w_a(\eta(\mu)) = \text{sinc}^2\left(\frac{0.886\mu}{B_{3dB}}\right),$$

where, B_{3dB} is the HPBW, and μ is the angle between the scattering cell and the antenna. During sampling, the lower bound of the PRF should satisfy the Nyquist sampling rate to avoid ambiguities caused by aliasing

of Doppler frequencies. For pulsed radar systems, the PRF should be small enough to allow for unambiguous range sampling. This is not a concern for FMCW radars, as the unambiguous range interval depends on the sampling frequency.

2.3.3. Azimuth Resolution

Before processing, the resolution of a physical antenna with beamwidth B_{3dB} is given by [12]

$$\rho'_a = R(\eta_c)B_{3dB} = \frac{0.886R(\eta_c)\lambda}{L_a}, \quad (2.12)$$

where L_a is the length of the physical antenna, and λ is the wavelength of the carrier signal. Similarly to range compression, the resolution after azimuth compression is approximately given by the reciprocal of the bandwidth. In this case, the Doppler bandwidth, B_d .

$$\rho_a = \frac{1}{B_d} \quad (2.13)$$

2.3.4. Doppler Centroid

The Doppler frequency at the antenna beam center is called the Doppler centroid and is proportional to the rate of change of phase at this position. From (2.3), the carrier-dependent phase of a point target is given by

$$\phi(\eta) = \frac{4\pi f_0 R(\eta)}{c} = \frac{4\pi}{\lambda} R(\eta). \quad (2.14)$$

The Doppler shift at the antenna beam center crossing time, η_c is defined as the rate of change of this phase term and can be written as

$$f_D(\eta_c) = \frac{1}{2\pi} \frac{d\phi(\eta)}{d\eta_c} = \frac{2}{\lambda} V_r(\eta_c).$$

Where V_r is the radial velocity of the platform with respect to the target. The observable Doppler spectrum of a signal depends on the PRF[12]. The maximum unambiguous velocity is given by

$$|V_a| = \frac{\lambda \cdot \text{PRF}}{4}. \quad (2.15)$$

Because of this, the Doppler centroid can be expressed in two parts

$$f_{DC} = f'_{\eta_c} + M_{amb} \cdot \text{PRF}, \quad (2.16)$$

where f'_{η_c} is the fractional part, and M_{amb} is the ambiguity number. Instead of a group of strong targets, the reflected SAR signal typically consists of a distribution of many targets. If these targets are mostly stationary, the Doppler centroid is a projection of the antenna pattern which is centered around the radial velocity of the radar in the steering direction. This will also be demonstrated in Chapter 4.

2.3.5. Backprojection Imaging

While there are numerous imaging algorithms available for synthetic aperture radar, their underlying principle is very similar. The main goal is to map the contributions from a given pulse to the correct location on an imaging grid. This can be achieved using either frequency domain, or time-domain algorithms.

Since the drone-radar should be able to image in any direction for arbitrary flight paths, an algorithm without approximations is desired. This can be achieved using the backprojection algorithm [55]. The backprojection algorithm adds the contribution from each range bin to the appropriate pixels in an imaging grid, while also applying the appropriate phase shift such that the contributions add coherently. Provided that the radar's position is accurately known, motion compensation is implicit in the algorithm itself. When using a radar system with multiple antenna elements, each element can be considered as a separate location in the synthetic aperture. By doing this, beamforming is also implicit in the algorithm itself. This idea is illustrated in Figure 2.6, where both the relative positions of the real antenna elements, and positions along the synthetic aperture contribute to the angular resolution.

Backprojection goes a few steps further than conventional beamforming by calculating the required phase correction for each range bin separately. This is useful in near-range scenarios where a plane wave assumption might not be valid.

The backprojection algorithm can be summarized as follows:




Figure 2.6: SAR imaging for an arbitrary flight path. Both the real array and the virtual array are treated as being part of the total aperture.

- For a given pulse location, the two-way range to each pixel in an imaging grid is calculated. eg. the range from the transmitter to the pixel, and back to the receiver.
- The pulse contributions corresponding to the calculated ranges are determined through interpolation and added to the imaging grid.
- A range dependent phase correction term is applied to each pixel. This ensures that signal contributions from different pulse locations add coherently when they correspond to the same location on the imaging grid.
- This process is repeated for each pulse location and antenna element until the entire synthetic aperture has been processed.

The interpolated, phase-corrected signal at the n^{th} pulse location, for a given pixel location ρ_n is given by

$$P_n(\rho_n) = \text{IP}(r_c(n), \rho_n) \cdot \exp j(\phi_{corr}(\rho_n)), \quad (2.17)$$

where IP represents the non linear interpolating function, which maps the range compressed signal r_c to the correct ranges as specified by the coordinates of the imaging grid. The pixel location ρ_n is a vector from the antenna location to the pixel of interest. ϕ_{corr} is the phase correction factor which is determined by the range to a pixel in the imaging grid, and is given by

$$\phi_{corr}(\rho_n) = -\frac{4\pi f_0 R(\rho)}{c}. \quad (2.18)$$

The imaging process can then be expressed by

$$I_n = I_{n-1} + P_n, \quad (2.19)$$

where I_n represents the SAR image after adding the n^{th} interpolated, phase corrected pulse. It should be noted that the pulse contributions do not need to be added sequentially. If permitted by memory limits, multiple pulses can be processed in parallel for increased performance. A more detailed explanation of the backprojection algorithm, along with a basic Matlab framework, is given by Gorham et al. in [21].

3

SAR Motion Compensation

The previous chapter outlined the signal processing principles behind FMCW radar and introduced the basics of SAR imaging. The flight path in most airborne SAR systems is not ideal and must be appropriately compensated by estimating the motion of the radar platform. This is especially true for agile platforms such as multi-copters. The goal of SAR motion compensation is to estimate the motion of the radar platform with enough accuracy to improve overall imaging quality. Aside from techniques specifically designed for SAR motion compensation, there are several proposed methods in literature from a broad spectrum of radar systems. In most cases, these approaches are based on reference points in the radar return signal, utilizing constant false alarm rate (CFAR) detection and image processing techniques to determine the relative change in position and velocity of strong reference points. The relative position of the radar platform is determined by tracking the changes of visual features in range compressed radar data.

One example of this is scan-matching [20], where the relative motion of a scene from scan to scan is used to determine the position of the radar platform. The drawback of a scan-matching approach is the requirement for scene contrast. Also, successive scans need to be sufficiently correlated to determine relative motion. Improved scan matching algorithms have also been presented. In [8], Cen et al. propose a keypoint extraction algorithm for more accurate scan matching. Another approach is proposed by Barnes et al. [5], where machine learning is used to detect suitable keypoints and suppress multipath effects for improved correlation performance.

Alternatively, artificial keypoints with known locations can be used. In [10], Clark et al. use a set of omnidirectional reflectors with known locations as beacons. Their relative locations are then used as observations in a Kalman filter to obtain an estimate of the radar's position. An addition to this approach is presented in [9], where a method for navigation using natural features is presented.

Most of the previously mentioned motion estimation techniques are designed for automotive radar systems. Radar-aided positioning for multi-copter platforms is less common in literature. However, there are some proposed techniques. An example is a radar-aided positioning system for GNSS-denied environments by Mostafa et al. [37], which estimates the forward velocity of a multi-copter platform from radar data and fuses it with other sensors. A similar approach is presented by Scannapieco et al. [43], where the motion for a lightweight drone is estimated using the range and bearing to targets. A key difference between the systems used in these methods and the study presented in this thesis, is the geometry of the radar system. The drone-radar provides an instantaneous 360-degree field of view while the methods presented in the literature are forward-looking linear arrays.

The phase information in a radar pulse also contains useful information about the state of the radar platform, specifically, its velocity. By measuring the Doppler spectrum of the reflected ground surface, an estimate of velocity can be made. The accuracy of the velocity estimate can be increased by incorporating multiple sensors facing in different directions. This was also the working principle behind Doppler navigation systems used in commercial aviation in the 1960's [19] (Figure 3.1b). In more recent literature, this technique has also been applied to automotive radar sensors that measure the road surface [23, 29, 54] (Figure 3.1a). In principle, this technique can also be applied to the drone-radar. However, for a drone platform, analysis of the Doppler spectrum is complicated by non-homogeneous clutter and terrain elevation. Moreover, quick changes in position and attitude may also deteriorate the estimation accuracy. On the other hand, the multi-channel architecture of the drone-radar allows for a large number of digital beams which can improve the

overall estimation accuracy.

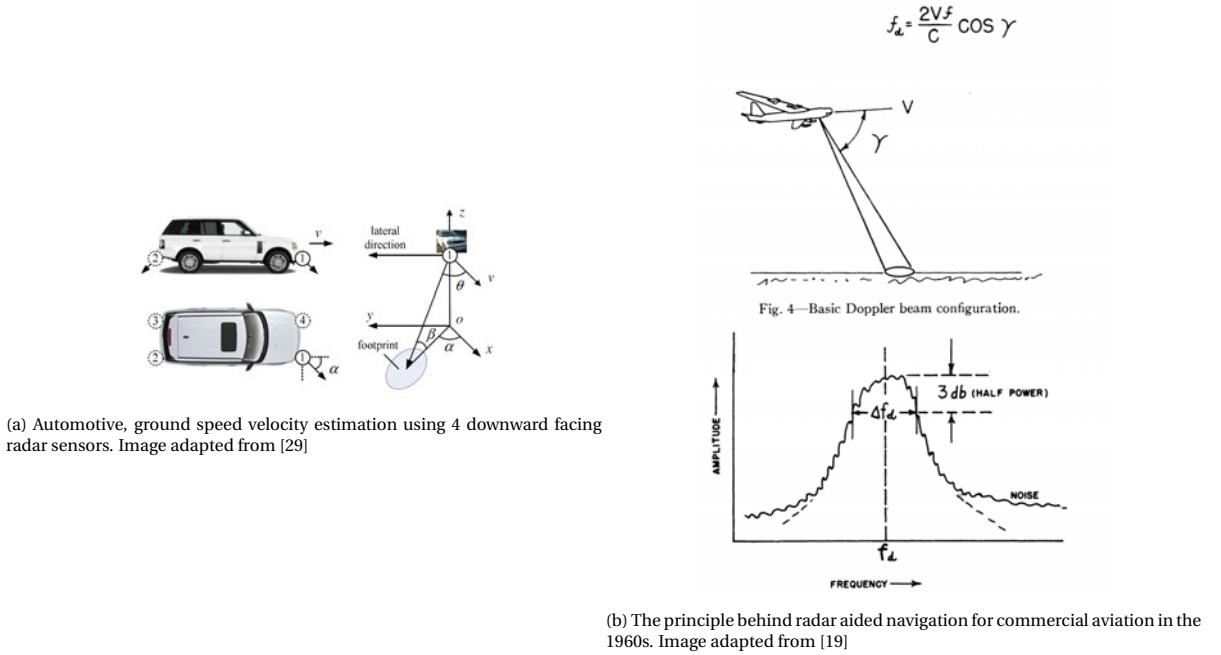


Figure 3.1: Doppler velocity estimation in automotive and airborne radar systems.

As improving SAR imaging quality is one of the main motivators of this thesis, conventional SAR motion compensation techniques will be used as a starting point for the techniques developed in later chapters. To a certain extent, this is also a form of motion estimation. The difference being that conventional SAR motion compensation techniques typically only estimate one-dimensional motion along the radar line of sight. Regardless, these techniques can provide some insight into the attainable positioning accuracy for airborne radar systems. There are two main categories of motion compensation in SAR imaging. These are Doppler centroid estimation (DCE) and autofocus. This chapter will give an overview of different techniques for DCE and autofocus, and explore their applicability to the drone-radar system.

3.1. Doppler Centroid Estimation

DCE techniques perform first-order motion compensation by estimating the centroid of the Doppler spectrum in SAR data. This gives an estimate of the velocity of the radar in its line of sight. For a traditional side-looking SAR system, the Doppler centroid may be non-zero due to the unintentional squint of the antenna. For the drone radar, the Doppler centroid varies according to the velocity of the platform and is different for each direction, and each digital beam. In this case, DCE can be used to estimate the velocity vector of the radar platform by combining the centroid estimates from different antenna directions. There are many different centroid estimation algorithms in various signal domains. These are time-domain, frequency-domain, and image-domain.

In [4], M. Bamler shows that frequency domain DCE methods can be grouped into a framework of "Correlation-based estimators". In this framework, the Doppler centroid is found by correlating the azimuth power spectrum with some weighting function

$$D(f) = \int A(g)B(g-f)dg, \quad (3.1)$$

where, A is the azimuth power spectrum, and B is the weighting function. The DC estimate is then found at the zero point of the correlation result

$$D(\hat{f}_{dc}) = 0.$$

The estimation techniques then differ by the weighting function used for estimation, with the most notable being energy balancing, matched-correlation and maximum-likelihood.

The Doppler centroid of a signal can also be estimated by measuring the change of phase from pulse to pulse in the time domain [12]. For a given azimuth sample $s(\eta)$, the phase difference between another sample separated by $\Delta\eta$ is given by the angle of the average cross correlation coefficient at lag 1

$$C(\Delta\eta) = \sum_{\eta} s^*(\eta)s(\eta + \Delta\eta).$$

The fractional part of the Doppler centroid is then given by

$$f'_{\eta_c} = \frac{F_a}{2\pi} \angle(C(\eta)).$$

In literature, this method is referred to as correlation Doppler estimation (CDE).

Aside from frequency and time domain methods performed on echo data. There are also several Doppler centroid methods that are based on approaches related to image processing. Namely, a method based on slope estimation by Young-Kyun et al. [56], and an improved version by I. Cumming and S. Li [11]. Here, the Doppler centroid is estimated using the radon transform and is related to the slope in the range profile of strong targets. Another example is morphological edge detection by Mao et al. [35], where the forward velocity is estimated using the clutter edge. Long et al. propose a method based on image entropy as a cost function based approach [32]. A more detailed overview of different centroid estimation techniques can be found in Appendix A.1.

Frequency domain techniques are more flexible for the drone-radar and can be easily manipulated to account for effects from the antenna pattern and high squint angles. This will be shown in Chapter 5. In addition, the matched-correlation doppler centroid estimator achieves the lowest variance in experimental scenarios as shown by W. Yu, and Z. Zhu in [52]. In Chapter 5, we will modify and combine the ideas of matched-correlation and clutter edge detection to develop a generalized velocity estimation technique for the drone-radar.

3.2. Autofocus

Autofocusing of SAR images is generally used to correct phase errors from pulse to pulse. As the contribution from each pulse is a complex number with a magnitude and phase, small motion errors can cause these pulses to add incoherently. With autofocus, the image quality is increased by estimating the phase error and applying the corresponding phase correction factor. A drawback of most autofocus techniques is that the total phase error should be smaller than a single resolution cell.

An intuitive, generalized formulation of the autofocus problem is presented by both J. Ash [3], and Dueresch et al. [14], in their papers on autofocus for the backprojection algorithm. The problem can be understood as follows. Given a group of M pixels

$$A = \{a_1, a_2, a_3, \dots, a_M\},$$

each pixel can be expressed as the sum of N pulses

$$a_m = \sum_{n=1}^N R_{nm} e^{j\delta_n},$$

where R_{nm} is a complex term denoting the complex sample corresponding to the given pixel and δ_n is a phase error term present in the n_{th} pulse. Assuming that the same phase error is present in each pixel for a given pulse, the phase error can be estimated by minimizing some cost function $F(A)$. One example of such a cost function is the image sharpness which has been shown to be well suited for SAR images by T. Schulz [46]. The optimal sharpness function is defined as the sum of squared pixel magnitudes

$$F(\phi_{\delta_N}) = - \sum_{m=1}^M a_m a_m^*, \quad (3.2)$$

where $*$ denotes complex conjugation. The set of estimated phase errors is then given by

$$\hat{\delta}_N = \underset{\delta_N}{\operatorname{argmin}} F(A). \quad (3.3)$$

There are several proposed solutions to solve the autofocus problem. With the most notable ones being phase gradient autofocus (PGA) [49], and map drift autofocus (MDA) [7]. PGA is a non-parametric autofocus

technique and can estimate arbitrary phase errors, while MDA is a model-based technique that estimates linear or quadratic phase error between different sub-apertures. PGA operates on defocused SAR images and assumes a Fourier relationship between the image domain and range compressed data [3]. The optimization problem in (3.3) can also be solved on a per pulse basis. As it has no closed-form solution it must be solved using iterative methods such as coordinate descent. In [3], J. Ash derives a closed-form solution for each coordinate using a geometrical interpretation. The approach by Duersch et al. [14], also uses coordinate descent. However, a line search is used to optimize each coordinate.

Autofocus algorithms do not explicitly estimate the motion of the radar platform. However, they can be used to estimate 3D trajectory deviations as proposed by Ran et al. [41]. Here, phase estimates from several local regions in a SAR image are combined to obtain the 3D phase error from pulse to pulse. Another promising approach is presented by A. Sommer and J. Ostermann [48], where estimated phase errors are used to calculate equivalent range errors, which are propagated along the trajectory at each pulse location. This allows for the estimation and correction of much larger errors.

In Chapter 5, we will expand on the principle of local images for trajectory deviation to exploit the omnidirectional field of view, and multi-channel architecture of the drone-radar. We will also combine this with the idea of error propagation to develop an omnidirectional autofocus algorithm that can both focus a SAR image, and estimate the sensor path.

3.3. Kalman Filtering

While not typically applied in conventional SAR processing, the goal of this thesis is to combine estimates from both the INS and radar in an optimal way. When dealing with observations from multiple sources that all contain information about a system's state, some kind of "sensor fusion" is required to obtain the final state estimate. When the parameters are corrupted by additive Gaussian noise, the maximum likelihood estimator of the parameter is given by a covariance weighted average of the observations [28].

In practical scenarios, there may also be prior information on the state of a system. In such cases, a bayesian approach can be adopted. For example, using a model, the state of a system at the next time step can be predicted with some uncertainty depending on the model's accuracy. The posterior likelihood of the system state is then obtained by multiplying the likelihood function given by the measurements, and the likelihood function given by the prediction. The optimal estimate is given at the maximum of this posterior likelihood function.

In essence, this is also what a Kalman filter does under the assumption of Gaussian probability distributions. In this section, the principle behind the Kalman filter for sensor fusion will be introduced using a simple linear model.

3.3.1. Standard Kalman Filter

Given a system model without inputs, the system state equation is given by

$$\mathbf{x}_{t|t-1} = \mathbf{F}_t \mathbf{x}_{t-1|t-1} + \mathbf{w}_t, \quad (3.4)$$

where \mathbf{x} is the state vector, \mathbf{F} is the state transition matrix which predicts the current state using the previous state, and \mathbf{w}_t is the process noise which can be used to model the uncertainty of the system model. The subscripts in the equation specify the current time step, and which time step a state estimate is based on. $\mathbf{x}_{t|t-1}$ denotes the state estimate at time step t based on a prediction from time step $t-1$. On the other hand $\mathbf{x}_{t-1|t-1}$ denotes the state estimate in time step $t-1$ after updating with measurements in time step $t-1$.

If the 3D position and velocity are included, the state vector is given by

$$\mathbf{x} = [x \quad \dot{x} \quad y \quad \dot{y} \quad z \quad \dot{z}]^\top. \quad (3.5)$$

In addition, there are also a number of measurements which provide information on the state of the system. These are described by the measurement equation

$$\mathbf{z}_t = \mathbf{H} \mathbf{x}_t + \mathbf{v}_t, \quad (3.6)$$

where \mathbf{z}_t is the measurement vector, \mathbf{v}_t is the measurement noise, and \mathbf{H} is the measurement matrix which maps the observations to the system states. The process noise and measurement noise are distributed according to covariance matrices \mathbf{Q}_t and \mathbf{R}_t respectively. Kalman filtering is a sequential algorithm which can

be roughly divided into two parts at each iteration. Namely, a prediction step, and an update step. The prediction step calculates the prior state estimate using the state transition matrix (3.4). In addition, a prediction of the covariance matrix of the state estimate is also done through

$$\mathbf{P}_t = \mathbf{F}_t \mathbf{P}_{t-1} \mathbf{F}_t^\top + \mathbf{Q}_t. \quad (3.7)$$

The update equations then incorporate the measurements to compute the final state estimate:

$$\mathbf{x}_{t|t} = \mathbf{x}_{t|t-1} + \mathbf{K}_t (\mathbf{z}_t - \mathbf{H} \mathbf{x}_{t|t-1}), \quad (3.8)$$

$$\mathbf{P}_{t|t} = \mathbf{P}_{t|t-1} - \mathbf{K}_t \mathbf{H} \mathbf{P}_{t|t-1}, \quad (3.9)$$

where \mathbf{K}_t is the Kalman gain matrix, which quantifies the magnitude of the update step and weighs the contributions from each sensor on the state estimate. The Kalman gain is given by

$$\mathbf{K}_t = \mathbf{P}_{t|t-1} \mathbf{H}^\top (\mathbf{H} \mathbf{P}_{t|t-1} \mathbf{H}^\top + \mathbf{R}_t)^{-1}. \quad (3.10)$$

An intuitive derivation of the update equations (3.8) and (3.9), is given by R. Faragher in [18].

3.3.2. Extended Kalman Filter

In radar systems in particular, the kinematic states are not observed in Cartesian coordinates. In most cases, the position of a target is expressed in either polar or spherical coordinates. It will become apparent in later chapters that this is also true when estimating the states of the radar system. Instead of a linear relationship between the measurements and the states through \mathbf{H} , a non-linear function $h(\mathbf{x}_t)$ now holds. Consider the following observations:

$$\mathbf{z}_t = [|\mathbf{v}_h| \quad \theta_C]^\top \quad (3.11)$$

where $|\mathbf{v}_h|$ is the velocity magnitude, and θ_C is the course angle. If the measurements are converted into the state vector coordinate system, the measurement covariance matrix might not be Gaussian. The relationship between the observations and the state in (3.8), is now given by

$$\mathbf{z}_t = h(\mathbf{x}_t) = \begin{bmatrix} \sqrt{\dot{x}^2 + \dot{y}^2} \\ \arctan \frac{\dot{y}}{\dot{x}} \end{bmatrix}. \quad (3.12)$$

In the other equations, the measurement matrix is replaced by a time dependent matrix \mathbf{H}_t , given by the Jacobian of h at the predicted state:

$$\mathbf{H}_t = \begin{bmatrix} 0 & \frac{\dot{x}}{|\mathbf{v}_h|} & 0 & \frac{\dot{y}}{|\mathbf{v}_h|} & 0 & 0 \\ 0 & -\frac{\dot{y}}{\dot{x}^2 + \dot{y}^2} & 0 & \frac{\dot{x}}{\dot{x}^2 + \dot{y}^2} & 0 & 0 \end{bmatrix}. \quad (3.13)$$

As the variance of the course estimate and the variance of the velocity magnitude will be based on different metrics, we will employ the extended Kalman filter framework. This will be detailed in Chapter 6.

II

Problem Analysis and System Design

4

Observable States

In the previous chapter, we discussed the motivation and background behind the problem, and some insight was given into the signal processing aspects of an FMCW imaging radar. In addition, several different tools for SAR motion compensation were presented. This chapter goes into more detail by examining the kinematic states of the radar, and determining which states can be observed through radar measurements.

There are two main properties of the radar signal that can be exploited, namely, the magnitude and phase of the range compressed signal. The magnitude is typically used for range measurements from pulse to pulse, while the phase of the signal is typically analyzed over multiple pulses to obtain an estimate of velocity.

Throughout this chapter, simulated range-Doppler images are shown to illustrate and present different ideas. The data used to create these images was simulated by a developed raw signal simulator, which is presented in detail in Appendix B. An overview of the reference frames and angles used to describe position and attitude is given in Appendix D.

4.1. Velocity

The velocity of the radar can be estimated through measurements of the radial velocity in each direction. The range-Doppler domain describes the projection of the radial velocity at a particular steering angle. When observing the Doppler spectrum of the radar signal, its distribution depends on both the motion of the targets in the scene and the motion of the radar platform itself. If the illuminated scene is relatively homogeneous, the observed spectrum is a projection of the antenna pattern. This projection contains information on the platform velocity and is the working principle behind frequency domain Doppler centroid estimators.

In conventional airborne SAR systems, the radar platform is relatively far from the scene being imaged. Because of this, the Doppler centroid does not vary significantly as a function of range. For a drone platform, however, the radar is relatively close to the scene of interest. As a result, the Doppler centroid varies significantly as a function of both range and cross-range. To understand how the velocity and orientation of the radar influence the measured Doppler centroid, we will use a simplified geometrical framework based on the following assumptions:

- The elevation beamwidth along the radar line of sight is large enough such that the scene is illuminated from directly below the radar to the largest range of interest.
- The azimuth beamwidth after steering is small enough such that there is a distinct beam center at each range.

Based on these two assumptions, the center of the beam footprint can be modeled as the intersection of the ground plane with the plane defined by the axis pointing along the radar line of sight and its z-axis. This is illustrated in Figure 4.1. The measured radial velocity is the scalar projection of the velocity vector and is given by

$$V_r = \mathbf{v} \cdot \frac{\mathbf{r}}{|\mathbf{r}|}. \quad (4.1)$$

Where \mathbf{v} is the velocity vector, and \mathbf{r} is the vector from the radar to a point on the intersecting line. It is also important to note that the projection also depends on the angle of the steering direction with respect to the

angle of the velocity vector. This is illustrated in Figure 4.2, and shown through simulation in Figure 4.3. When the difference between the steering direction and the velocity vector is smaller than half the beamwidth, as in the forward-looking case, the spectrum overlaps. This is because the projected velocities on both sides of the beam have the same sign. Because of this, the beams aligned with the velocity vector contain less information on the course of the radar and will have a biased centroid. This must be taken into account when estimating the velocity.

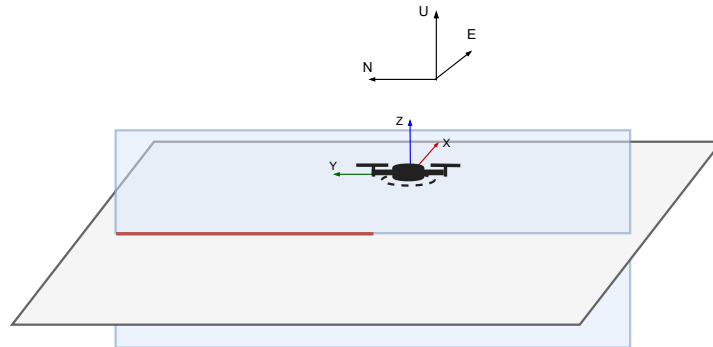


Figure 4.1: This figure illustrates the footprint of the antenna beam center as the intersection between two planes. The radar line of sight is aligned with the y axis and the attitude is aligned with the inertial East-North-Up(ENU) reference frame.





Figure 4.2: Geometry of the radar beam when steering in different directions.




(a) Range Doppler image when the beam is steered to 0 degrees (East North) (b) Range Doppler image when the beam is steered to 90 degrees (East North)

Figure 4.3: A comparison between the range-Doppler image when the radar has a constant velocity in the x-direction. In this simulation, the velocity of the radar is 5 m/s, and all attitude angles are 0. For beamforming, 9 antenna elements are used. 192 pulses are used to create the Doppler image.

4.2. Attitude

While the attitude of the radar is not directly observed, a non-zero attitude affects the perceived value of the radial velocity. To show this, we consider the same scenario as shown in Figure 4.3. This time, however, the radar has a positive non-zero pitch of 15 degrees. The effect of non-zero pitch can be seen in Figure 4.4. The range-Doppler spectra show that a non-zero pitch has almost no impact on the perceived radial velocity aligned with the velocity vector. On the other hand, a non-zero pitch has a significant, range-dependent impact on the perceived radial velocity orthogonal to the velocity vector. This can be explained using the simplified geometrical framework.



(a) Range Doppler image when the beam is steered to 0 degrees (East North) (b) Range Doppler image when the beam is steered to 90 degrees (East North)

Figure 4.4: A comparison between the range-Doppler image when the radar has a constant velocity in the x-direction, and non-zero pitch. In this simulation, the velocity of the radar is 5 m/s, and the pitch angle is 15 degrees. For beamforming, 9 antenna elements are used. 192 pulses are used to create the Doppler image.

When the radar has a non-zero attitude, the intersecting plane is rotated as well. This causes the position and orientation of the centerline to change. For a rotation around the axis normal to the plane (x-axis in Figure 4.1), the intersecting line does not change. However, a rotation around the axis defined by the steering direction causes the intersection to shift as a function of the rotation angle and radar altitude. This is shown in Figure 4.5.

To determine the offset in estimated velocity caused by a non-zero attitude, we first consider the case where the radar is moving in the East direction, and the heading is zero (Figure 4.5). For a beam steered in the y-direction, the observed velocity should be zero. However, there is an offset caused by a rotation around the y-axis. In this case, the offset is given by

$$v'(\theta_{st}, r) = v_x \frac{h \tan \beta}{r} \Big|_{\theta_{st} = \frac{\pi}{2}}, \quad (4.2)$$

where β is the rotation around the y-axis, h is the height of the radar above the ground plane, and r is the

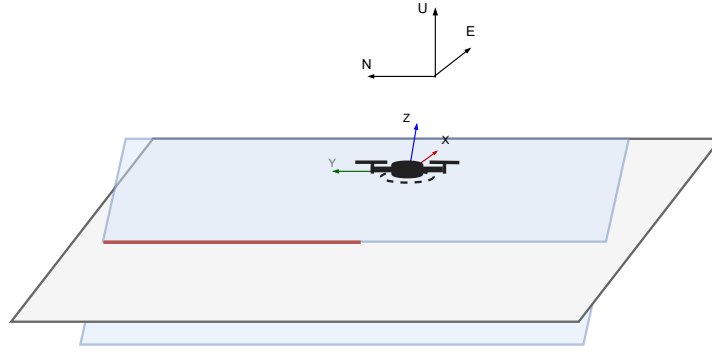


Figure 4.5: This figure illustrates the center of the antenna footprint when the radar is rotated along the y-axis. The centerline shifts as a function of the rotation angle and radar altitude.

range. This result can be generalized to a velocity offset for an arbitrary attitude by defining a velocity perpendicular to the steering angle v_p , and a rotation β_{st} around the axis aligned with the steering angle θ_{st} . For a given horizontal velocity vector \mathbf{v}_h , The offset caused by the pitch and roll in a given steering angle, and range bin, can be found by analyzing the rotation matrices

$$\begin{aligned} \mathbf{R}_{yx'} = \mathbf{R}_y(\beta)\mathbf{R}_{x'}(\gamma) &= \begin{bmatrix} \cos\beta & 0 & \sin\beta \\ 0 & 1 & 0 \\ -\sin\beta & 0 & \cos\beta \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\gamma & \sin\gamma \\ 0 & -\sin\gamma & \cos\gamma \end{bmatrix} \\ &= \begin{bmatrix} \cos\beta & \sin\beta\sin\gamma & \sin\beta\cos\gamma \\ 0 & \cos\gamma & -\sin\gamma \\ -\sin\beta & \cos\beta\sin\gamma & \cos\beta\cos\gamma \end{bmatrix}. \end{aligned} \quad (4.3)$$

where the subscript indicates the order of rotation. In this case, the first rotation is around the x-axis, and the second rotation is around the rotated y-axis (y'). When the beam is steered in a particular direction, this can be seen as a yaw rotation in the local reference frame around the rotated z-axis z'' . This can be described by the rotation matrix

$$\mathbf{R}_{yx'z''} = \mathbf{R}_{yx'} \cdot \begin{bmatrix} \cos\theta_{st} & -\sin\theta_{st} & 0 \\ \sin\theta_{st} & \cos\theta_{st} & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (4.4)$$

The angle β_{st} is equal to the angle between the rotated y-axis y'' and the ground plane. This can be found by computing the angle between the second column of $\mathbf{R}_{yx'z''}$, and the normal vector of the ground plane. β_{st} is found to be

$$\beta_{st} = \frac{\pi}{2} - \arccos(\sin\beta\sin\theta_{st} + \cos\beta\sin\gamma\cos\theta_{st}). \quad (4.5)$$

The velocity v_p on the other hand, requires the velocity vector itself to be known, and is given by,

$$v_p = |\mathbf{v}_h| \sin\delta \quad (4.6)$$

where δ is defined as the angle between the steering direction and the velocity vector, and $|\mathbf{v}_h|$ is the absolute value of the velocity vector in the horizontal plane. This implies that the actual non-offset velocity must be known to estimate the offset caused by the attitude. Fortunately, we will show in Chapter 5 this can be substituted by using the initial biased centroid estimate to calculate the velocity vector for a negligible loss of accuracy. When moving forward, drones typically pitch down. Because of this, this effect must be taken into account when estimating velocity. Chapter 5 will also show how the omnidirectional array can be used to suppress the effects of attitude.

4.3. Relative Horizontal Position

In typical scenarios, it is not possible to determine the horizontal position of the radar in a global reference frame. Doing so requires all pulses to be sufficiently correlated to an initial reference pulse. What is more

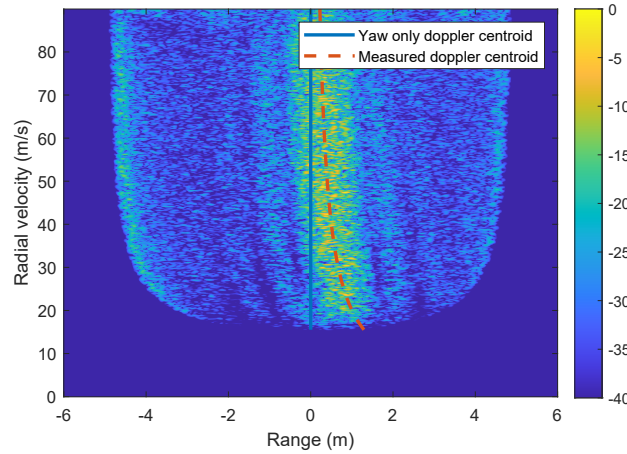


Figure 4.6: The actual perpendicular velocity component when moving forward versus the measured Doppler centroid due to non-zero pitch plotted onto the range-Doppler image. For this simulation, the altitude, pitch, and forward velocity are 15m, 15 degrees, and 5m/s respectively.

commonly done in literature, is the estimation of relative motion from pulse to pulse. This is done using the magnitude in scan matching approaches, and using the phase in autofocus. In essence, this is also a form of measuring velocity. A requirement for this approach to achieve adequate performance is contrast in the scene. Contrast can be used as reference points for the phase and magnitude from pulse to pulse. This thesis proposes an alternative relative positioning method termed multi-beam autofocus, for positioning. This approach will be presented in the next chapter.

4.4. Altitude

If the terrain is relatively flat, it is possible to determine the absolute position in the upwards direction, as the reference point stays the same. This technique is also known as radar altimetry and is widely used in satellites and airplanes. Most literature on radar altimetry deals primarily with long-range altimeters and the complications associated with it. Short-range altimetry is not a difficult problem by itself. However, being mostly side-looking, the drone-radar is not designed as an altimeter and requires a slightly different processing approach.

In conventional satellite altimeters, the impulse response of the surface is characterized by the so-called "Brown model", which was introduced by G. Brown [6]. This model parametrizes the leading edge and trailing edge of the surface reflection based on the pointing angle and the surface roughness to form an optimal matched filter. As the drone-radar has a wide beam and is not strictly down-looking, the leading edge has a relatively gentle slope. Figure 4.7 shows the impulse response predicted by the brown model for different nadir angles compared to the range profile of a pulse from the drone radar.




Figure 4.7: The left image shows the surface response as predicted by brown for a typical nadir looking pulsed radar system. In the image the effect of changing the nadir angle ξ is also shown. The right image shows the range profile of an arbitrary pulse from the drone-radar. As expected, the wide beam and side looking architecture result in a relatively gentle slope.

It is difficult to design a matched filter for the range profile due to the noisy response. However, there is still a clear difference in power level. If we instead observe the range profile on a log-scale, the change point is much sharper and resembles a step function, this is shown in Figure 4.8. In Chapter 5, we will show that the change point detector proposed by M. Lavielle [30], produces a stable estimate of the height above the surface for the drone-radar.




Figure 4.8: The range profile of a single radar pulse presented on a log-scale. The first reflecting target can easily be recognized by an order of magnitude change in received signal power.

While the altitude of the radar is technically observable, it will become clear that the terrain elevation is not negligible and has a significant impact on the quality of the estimate.

4.5. Summary

In this chapter, we analyzed the observability of the kinematic states of the radar platform in the radar data. More specifically, this chapter observed how velocity, attitude and altitude are represented in the radar data. The radar can be used as a velocity sensor, estimating the velocity using either the Doppler shift of the radar over multiple pulses, or relative changes in reference points from pulse to pulse. In addition, the radar sensor can also be used as an altimeter by analyzing the range profile and determining the closest reflection. Table 4.1 gives an overview of the states of the system and specifies which states are observable using radar data.

Table 4.1: The kinematic states of the radar and their observability using either radar data or INS measurements.

State	Radar		INS	
	Observable	Method	Observable	Method
Position (E, N)	No	-	Yes	GNSS,IMU
Position (U)	Yes	Radar altimeter	Yes	GNSS,IMU,barometer
Velocity (E, N)	Yes	Doppler centroid estimation, scan-matching, autofocus	Yes	GNSS,IMU
Velocity(U)	Yes	Radar altimeter	Yes	GNSS,IMU
Yaw (α)	No	-	Yes	Magnetic compass, GNSS
Pitch(β)	Partly	Doppler centroid	Yes	IMU
Roll(γ)	Partly	Doppler centroid	Yes	IMU

5

State Estimation Using Radar Data

The previous chapter performed an analysis on the observability of the various kinematic states of the radar and analyzed the effect of non-zero attitude on the velocity estimate. This chapter will further examine the radar data, and derive the appropriate methods to extract the desired parameters. The chapter will be organized as follows.

The first section will focus on estimating the height of the radar above ground by formulating the problem as a change point detection problem. In the second section, a velocity estimator is proposed to estimate the velocity of the radar in the horizontal plane. Simulated data will be used to verify the performance of the velocity estimator for different antenna configurations. Using experimental data, a qualitative comparison between the INS velocity estimate, and radar velocity estimate will also be given. A detailed description of the radar system and the experimental setup can be found in Appendix C.

In the final section, we will present a novel autofocus method, specifically designed for the omnidirectional radar, which can estimate relative position using phase changes from pulse to pulse.

5.1. Height Estimation


While the drone radar is not designed as an altimeter, it can still be used to estimate the height above the ground by exploiting the characteristics of the received signal. Since the returns from directly below the radar are usually observed first in the range profile, they are easily distinguishable through a sharp power level change in the range profile. This can be used to estimate the height of the radar with respect to the ground surface.

The range profile obtained from the sum of antenna elements can be seen in Figure 5.1 for simulated data. In this figure, a comparison is shown between a non-coherent summation and a coherent summation of the receiver antenna elements. While coherent summation is a powerful tool in radar signal processing, it is not always the optimal approach. Because of the random phase distribution of the ground reflections, a coherent summation randomly adds contributions both constructively and destructively, resulting in a noisier edge. On the other hand, a non-coherent addition lacks interference and only adds magnitudes. Because of this, the channels are added non coherently when estimating height.

5.1.1. Change-point Detection

Since the ground reflection is observed as a change in power level instead of a peak, the location of the change point between power levels needs to be found. One possible way of doing this is through the gradient, similar to an edge detection problem. The problem with this approach is that the performance of the gradient method depends on the sharpness of the edge and the amount of noise. In this case, a better approach is to find the change point in the mean by considering the entire data segment. The change point is the point at which some statistical property (in this case, the mean) changes most significantly. One popular means for change point detection is through minimization of a quadratic cost function as proposed by M. Lavielle [30]

$$J_{cp} = \sum_{k=1}^K \sum_{n=\tau_{k-1}+1}^{\tau_k} (y_n - \bar{y}_k)^2, \quad (5.1)$$



(a) Range profile when antenna elements are coherently summed.

(b) Range profile when antenna elements are non-coherently summed.

Figure 5.1: A comparison between the range profile of the sum of antenna elements for both a coherent, and non-coherent summation.

where K is the number of sections in the series, and τ_k represents the change point of the k^{th} segment. \bar{y}_k is the empirical estimate of the mean of the k^{th} segment. The value of the cost function represents the total squared error of each detected region from its estimated mean. Given a prior number of sections K , the goal is to find the values of τ_k that minimize this cost function. For the estimation of height, there is only one significant change point that needs to be estimated. The problem then simplifies to minimizing,

$$J_{cp}(\tau, y) = \sum_{n=1}^{\tau_k-1} (y_n - \bar{y}_k)^2 + \sum_{n=\tau_k}^N (y_n - \bar{y}_{k+1})^2. \quad (5.2)$$

Minimizing this equation finds the point τ_k that minimizes the sum of the errors between the means of the estimated segments. Substituting y_n with the range compressed signal, τ_k is now the range index of the height above ground. An example of a detected change point for a single pulse can be seen in Figure 5.2. For the reasons specified in the previous chapter, the change point detector is applied to log-scale data. This makes it easier to detect the order of magnitude change in signal strength between the background noise and the closest target.

It is not immediately apparent from the range profile of the simulated data-set (Figure 5.1) that the change point detector achieves better performance. To show the effectiveness of a cost function approach compared to a gradient approach, an experimental data segment is used. While the gradient is quite noisy and produces noisy estimates as a result, the change point detector is more stable due to incorporating information from the entire data segment in its cost function. This can be seen in Figure 5.3, where the result from the two methods is compared.

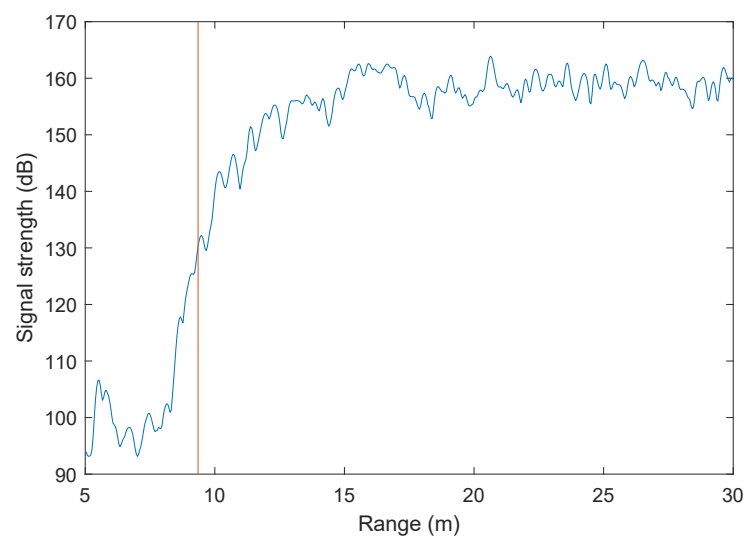


Figure 5.2: This figure shows the estimated change point using an experimental data set.

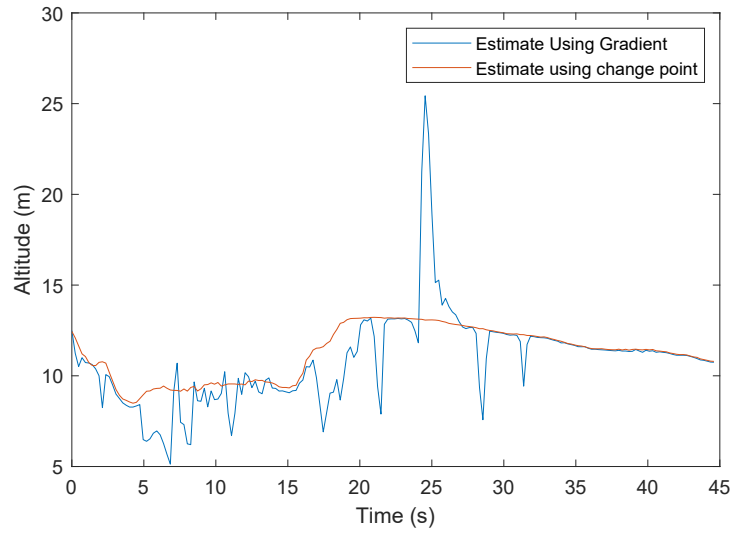


Figure 5.3: The estimated altitude using the change point detector compared to the estimated altitude using the gradient.

To get an indication of the performance of the altitude estimate, the method was tested on 35 s of simulated data. For the simulations, all available channels are used. The simulations are done using various maneuvers both with pitch and without pitch. The estimated standard deviation was found to be **0.018 m**, and the bias was found to be **0.32 m**. This means that the estimator has a significant bias. This bias is likely caused by the slope in the leading edge of the range profile. Unfortunately, this slope depends on the roughness of the terrain, which is an unknown parameter. A similar problem occurs in satellite altimetry due to varying penetration depth over surfaces such as land, water, and snow. The technique used to solve this problem is termed retracking. The most common method is the threshold first maximum retracking algorithm (TFMRA)[47]. This algorithm sets the "tracking-point" to a percentage of the detected peak echo power. The percentage is empirically set based on the terrain, with common values being 40%, 50%, and 80%.

For the drone-radar, the slope is not only caused by surface penetration, but also because the radar system has a wide beamwidth and relatively low gain in the downward direction. Furthermore, the power from the ground echo in experimental data will be higher due to the perpendicular grazing angle. Since the simulated data is not an accurate representation of the terrain in experimental scenarios, this bias is left as-is. A tunable parameter can be defined to correct the actual range to a percentage of the initial detected range, similarly to TFMRA.

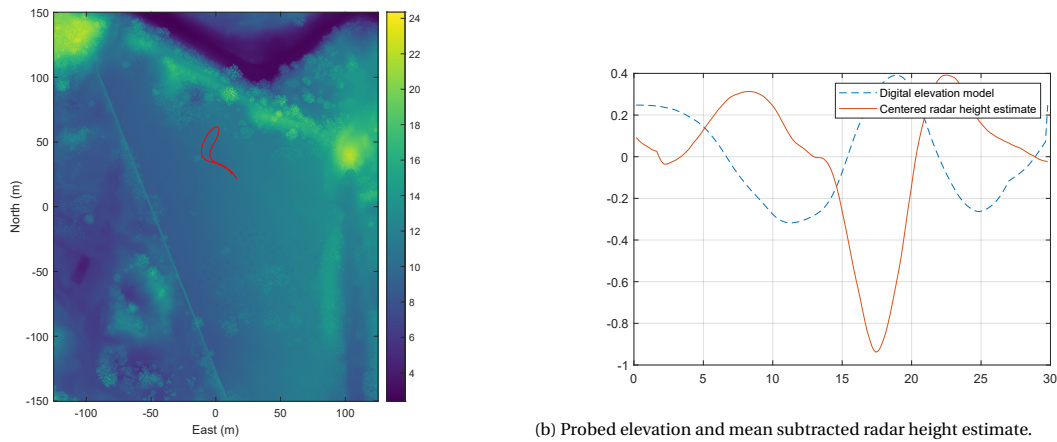
5.1.2. Considerations For Experimental Data

Two additional factors need to be taken into consideration when dealing with experimental data. First, while the estimate may be accurate, it should be noted that the obtained estimate is the distance between the radar and the closest reflecting object to the radar. This means that the estimated height does not always correspond to the reflection directly below the radar. As such, nearby tall objects will deteriorate the estimation accuracy. Moreover, a hilly terrain makes it impossible to estimate the radar's height with respect to a fixed reference point.

The first problem can be solved by performing the estimate in the range-Doppler domain. By performing the estimate around the zero-Doppler line, a true estimate of the distance directly below the radar is obtained. A possible problem with this approach is that fewer samples are used to create the estimate, possibly leading to more outliers. Also, when the altitude of the radar is increasing or decreasing, the true altitude is no longer around the zero Doppler line. As a result, the estimation method should be selected based on the scene and the trajectory.

Another way to increase the altimeter resolution is through delay-Doppler or SAR altimetry [42]. Similar to conventional SAR, the idea behind SAR altimetry is to create a finer spot resolution by synthesizing a larger antenna aperture. This is done by coherently adding ground returns using a rough initial estimate of velocity and height. We implemented SAR altimetry in the current study. However, the ground returns lacked enough coherence to provide additional benefits. In the end, the noncoherent summation of the range compressed signal was used.

The second problem can be solved by using a digital elevation model (DEM) of the scene. In The Netherlands, the freely available AHN3 DEM can be used which provides a 0.5 m resolution DEM [2]. It should be noted that this requires the GNSS location of the radar to be known with reasonable accuracy. However, it should be sufficient to eliminate slow variations in terrain elevation. To test this, the AHN3 dataset of the measurement area was downloaded and processed to match the coordinate system. An image of the processed DEM can be seen in Figure 5.4. The elevation is referenced to the selected origin of the ENU reference frame. At each point in the radar trajectory, the DEM is probed to obtain the elevation at this coordinate. The DEM profile is adequately smoothed and subtracted from the radar height estimate. The result can be seen in Figure 5.5. The compensated radar height is compared to the integrated GNSS velocity to get a qualitative indication of performance.



(a) Digital elevation model of the measurement area. The radar trajectory for this test is also shown in red.

Figure 5.4: An image of the DEM is shown on the first figure along with the radar trajectory from one of the data segments. The second figure shows the probed elevation along the radar trajectory along with the mean subtracted height estimate.

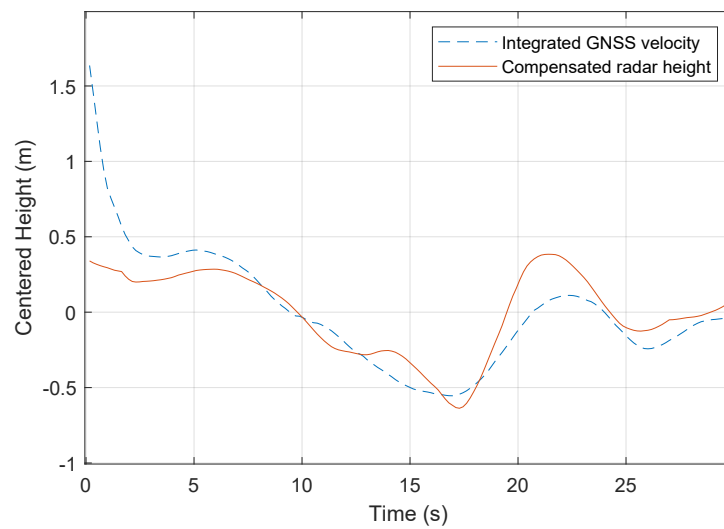


Figure 5.5: The compensated radar height estimate alongside the integrated upwards GNSS velocity. In this figure, the mean from both estimates is subtracted to aid visual comparison.

5.1.3. Overview

The processing steps of the proposed altimeter can be summarized as follows:

- The data from each channel is pre-processed, and range compressed. In addition, a range segment of interest with a minimum allowable range is selected to remove the direct cross-feed between receivers and transmitters.
- The pre-processed data from all channels is non-coherently summed.
- Using the change point detector described by (5.2), the point at which the signal strength changes most significantly is found. This point represents the altitude estimate.
- If needed, the obtained altitude estimates are compensated with a DEM referenced by the GNSS.

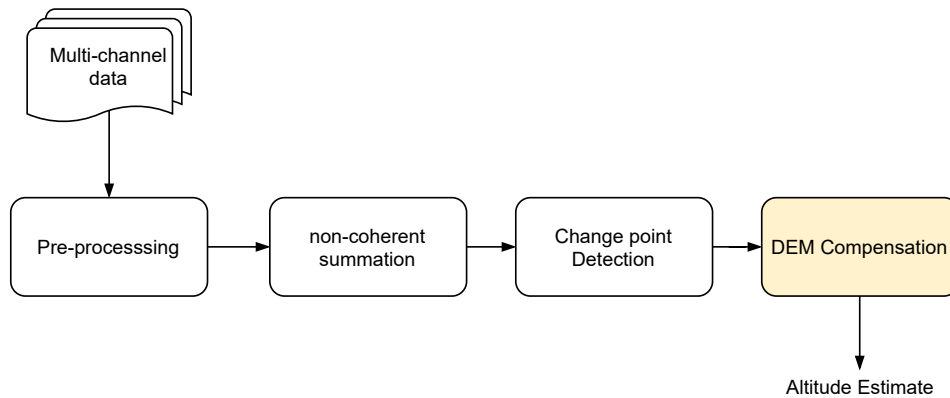


Figure 5.6: Flowchart of the altimeter processing steps. The highlighted DEM compensation step is not strictly required and depends on the terrain.

5.2. Velocity Estimation

In the previous section, an altimeter processing chain is proposed to estimate the altitude of the radar. This section will look at aspects of the radar data which can be used to estimate its velocity. To estimate the horizontal velocity, two distinct features of the range Doppler map can be exploited. The Doppler centroid provides an estimate of the radial velocity in the steering direction. The absolute value of the velocity vector, on the other hand, is reflected by the distinct edge on both sides of the spectrum. The clutter edge corresponds to the returns which are aligned with the velocity vector, as these have the highest Doppler component. Both of these features are clearly visible in Figure 5.7, where the centroid and clutter edge are highlighted for a simulated beamformed range-Doppler image.

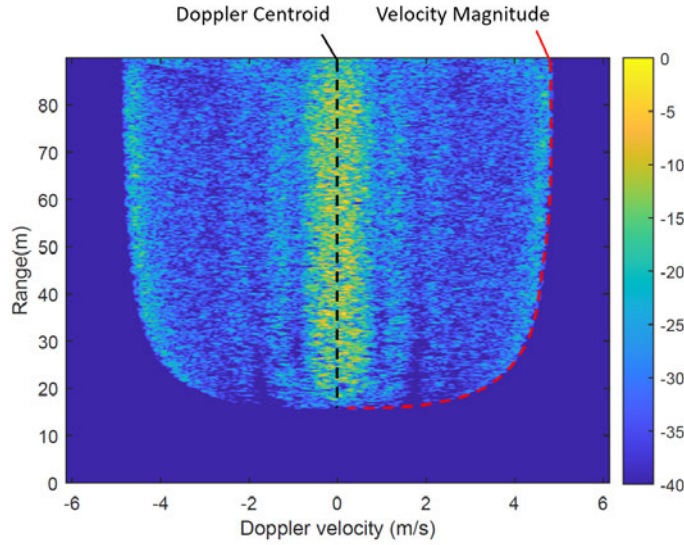


Figure 5.7: A simulated range Doppler image for a horizontal velocity of 5 m/s with the beam steered 90 degrees from the local x-axis. The clutter edge and Doppler centroid are also highlighted.

5.2.1. Velocity Magnitude

To estimate the absolute value of the horizontal velocity vector, we propose to use a parameterized edge detector. A similar approach is proposed by Mao et al. [35] for Doppler beam sharpening. Here, the clutter edge is detected using a threshold, followed by a series of morphological filtering operations. The velocity parameters are then extracted using a least-squares fit through the detected edge. While thresholding might work in cases where the main lobe is pointing forward, this would not work for a semi-side looking case as shown in Figure 5.7 due to the presence of the beam pattern. Instead, we can parameterize the edge before detection to simplify the problem. For a given Altitude h , and absolute horizontal velocity $|\mathbf{v}_h|$, the radial velocity at a range R is given by,

$$v_r(R) = |\mathbf{v}_h| \sqrt{1 - \left(\frac{h}{R}\right)^2}. \quad (5.3)$$

The values along this curve can be calculated for the range of possible velocities by interpolating the velocity axis for each range-bin, and summing along the range dimension. For each range bin, the new velocity axis \tilde{v} is given by

$$\tilde{v} = v \frac{1}{\sqrt{1 - \left(\frac{h}{R}\right)^2}}. \quad (5.4)$$

The interpolation is visually shown in Figure 5.8 using a range-Doppler image from an experimental data segment.

After summing, the problem is reduced to a 1-dimensional edge detection problem as shown in Figure 5.9. To increase the Signal to Noise Ratio (SNR) of the edge, the spectrum is folded over along the centerline to add both edges together. Since many range bins are averaged, the resulting edge is not very noisy and can be easily detected by its gradient. The estimator is applied to each digital beam separately. To gauge the quality




Figure 5.8: Transformation of the range-Doppler spectrum to a Doppler spectrum representing the horizontal velocity at each range bin. The spectrum is transformed by interpolating the Doppler axis at each range bin.

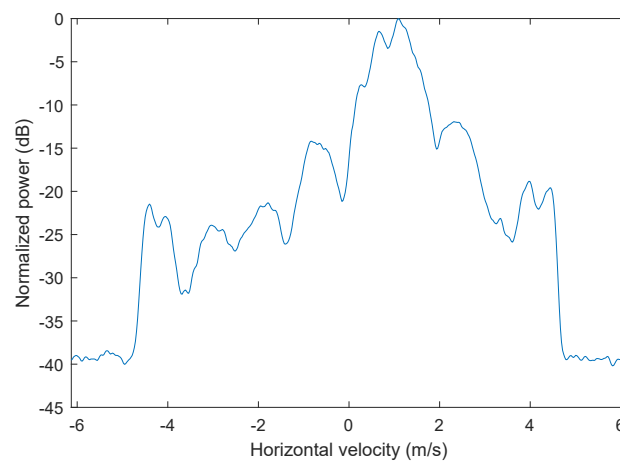


Figure 5.9: The 1D spectrum obtained after interpolating the Doppler axes in the original spectrum, and summing along the range dimension.

of each estimate, the peak-to-background ratio (PTBR) of the gradient magnitude is used as a metric. The PTBRs are normalized and used to weigh the contribution to the final estimate.

The estimator is first tested on 30 seconds of simulated data for a variety of maneuvers. The estimation was performed multiple times using different sets of parameters. The parameters that were changed were the number of antenna elements per beam and the step size between beam steering angles. The standard deviation, bias, and RMS errors can be seen in Table 5.1. The error does not vary significantly as a function of these parameters. However, in general, a larger number of antenna elements and a smaller step size yield a slightly better RMS error. This can be explained by the fact that there will be more estimates from the forward and rear looking regions with high SNR, decreasing the overall error.

The main difference between the simulated data and experimental data is caused by thermal noise and scene contrast. This can be observed when comparing the experimental range-Doppler spectrum with the simulated one (Figures 5.8 and 5.7). As contrast does not affect the edge, thermal noise is the most significant factor. The obtained velocity estimates are compared to the INS velocity over several experimental data segments for a qualitative indication of performance. Details of the experimental setup can be found in Appendix C. Two data segments consist of straight passes with a constant course and are 45 s in length each. In the third data segment, which is 30 s in length, the movement of the drone is more dynamic. The results are shown in Figures 5.10 to 5.12.

Table 5.1: The standard deviation, bias, and RMS error of the edge velocity estimator for simulated data. A number of different values are used for the steering step size and the number of antenna elements to test the dependence on different processing configurations.

stepsize	15			30			45		
#Rx	5	7	9	5	7	9	5	7	9
σ (m/s)	0.0168	0.021	0.018	0.021	0.024	0.021	0.020	0.030	0.028
$\bar{\epsilon}$ (m/s)	-0.015	-0.010	-0.012	-0.018	-0.010	-0.011	-0.012	-0.011	-0.014
RMS (m/s)	0.023	0.024	0.022	0.027	0.026	0.024	0.022	0.031	0.031

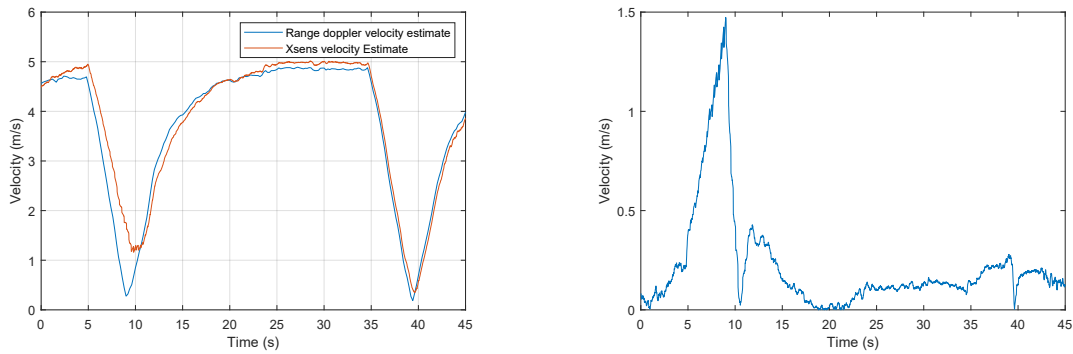


Figure 5.10: The estimated velocity magnitude and the INS velocity estimate. A data segment of 45 seconds is used, with 8 active side-looking receive elements. The RMS value of the error for this segment is 0.37 m/s.

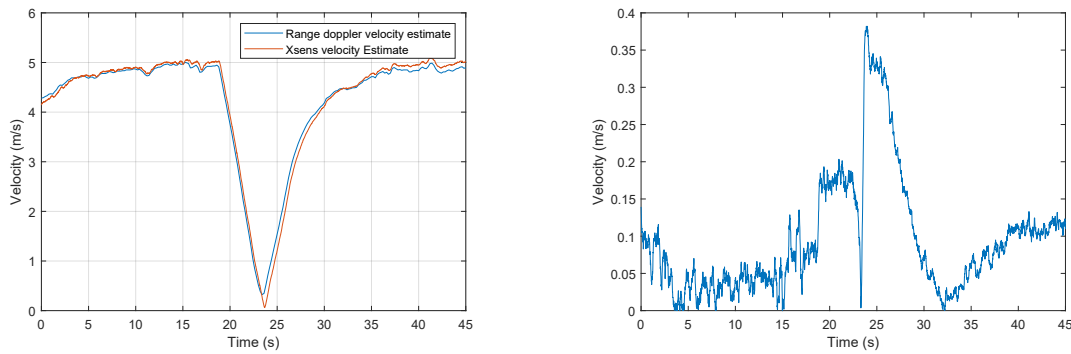


Figure 5.11: The estimated velocity magnitude and the INS velocity estimate. A data segment of 45 seconds is used, with 8 active side-looking receive elements. The RMS value of the error for this segment is 0.14 m/s.

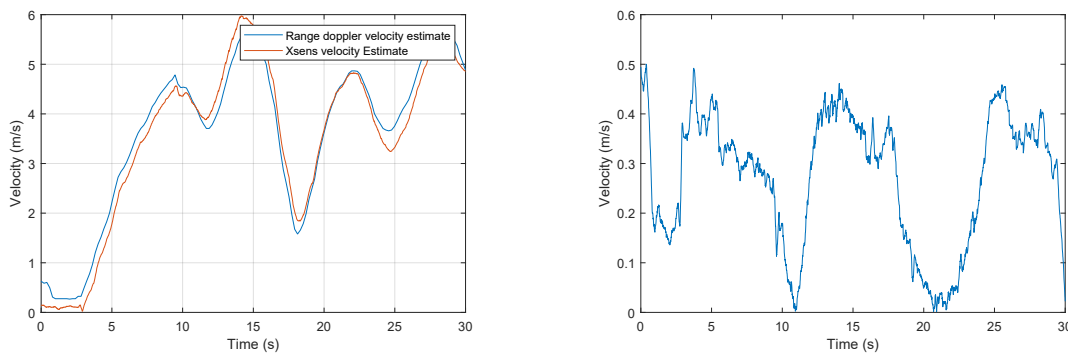


Figure 5.12: The estimated velocity magnitude and the INS velocity estimate. A data segment of 30 seconds is used with 8 active receive elements. For this data segment, an arbitrary flight path is flown. The RMS value of the error for this segment is 0.32 m/s.

For each estimate, a burst of 192 slow-time samples is used to compute the Doppler spectrum. The effective velocity resolution is given by

$$\frac{\text{PRF}}{N_{\text{FFT}}} \frac{c_0}{2f_c}. \quad (5.5)$$

For a PRF of 814 Hz, the velocity resolution is approximately 0.07 m/s. Qualitatively, the velocity estimates match the estimates produced by the INS quite well. Producing errors which are close to the specified performance of the INS. Something that can be noted from visual inspection, is that the radar velocity estimate is not able to estimate zero-velocity very well. This can be seen in Figure 5.12 at the start. It can also be observed in the other figures when the radar slows down and changes direction. This effect is caused by the zero Doppler components having a width larger than zero when the radar is not moving. This should be taken into account when fusing sensor data.

Overview

A flowchart of the estimator can be seen in Figure 5.13. The processing steps can be summarized as follows:

- For a given data segment, the range-compressed data is coherently added to create a number of digital beams.
- For each digital beam, the velocity is estimated by finding the velocity for which the gradient along a curve parameterized by the velocity is maximized
- Estimates from each beam are weighed by their normalized peak-to-background ratio to form the final estimate.

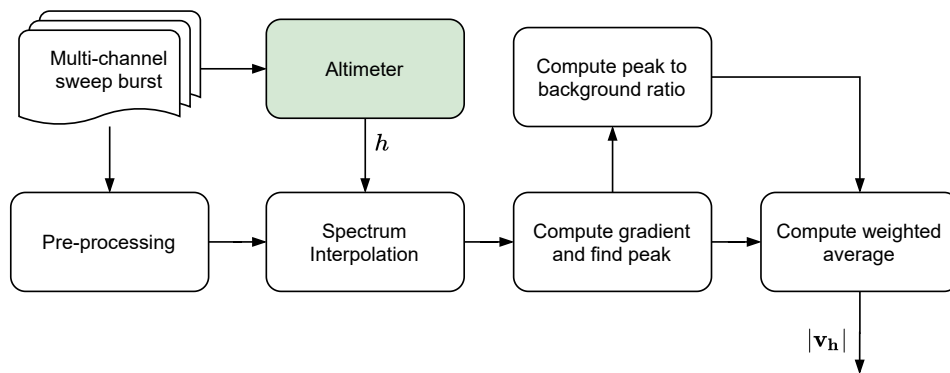


Figure 5.13: A flowchart of the steps used to compute the magnitude of the horizontal velocity vector.

5.2.2. Doppler Centroid

By estimating the edge of the Doppler spectrum, a constraint is obtained for the magnitude of the velocity vector. The angle of the velocity vector with respect to East (x-axis of the inertial reference frame) can be estimated from the Doppler centroid. The Doppler centroid at a certain range bin is an estimate of the radial velocity in that direction. There are several different Doppler centroid estimation algorithms available which are designed to deal with different scenarios. However, the majority of algorithms can be grouped into a framework of correlation-based estimators, which correlate the spectrum with a weighing function. The Doppler centroid is then found at the zero points (or peak, depending on the weighing function). The weighing function can incorporate prior information about the antenna pattern or noise distribution [4, 52].

The accuracy of the centroid estimate depends on several factors, which are mostly present in experimental data. The most notable being the distribution of targets in the illuminated scene. Strong, point-like targets cause unpredictable spikes in the spectrum and can make it difficult to detect its peak. Shadowed areas like the one present in Figure 5.8, also distort the spectrum further. Aside from the target distribution, there is also distortion in the antenna pattern when the squint angle is non-zero. This is because there is a non-linear relationship between the Doppler centroid velocity and the squint angle. The relationship is given by

$$|\theta_{DC}| = \arccos \frac{v_{DC}}{v_h}, \quad (5.6)$$

where θ_{DC} is the angle between the velocity vector and the steering direction, v_{DC} is the Doppler centroid velocity, and v_h is the magnitude of the horizontal velocity vector. The result of this is that the spectrum is compressed on one side at high squint angles. This was already introduced in the previous chapter. Typical centroid estimators in literature do not address this problem, as the relationship is approximately linear in the side looking region used for SAR imaging. However, due to its dynamic motion, the drone platform has rapidly changing Doppler centroid values spanning the entire bandwidth.




Figure 5.14: An illustration of the distortion in the Doppler spectrum caused by velocity projection from high squint angles. When the steering direction is close to the velocity vector, the spectrum overlaps, which is illustrated by the dotted line.

To solve the first problem and reduce scene contrast, a 2D cell averaging constant false alarm rate (CA-CFAR) detector is used to remove targets with very high Signal to Clutter (SCR) values ($> 10\text{dB}$). The CFAR window is narrow in azimuth to estimate the clutter level of a particular part of the antenna pattern. The range-Doppler image for an experimental data segment after CFAR thresholding is shown in Figure 5.15. While there is still contrast caused by scene features such as the shadowing seen on the left of the image, the overall "spikiness" of the spectrum is reduced. This is also visible in the main lobe of the 1D spectrum in Figure 5.16. The CFAR thresholding step is applied to the interpolated range-Doppler spectrum as defined for the edge detector.




Figure 5.15: The interpolated range Doppler image before (left) and after (right) removing strong targets with a CFAR detector. Notable removed targets are the two corner reflectors in the nulls of the antenna pattern, and the strong reflections in the main lobe around 80 m.




Figure 5.16: A comparison between the computed 1D spectrum without removing strong targets (left), and after removing strong targets using a CFAR detector (right).

It has been shown that the centroid estimator achieving the lowest variance in practice, is a correlation-based estimator that uses the nominal antenna pattern as its weighting function [52]. The expected antenna pattern can be calculated using the generalized array factor as defined in (2.9). As the spectrum axis is currently represented in units of velocity, the axis is first interpolated according to (5.6) to transform each velocity to an angle with respect to the velocity vector. This projects the spectrum back to a polar representation, minimizing the distortion of the spectrum at higher squint angles. The spectrum before and after interpolating, for an experimental data segment can be seen in Figure 5.17. After interpolating, the spectrum is correlated with the nominal antenna pattern computed using the array factor and the antenna configuration (Figure 5.18). The peak of the correlation result is then selected as the centroid estimate.

The estimated Doppler centroid can now be interpreted as a centroid angle that quantifies the angle between the steering direction and the velocity vector. Figure 5.19 shows the estimated centroid for simulated linear motion. For this simulation, a steering step size of 15 degrees is chosen. This means that the centroid spacing should remain constant as a function of steering direction. It can be seen that the spacing between subsequent estimates is mostly constant except for very high squint estimates (< 20 deg and > 160 deg). Since these estimates will always exhibit bias, they can be disregarded if enough estimates within a reliable range exist.

The centroid angle from a single steering direction provides only an estimate of the absolute angle between the steering direction and the velocity vector. The relationship can be described by

$$|\theta_{st,n} - \theta_{c,n}^x| = |\theta_{DC,n}|, \quad (5.7)$$

where $\theta_{st,n}$ represents the steering direction of the n^{th} estimate, $\theta_{c,n}^x$ represents the angle of the velocity vector with respect to the local x-axis, and $\theta_{DC,n}$ represents the n^{th} centroid estimate. This is also visually illustrated in Figure 5.20.




Figure 5.17: A comparison between the 1D spectrum after summing along the range dimension (left), and the 1D spectrum after interpolating the axis to an equivalent angle from the velocity vector (right).




Figure 5.18: The measured spectrum along with the calculated nominal antenna pattern. The antenna pattern is shifted on top of the measured spectrum in this image to visually show their correlation.




Figure 5.19: The figure on the left shows the estimated Doppler centroid for each steering direction as a function of the sample number, where each sample represents the centroid angle obtained for a given sweep burst. The figure on the right shows the estimated centroid for each steering direction averaged over the time samples. The centroid estimates are obtained using a simulated data-segment containing linear motion.




Figure 5.20: Illustration of the angles measured by the centroid estimator. The local reference frame is offset from the inertial reference frame in this illustration to stress that the radar has a non-zero heading, α in the inertial reference frame.

To obtain the sign of the estimate, at least two independent centroid estimates are required. This problem can be reformulated as a least-squares problem in the following way. Each steering direction can be represented by a unit vector h_n . Since the estimated centroid angle is the absolute value between the steering vector and the velocity vector in the local reference frame, the cosine of the centroid angle is equal to the scalar projection of the unit velocity vector onto the unit steering vector

$$\hat{\mathbf{h}}_n^\top \cdot \hat{\mathbf{v}}_h = \cos(\theta_{DC,n}). \quad (5.8)$$

This can be generalized to any number of centroid estimates for a single data section. By defining a matrix \mathbf{H} of steering directions, and a vector \mathbf{y} of centroid cosines:

$$\begin{bmatrix} \cos(\theta_{DC,1}) \\ \vdots \\ \cos(\theta_{DC,N}) \end{bmatrix} = \begin{bmatrix} h_{n,\hat{x}} & h_{n,\hat{y}} \\ \vdots & \vdots \\ h_{n,\hat{x}} & h_{n,\hat{y}} \end{bmatrix} \cdot \begin{bmatrix} v_{\hat{x}} \\ v_{\hat{y}} \end{bmatrix} \Leftrightarrow \mathbf{y} = \mathbf{H} \cdot \hat{\mathbf{v}}. \quad (5.9)$$

if the noise in the centroid estimates is assumed to be Gaussian, a linear estimation problem can be formulated

$$\mathbf{y} = \mathbf{H} \cdot \hat{\mathbf{v}}_h + \mathbf{w}, \quad (5.10)$$

where \mathbf{w} is an additive Gaussian noise vector distributed according to $\mathcal{N}(0, \mathbf{W})$. Here, \mathbf{W} is the covariance matrix of the system containing the variances of each measurement. The least-squares estimator for $\hat{\mathbf{v}}$ is then

$$\hat{\mathbf{v}}_h = (\mathbf{H}^\top \mathbf{W}^{-1} \mathbf{H})^{-1} \mathbf{H}^\top \mathbf{W}^{-1} \mathbf{y}. \quad (5.11)$$

In this equation, the covariance matrix weighs each centroid estimate. In practice, the variance of each radial centroid estimate should be determined according to some quality criterion. The covariance matrix of the final velocity estimate is given by

$$\mathbf{C} = (\mathbf{H}^\top \mathbf{W}^{-1} \mathbf{H})^{-1},$$

where each diagonal element represents the variance of a spatial component. From the above equation, it can be seen that if more independent, low-variance centroid estimates are incorporated, the variance of the

overall estimate can be reduced. The course with respect to the local reference frame is then

$$\theta_c^x = \arg(v_{\hat{y}} + j \cdot v_{\hat{x}}), \quad (5.12)$$

To compute the course with respect to the inertial (East-North) coordinate system, the estimated centroid must be referenced to an inertial heading. The course of the radar, which is defined as the angle between the horizontal projection of the velocity vector and the East direction, can be calculated using

$$\theta_c = \alpha + \theta_c^x, \quad (5.13)$$

where α is the heading, which is defined as the angle between the x-axis of the INS and the East direction.

For a simple case, the matrix \mathbf{W} in (5.11) can be assumed to be equal to the identity matrix. However, it has been shown that the variance of a centroid estimate is directly proportional to scene contrast, which is defined by

$$SC = \frac{\langle I^2 \rangle}{\langle I \rangle^2}, \quad (5.14)$$

where I is the pixel intensity and $\langle \cdot \rangle$ represents an averaging operator. In practical scenarios, the difference in contrast between different sides of the radar might be substantial. In [4], R. Bamler shows that the variance of a correlation-based Doppler centroid estimator can be approximated by

$$\text{var}(f_{DC}) = \frac{PRF \cdot SC}{N} \cdot \frac{\sum_{-B/2}^{B/2} (A(f) \cdot B(f))^2}{\left(\sum_{-B/2}^{B/2} A(f) \cdot B'(f)\right)^2}, \quad (5.15)$$

where N represents the total number of samples contributing to the estimate of f_{DC} , and B is the processing bandwidth. $A(f)$ and $B(f)$ represent the nominal antenna pattern and correlating function respectively.

The value of N is mainly determined by the number of slow-time samples and range bins used for estimation. In addition, when there is contrast in the scene, the number of effective samples is divided by the contrast factor SC [4, 34, 52]. The expression in (5.15) can also be written in terms of the centroid angle θ_{DC} , as

$$\text{var}(\theta_{DC}) = \frac{PRF \cdot SC}{N_B} \cdot \frac{\sum_{-B/2}^{B/2} (A(\theta) \cdot B(\theta))^2}{\left(\sum_{-B/2}^{B/2} A(\theta) \cdot B'(\theta)\right)^2}, \quad (5.16)$$

where N_B represents the number of samples corresponding to the range from $-\pi/2$ to $\pi/2$, which is determined by the magnitude of the velocity vector. Without going into the details of the antenna pattern and correlation function, two important observations can be made. First, the variance for a single estimate is directly proportional to the scene contrast. And second, the variance is also approximately proportional to the magnitude of the velocity vector. The contrast can be used to weigh individual course estimates, while the second observation can be used when fusing the radar estimates with INS estimates.

Accounting For Non-zero Attitude

There are two main ways to combat offset caused by non-zero attitude. The easiest way is by simply using beams pointing in opposite directions. This is also called a Janus configuration and is commonly found in acoustic Doppler logs on navy ships. The idea behind a Janus configuration is that the offsets caused by attitude on opposite sides of the sensor will cancel each other out. This is because while the horizontal velocities have opposite signs, the offset caused by attitude is the same on both sides. In the same way, the vertical velocity component is also suppressed.

If only one side of the radar is available, the offset caused by the attitude can still be compensated. However, this must be done using the attitude estimates obtained from the INS. (4.5) in the previous chapter, describes the velocity offset caused by a particular attitude. To simplify the analysis, the average offset over a specific window can be observed. Recall that the spectrum is integrated along the range lines for a zero attitude assumption. As the offset is proportional to $1/R$, the average offset for a range from R_1 to R_2 is given by the integral of the offset

$$\langle v_s(R_1, R_2) \rangle = \frac{1}{R_2 - R_1} \int_{R_1}^{R_2} v_s(R) dR = v_p h \tan \beta_{st}(\theta_{st}, \beta, \gamma) \frac{\log \frac{R_2}{R_1}}{R_2 - R_1}. \quad (5.17)$$

The average centroid offset is then

$$\langle \theta_{c,o} \rangle = \arcsin \frac{\langle v_s(R_1, R_2) \rangle}{v_h}. \quad (5.18)$$




Figure 5.21: Illustration of a Janus configuration using 4 beams on the drone radar.

The velocity v_p used to calculate the offset, depends on the unbiased course itself. However, if the biased course is used as an initial estimate, the relative error made in correcting the offset is small. To show this, consider the case where the radar is moving linearly at 5/ms, has a positive pitch of 15 degrees, and the beam is steered to 90 degrees in the local reference frame.

The centroid offset as a function of course angle can be seen in Figure 5.22. The figure shows the average offset over a range from 30 to 80 m at a height of 15 m with a pitch of 15 degrees. It can be seen that while an initial estimate of the course is necessary to determine the centroid offset, the error made when using the initial biased estimate is small. For larger course angles, the difference in offset between subsequent angles becomes larger. However, the process could be iterated several times to converge at the correct estimate.




Figure 5.22: The average offset in course for a pitch of 15 degrees in a range from 30 to 80 m as a function of the course angle. The left data-tip highlights the measured course offset when the actual course is 90 degrees. The right data-tip highlights the error made when the biased course is used to calculate the offset.

Evaluation

To get an indication of the performance of the centroid estimator, the centroid estimator is first tested on the set of simulated data previously also used for the edge velocity estimator. The obtained results can be seen in Table 5.2. This shows that the accuracy of the estimate depends mostly on the number of beams that are used. In addition, beams using more antenna elements also reduce the error. This is an expected result as the overlapping effect shown in Figure 5.14 is reduced for narrow beamwidths.

Table 5.2: The average standard deviation, bias, and RMS error of the centroid estimator using the set of simulated data. The error statistics are calculated for a number of antenna configurations and steering step sizes.

stepsize	15			30			45		
#Rx	5	7	9	5	7	9	5	7	9
σ (deg)	0.38	0.30	0.27	0.56	0.38	0.32	0.92	0.71	0.75
$\bar{\epsilon}$ (deg)	0.05	0.05	0.02	0.04	0.09	0.04	0.13	0.14	0.17
RMS (deg)	0.38	0.29	0.27	0.56	0.39	0.32	0.93	0.72	0.77




Figure 5.23: The Course with respect to East (left) and offset from the INS (right) corresponding to the velocity graph in Figure 5.10.




Figure 5.24: The Course with respect to East (left) and offset from the INS (right) corresponding to the velocity graph in Figure 5.11.




Figure 5.25: The Course with respect to East (left) and offset from the INS (right) corresponding to the velocity graph in Figure 5.12.

After testing on simulated data, the obtained course estimate can be compared to the INS course estimate using the same data segments as with the velocity magnitude. The results can be seen in Figures 5.23 to 5.25. Visually, The obtained course estimates are relatively close to the estimates obtained from the INS, with the largest differences being seen at direction changes and low velocities (around 25 s in Figure 5.24 and from 0 to 3 s in Figure 5.25). However the estimation accuracy is still relatively stable considering there are much fewer samples used to estimate the course.

5.2.3. Overview Of Velocity Estimation

Now that all parts of the velocity estimator have been presented, the velocity estimation procedure can be summarized as follows:

1. For the current data segment, the height is estimated using the radar altimeter.
2. A number of digital beams are defined and the appropriate channels are coherently added together.
3. For each digital beam, the velocity axis of each range bin is interpolated to represent the velocity in the horizontal plane. The spectrum is summed in the range direction to obtain a 1D spectrum.
4. The magnitude of the horizontal velocity vector is determined by finding the edge of this spectrum. Individual estimates from different beams can be weighed using the peak-to-background ratio of the edge gradient.
5. Using the integrated spectrum, and previously estimated velocity magnitude, the spectrum is interpolated to a polar form.
6. The interpolated spectrum is correlated with the array factor of the constructed receiver array to estimate the centroid angle.
7. Using multiple centroid angles, and the heading in the inertial reference frame, the course in the horizontal plane can be determined.
8. If only one-sided data is available, the offset caused by a non-zero attitude can be compensated using the output of the INS.

A flowchart of the velocity estimator is shown in Figure 5.26. The final estimate consists of a magnitude $|\mathbf{v}_h|$ which specifies the horizontal velocity, and a course θ_c which specifies the angle of the horizontal velocity with respect to the East direction. Depending on the requirements, this polar form may also be transformed into a Cartesian representation.

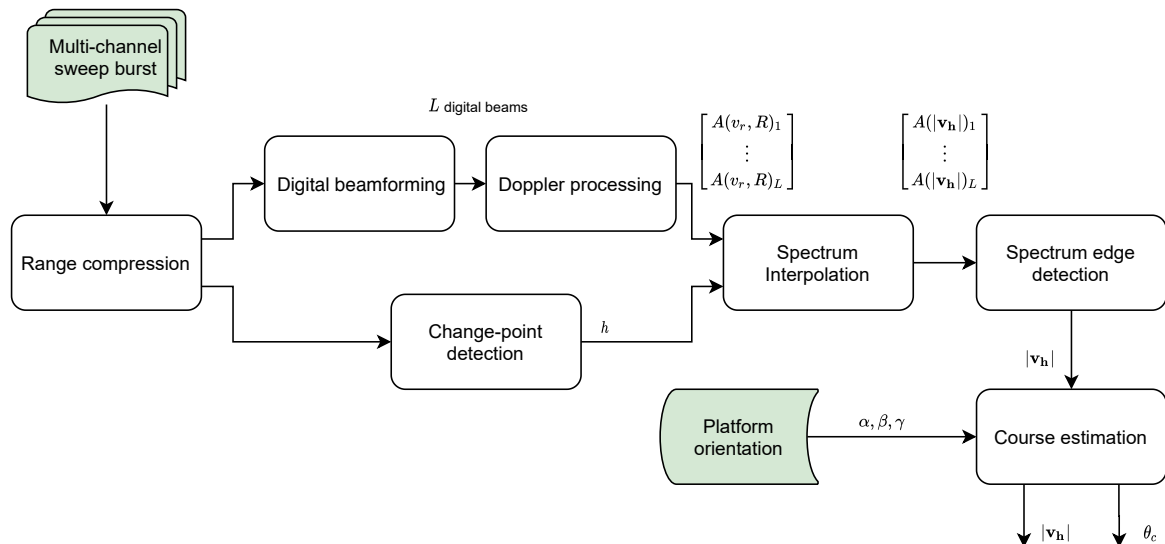


Figure 5.26: A flowchart of the entire velocity estimation procedure. A vector notation is used to indicate outputs from multiple steering directions. Here $A(v_r, R)$ corresponds to the range-Doppler spectrum, and $A(|\mathbf{v}_h|)$ corresponds to the interpolated 1D spectrum.

5.3. Multi-Beam Autofocus

In the previous section, an estimator was proposed for estimating the horizontal velocity vector. The estimator exploits the clutter distribution of the Doppler spectrum to obtain an estimate of course and velocity. The performance of the velocity estimate is optimal when scene contrast is minimal and deteriorates when contrast is present. Scene contrast, however, also provides useful information about platform motion. To exploit scene contrast, a different estimation method needs to be used. Contrast is the basis of many autofocus techniques in SAR imaging. Autofocus is a very powerful technique, in the sense that it provides wavelength level error correction provided that it is correctly implemented. In this section, we propose an extension to traditional autofocus techniques. The technique is specifically designed for an omnidirectional antenna array and is termed multi-beam autofocus. This section will describe the proposed implementation of autofocus for ego-motion estimation of the drone radar. First, the general autofocus problem will be presented. After which an equivalent convex problem will be derived which can be efficiently solved as a Semidefinite Program (SDP) using a row-by-row method. Various techniques presented in recent literature will be adapted and combined to develop a new relative positioning algorithm for an omnidirectional radar array.

5.3.1. The Autofocus Problem

Consider the autofocus problem as formulated in Section 3.2. A group of de-focused pixels can be expressed as an $N \times M$ matrix

$$A = \begin{bmatrix} R_{11} \exp j\phi_{\delta 1} & R_{12} \exp j\phi_{\delta 1} & \dots & R_{1M} \exp j\phi_{\delta 1} \\ R_{21} \exp j\phi_{\delta 2} & R_{22} \exp j\phi_{\delta 2} & \dots & R_{2M} \exp j\phi_{\delta 2} \\ \vdots & \vdots & \ddots & \vdots \\ R_{N1} \exp j\phi_{\delta N} & R_{N2} \exp j\phi_{\delta N} & \dots & R_{NM} \exp j\phi_{\delta N} \end{bmatrix}, \quad (5.19)$$

where each row of A represents a "pulse" contribution to the pixel tile. Note that it is assumed that the pixels are already projected to the imaging grid. This is also illustrated in Figure 5.27.




Figure 5.27: This figure shows the construction of the A matrix in (5.19) using radar geometry.

The group of pixels is assumed to be chosen such that the phase error $\phi_{\delta n}$ for a given pulse is the same across all pixels. Moreover, the phase error is assumed to be smaller than a single resolution cell such that R_{NM} does not change. To correct the phase error present in each pulse, the matrix A should be multiplied by a vector \mathbf{x} with entries equal to the conjugate of the estimated phase error in each pulse. Where \mathbf{x} is given by

$$\mathbf{x} = \begin{bmatrix} \exp j\hat{\phi}_{\delta 1} \\ \exp j\hat{\phi}_{\delta 2} \\ \vdots \\ \exp j\hat{\phi}_{\delta N} \end{bmatrix}. \quad (5.20)$$

The phase corrected pixels are then given by

$$\mathbf{A}_{\text{corr}} = \mathbf{x}^H \mathbf{A}. \quad (5.21)$$

Here, \mathbf{A}_{corr} is a vector of pixels which can be reshaped into the original pixel tile. In order to obtain an estimate of \mathbf{x}^H , the image sharpness can be used as a metric. This metric has shown to be well suited to SAR imagery [14]. The sharpness function is defined as the sum of squared pixel values and can be written as

$$G(\mathbf{x}) = (\mathbf{x}^H \mathbf{A})(\mathbf{x}^H \mathbf{A})^H = \mathbf{x}^H \mathbf{A} \mathbf{A}^H \mathbf{x}. \quad (5.22)$$

By defining

$$\mathbf{C} = \mathbf{A} \mathbf{A}^H,$$

the following constant modulus quadratic program (CMQP) is obtained:

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && -\mathbf{x}^H \mathbf{C} \mathbf{x} \\ & \text{subject to} && |\mathbf{x}_i| = 1, i = (1, \dots, N). \end{aligned} \quad (5.23)$$

While the objective function is convex, the set over which \mathbf{x} is defined, is not. CMQP's belong to the set of Non-Polynomial time (NP) hard problems. M. Duersch [14] solves the problem using a coordinate descent approach. Each entry of \mathbf{x} is optimized separately using a grid search. The problem can be solved more efficiently by solving a separate convex problem for which the same solution is obtained. A popular relaxation method for CMQP's is semi-definite relaxation (SDR) [33]. The idea of applying SDR to autofocus was first proposed by Lui et al [31]. Here, they apply it to phase gradient autofocus (PGA) and multichannel autofocus (MCA). It was also shown to be a tighter relaxation in terms of objective value than other methods. The Application of SDR to bistatic backprojection was shown by Evers et al. [17]. Here they show its application in an experimental on-rail setup using metal plates as focusing points. The semi-definite program (SDP) is obtained as follows. By noting that

$$\mathbf{x}^H \mathbf{C} \mathbf{x} = \text{Trace}(\mathbf{x}^H \mathbf{C} \mathbf{x}) = \text{Trace}(\mathbf{C} \mathbf{x} \mathbf{x}^H),$$

and defining

$$\mathbf{X} = \mathbf{x} \mathbf{x}^H,$$

the following equivalent SDP is obtained

$$\begin{aligned} & \underset{\mathbf{X}}{\text{minimize}} && -\text{Tr}(\mathbf{C} \mathbf{X}) \\ & \text{subject to} && \text{diag}(\mathbf{X}) = (1, \dots, 1), \\ & && \mathbf{X} \geq 0, \\ & && \text{rank}(\mathbf{X}) = 1, \end{aligned} \quad (5.24)$$

In this problem, the remaining non-convex constraint is the rank-1 constraint on matrix \mathbf{X} . The problem can be relaxed by eliminating this constraint, yielding the following convex optimization problem

$$\begin{aligned} & \underset{\mathbf{X}}{\text{minimize}} && -\text{Tr}(\mathbf{C} \mathbf{X}) \\ & \text{subject to} && \text{diag}(\mathbf{X}) = (1, \dots, 1), \\ & && \mathbf{X} \geq 0 \end{aligned} \quad (5.25)$$

SDPs are commonly solved using so-called interior point (IP) methods. However, these are fairly computationally complex for large scale problems. To ease the computational burden, we propose to use the row-by-row method as applied by Wai et al. [50] to multiple input multiple output (MIMO) detection, to efficiently solve the SDP. This is explained in Appendix A.2.

Once \mathbf{X} is obtained, the relative phase error with respect to the first pulse \mathbf{x} , is given by the first row or column of \mathbf{X} . The first off-diagonal of \mathbf{X} contains the phase difference from pulse to pulse.

5.3.2. Focusing A Single Target

Guo et al. show in [22], that autofocus can be integrated into an imaging algorithm, provided that the pulse-to-pulse phase error is smaller than $\frac{\lambda}{2}$. This is done by focusing each pulse as it is processed. This means that storing all pulses is not required. This principle can also be applied to the multi-beam autofocus algorithm.

Given an incomplete, but focussed backprojected image I_{n-1} , and the subsequent pulse S_n , the image after processing the n^{th} pulse is given by

$$I_n = I_{n-1} + S_n \exp(j\phi_{\delta n}), \quad (5.26)$$

where k is the wavenumber and $\phi_{\delta n}$ is the phase correction computed by the autofocus algorithm. The pulse contribution S_n is assumed to be interpolated to the correct location and phase-corrected according to the range of the radar.

To test the autofocus algorithm, a single simulated point target imaged using a trajectory error. For the trajectory error, a cross-range acceleration is introduced along a straight flight path. This is illustrated in Figure 5.28. The target before and after focusing can be seen in Figure 5.29.




Figure 5.28: The actual simulated radar trajectory shown in blue and the "measured" flight path shown in orange. The point target is positioned at the origin. The total trajectory comprises 1 second of data.

5.3.3. Ego-motion Estimation

By itself, autofocus provides only 1-D compensation along the radar line of sight. While this is adequate for far-range SAR imaging, there is too much range variation when dealing with near-range scenes. This means that only a single part of the image will be focused. Moreover, corrections are usually limited to within a single resolution cell since only the phase is corrected.

In [48], A. Sommer and J. Osterman address the problem of phase wrapping through the correction of flight trajectory while performing autofocus. They propose to propagate each phase correction to all subsequent pulses as an equivalent range error, allowing for a much larger compensation along the total flight path while only requiring that the maximum error from pulse to pulse is limited to $\frac{\lambda}{2}$. However, for near field imaging, the measured range error does not correspond well to the actual trajectory error. This is illustrated in Figure 5.30 where only a single point target is used for trajectory correction.

To account for the variability of the phase error over the entire image, Ran et al. [41] propose the use of multiple local images within a synthetic aperture to calculate a 3D phase adjustment. We propose to combine both the idea of 3D phase error correction and error propagation and generalize it to an omnidirectional array. If it is assumed that the phase error is only caused by motion errors, the two-way radial range error is given by

$$\Delta r = \frac{\lambda}{4\pi} \phi_{AF}$$




Figure 5.29: Autofocusing a single point target with a range acceleration error. The left image shows the unfocussed while the right image shows the focussed target.




Figure 5.30: Corrected trajectory using a single point target. An unmeasured acceleration is introduced in the range direction. Using autofocus and range error propagation, a corrected trajectory is calculated. Since only a single dimension is available, the trajectory is corrected in the wrong plane

The problem geometry is illustrated in Figure 5.31. Under the assumption that $n - 1$ pulses are focussed, the unit vector from the center of the i^{th} focussing image to the radar is given by \mathbf{p}_n^i . The measured radial range error can be expressed as a scalar projection of the 3D trajectory error on \mathbf{p}_n^i .

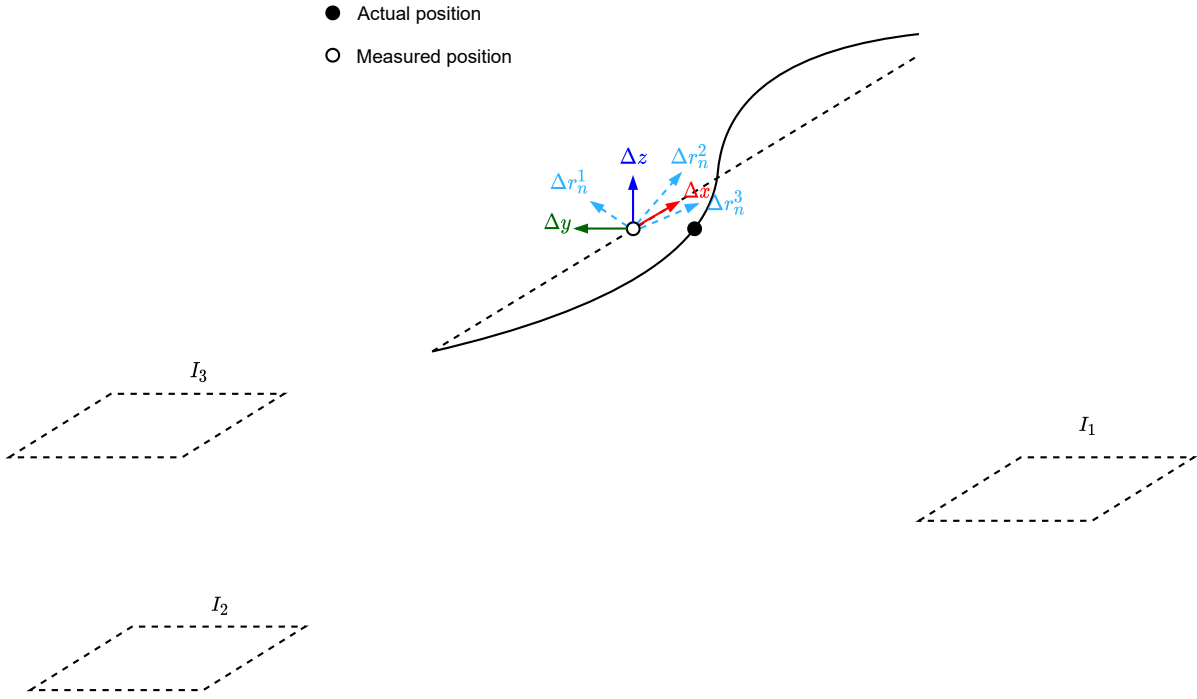


Figure 5.31: An illustration of the pulse by pulse range error for 3 focussing regions. The estimated relative range errors Δr_n^i are projections of the actual error $[\Delta x \ \Delta y \ \Delta z]$

$$\Delta r_n^i = \mathbf{p}_n^i \cdot \Delta \mathbf{x}_n \quad (5.27)$$

For I images, the problem can be written as

$$\Delta \mathbf{r}_n = \mathbf{P}_n \cdot \Delta \mathbf{x}_n \quad (5.28)$$

where $\Delta \mathbf{r}_n$ is the vector of radial errors and \mathbf{P} is the measurement matrix containing the unit vectors from each local image to the radar in its rows

$$\Delta \mathbf{r}_n = \begin{bmatrix} \Delta r_n^1 \\ \vdots \\ \Delta r_n^I \end{bmatrix}, \mathbf{P} = \begin{bmatrix} p_n^1(x) & p_n^1(y) & p_n^1(z) \\ \vdots & \vdots & \vdots \\ p_n^I(x) & p_n^I(y) & p_n^I(z) \end{bmatrix}.$$

If a simple additive noise model is assumed, the maximum likelihood estimator for $\Delta \mathbf{x}_n$ is given by the weighted least squares estimator

$$\Delta \hat{\mathbf{x}}_n = (\mathbf{P}^T \mathbf{W}^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{W}^{-1} \Delta \mathbf{r}_n, \quad (5.29)$$

where \mathbf{W} is the covariance matrix, containing the variances of each measurement in its diagonal.

The proposed positioning method combines both range error propagation and multiple local imaging, such that the entire trajectory is updated after each estimate of $\Delta \hat{\mathbf{x}}_n$. This makes it possible to obtain an estimate of relative position from pulse to pulse for arbitrary trajectories. However, a rough estimate of the actual trajectory should be available.

To exploit the capabilities of the omnidirectional array, a digital beam is steered to the appropriate direction as illustrated in Figure 5.32. Beam steering is performed implicitly through the imaging algorithm by selecting the appropriate antenna elements depending on the direction of the focusing region. We propose to use a total of 9 antenna elements per beam for increased resolution. Selecting the appropriate center antenna element ensures that the highest resolution is obtained by minimizing the amount of digital steering

that needs to be performed. For each pulse, the range compressed signal is then interpolated to the grid and the appropriate phase shift is applied, before adding the contribution from each antenna element. By doing this, the interference from nearby scattering regions is minimized. This process is repeated for each region and each pulse.




Figure 5.32: Illustration of beam steering to different focussing regions.

5.3.4. Proof Of Concept

To verify the feasibility of the proposed positioning method, a simulation of the radar system is used. The reflected signal of several point targets is simulated for an unknown trajectory. To simplify the simulation, the initial position with respect to the targets is assumed to be known and the location of the point targets is also assumed to be known. This is a reasonable assumption since the relative position of the targets can be determined from a single pulse for which the azimuth resolution is approximately 10 degrees. For relatively flat terrain, the altitude of the radar can be determined from the nadir return, while local regions to be used for focusing can be selected based on criteria such as reflectivity or contrast.

The setup used for the simulation is shown in Figure 5.33. Four Point targets are arbitrarily placed on a Cartesian grid. The radar trajectory is chosen to be a straight trajectory along the x direction with arbitrary perturbations (velocity error in x , oscillation in y and acceleration in z). The perturbations in each axis are given by

$$\begin{aligned}\epsilon_x &= \frac{n}{N} \\ \epsilon_y &= 0.3 \sin \frac{14n}{N} \\ \epsilon_z &= \frac{n^2}{N},\end{aligned}$$

where n is the pulse number and N is the total number of pulses. The velocity is chosen to be 6 m/s and the duration is 2 seconds. The unperturbed trajectory is considered to be the measured trajectory or initial guess. The trajectory is first estimated with a clean signal. The results are shown in Figure 5.34 and Figure 5.35




Figure 5.33: Geometry used for the simulation. The radar moves along the positive x direction with perturbations in each axis.

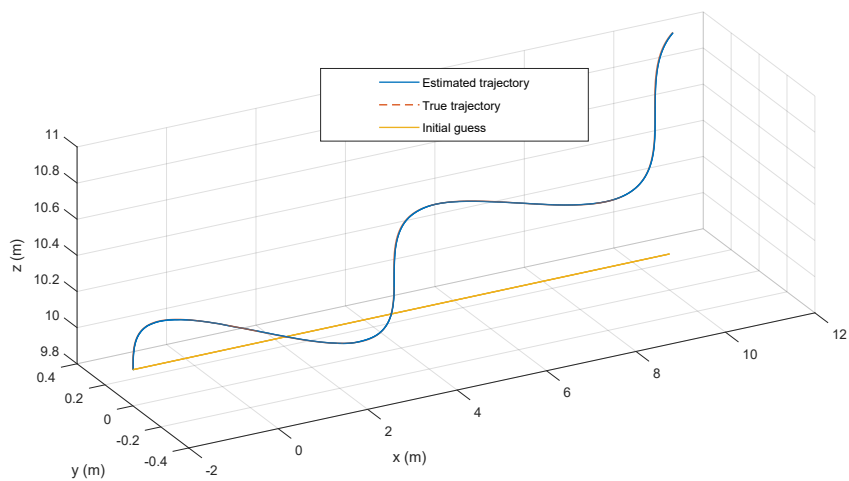


Figure 5.34: The estimated trajectory using the autofocus method with no added noise. The figure shows the initial guess, the estimated trajectory, as well as the true trajectory.




Figure 5.35: The estimated trajectory error in meters. The total length corresponds to a duration of 2 seconds.

shows the error between the estimated trajectory and the actual trajectory. To test the robustness in a more realistic scenario, additive noise is added to the range compressed signal. Down to an SNR value of -25 dB, the estimation performs reasonably well with only a slight increase in error. The estimation accuracy starts to degrade around -27 dB. This can be seen in Figures 5.36 and 5.37. The errors manifest as jumps on the order of a wavelength. These are most likely caused by phase wrapping.




Figure 5.36: The estimated trajectory with an SNR of -27 dB.

5.3.5. Conclusion

This section proposed a 3D positioning method based on autofocus. A flowchart of the proposed algorithm can be seen in Figure 5.38. On simulated radar data the proposed method performs very well, achieving sub-resolution accuracies for an SNR of -27 dB. It is important to note that the targets used for this simulation are point targets with a constant reflectivity across all illumination angles. In practical scenarios, this is not the case. The focussing regions could be adaptively selected in order to maintain a relatively constant reflectivity. The sidelobes of other areas might also pose a problem to autofocus performance. This is reduced




Figure 5.37: Estimation error for an SNR of -27 dB.

by beamforming in the appropriate direction. However, the actual performance needs to be validated with real data. The performance could be improved by taking the dynamics of the radar platform into account. By applying constraints to the maximum velocity and acceleration, noisy estimates can be filtered. It is also beneficial to estimate the covariance matrix, \mathbf{W} in (5.29). This will reduce the effect of noisy estimates, and provide information when incorporating the estimates with INS data.

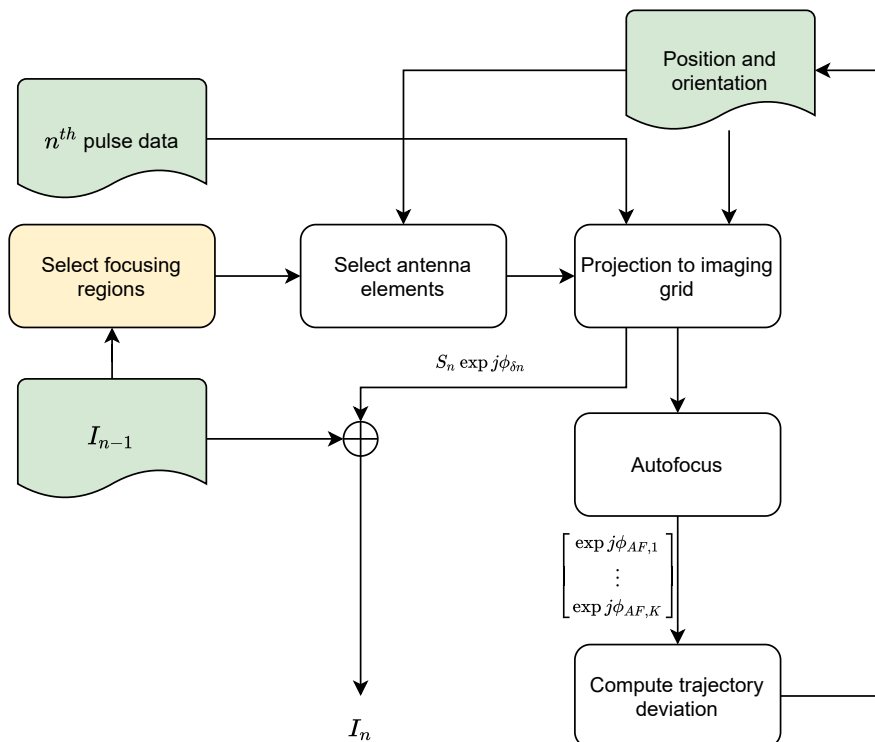


Figure 5.38: A top-level flowchart for the proposed multi-beam autofocus algorithm. The inputs are highlighted in green. The selection of focusing regions is highlighted in orange, as this is currently done manually but could be automated using a detection approach. Note that the projection to the imaging grid is done twice per pulse. Once to interpolate the range bins, and once to update the phase corrections based on the corrected trajectory.

5.4. Summary

In this chapter, we analyzed the characteristics of the radar signal to develop a number of motion estimation techniques for the drone-radar. The developed techniques built on the existing fields of radar altimetry, and SAR motion compensation, and introduced new, generalized algorithms for dealing with data from an omnidirectional radar array. The proposed velocity estimator is universal in the sense that it can be applied to a wide number of receiver configurations, both omnidirectional and fixed-look receive arrays, and is mostly invariant to the direction of the radar beam with respect to the velocity vector. In addition, the estimate does not depend on reference objects, allowing accurate estimates in many scenarios.

On the other hand, a method was also proposed to exploit scene contrast. The proposed multi-beam autofocus builds on the works of Ran et al. [41] and Sommer et al. [48] to provide a framework for simultaneous omnidirectional imaging and trajectory correction while taking advantage of the beamforming capabilities of the antenna array.

6

Sensor Fusion

In the previous chapter, several algorithms were developed to estimate kinematic states from radar data. We obtained estimates for the height and horizontal velocity and developed an extended autofocus algorithm for correcting relative motion from pulse to pulse by exploiting contrast in a number of local images. To summarize,

- The altitude was estimated using the ground reflection compensated by a DEM, providing an estimate of the height of the radar with respect to a fixed origin point in the scene.
- The horizontal velocity magnitude was estimated using the edge of the clutter spectrum.
- The course of the radar was estimated through the Doppler centroid from multiple steering directions. Combined with the velocity magnitude, this provides a velocity in the East - North plane.
- In ideal conditions (multiple high contrast areas), the presented multi-beam autofocus provides relative positioning from pulse to pulse in 3 dimensions.

To combine these estimates in a meaningful way, both with each other, and with the INS, some kind of sensor fusion needs to be applied. In this chapter, the available observations will be analyzed and used to decide on a sensor fusion framework for combining multiple observations. The first section will give a summary of the available observations from the radar and INS. Thereafter, based on practical considerations about the system's use cases, a sensor fusion algorithm will be selected that is suited for the task at hand.

6.1. Fusion Framework

While the terminology used in this report, and the presentation of the observations in the previous section already hint at a Kalman filter implementation, it is not immediately apparent that the Kalman filter is an appropriate fusion mechanism in this case. After all, several problems can present themselves when using a Kalman filter. The Kalman filter is a probabilistic filter that uses both measurements and predictions. The predictions are based on a kinematic model of the system being observed. This can be very advantageous if the system is modeled well, but also detrimental when this is not the case, as inaccurate predictions can deteriorate the estimates very rapidly.

An advantage of using a prediction based approach is that predictions are done based on information from the last estimated state. This means that a velocity observation will not cause the position estimate to drift over time if an absolute observation of position is also available, as would be the case if the velocity estimate were to be integrated from the initial time step before fusion.

Since, in this case, the system of interest is a moving object with on board sensors, its behavior is well modeled by a kinematic model. Thus, it makes sense to use a Kalman filter. Modeling the dynamics of multi-rotor aircraft is a well-studied problem, and there is a large amount of documentation and software available.

Ideally, the radar estimates should be incorporated into one of these existing frameworks. However, the design of an optimal system model and software implementation is an extensive research topic in itself. As this thesis focuses mostly on the signal processing aspects of the radar data, a simple model will be used to obtain the absolute position from velocity estimates. To observe the advantage of accurate velocity estimates,

this chapter will also include results from a tightly coupled Kalman filter which directly fuses the raw sensor readings with the radar estimates, showing the advantage of radar estimates in a more practical scenario.

6.1.1. System Model

For the system model, a constant-acceleration model is used. This means that the state in the next time step is predicted by linearly propagating the acceleration and velocity estimate. The state vector is given by

$$\mathbf{x} = [x \quad \dot{x} \quad \ddot{x} \quad y \quad \dot{y} \quad \ddot{y} \quad z \quad \dot{z} \quad \ddot{z}]^\top, \quad (6.1)$$

and the state transition matrix is given by

$$\mathbf{F} = \begin{bmatrix} 1 & d_T & \frac{d_T^2}{2} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & d_T & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & d_T & \frac{d_T^2}{2} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & d_T & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & d_T & \frac{d_T^2}{2} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & d_T & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 & d_T & \frac{d_T^2}{2} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & d_T \end{bmatrix}, \quad (6.2)$$

where d_T is the size of the time step. The prediction equation is then given by

$$\mathbf{x}_t = \mathbf{F}\mathbf{x}_{t-1} + \mathbf{w}_t, \quad (6.3)$$

where \mathbf{w}_t is a random vector which introduces noise to the state estimate, and is described according to the process noise covariance matrix \mathbf{Q} as presented in Chapter 3.

6.1.2. State Observations

An overview of the radar state observations can be seen in Table 6.1.

Table 6.1: Observations obtained from radar data.

State	State
Position (U)	Radar altimeter
Velocity (E, N)	Doppler centroid estimation, Multi-beam autofocus
Velocity (U)	Multi-beam autofocus

The measurement vector of radar estimates is given by

$$\mathbf{z}_R = [v_h \quad \theta_C \quad P_u]^\top, \quad (6.4)$$

where the subscript R indicates that the observations originate from the radar data. Note that the observations from the multi-beam autofocus are omitted in this observation vector, because they rely on an initial estimate. Because of this, these will be fused in a later stage. The output of the Xsens MTi-G-710 used on the experimental setup (Appendix C) consists of an already fused, GNSS and IMU position estimate. In addition, there is also a GNSS velocity estimate. Since these estimates are the output of the Kalman filter present in the INS, they can be seen as a separate state vector

$$\mathbf{x}_I = [x \quad \dot{x} \quad y \quad \dot{y} \quad z \quad \dot{z}]^\top, \quad (6.5)$$

where the subscript I indicates that the state estimates originate from the INS. The INS observations measure the states directly, while the radar observations are non-linear functions of the states. Although it is not specifically stated in the user manual of the Xsens g710, judging from the specified accuracies, the velocity vector is obtained using the instantaneous Doppler shift of the GNSS signal. It is also unclear whether the velocity estimate is already incorporated into the position estimate output by the INS. We will look at this in more detail in Chapter 7 when evaluating the accuracy of the obtained position estimates using quality metrics for SAR imaging.

At this point, there is a choice between converting the measurements to Cartesian coordinates and implementing a standard Kalman filter or leaving the measurements in their original domain and linearizing the measurement matrix at each state. Recall from the previous chapter that the radar velocity estimate is comprised of a magnitude and an angle. The magnitude and course are also estimated using different mechanisms, indicating that their noise is independent. This means that there is a coupling between the x and y directions which changes for different course angles. As the jacobian of the measurement matrix is not too difficult to compute and has already been shown in Chapter 3, the Extended Kalman Filter (EKF) will be used. The relationship between the radar measurements and the system states is given by

$$h_R(\mathbf{x}) = \begin{bmatrix} \sqrt{\dot{x}^2 + \dot{y}^2} \\ \arctan(\frac{\dot{y}}{\dot{x}}) \\ z \end{bmatrix}, \quad (6.6)$$

and its Jacobian is given by

$$\mathbf{H}_{Rt} = \begin{bmatrix} 0 & \frac{\dot{x}}{\sqrt{\dot{x}^2 + \dot{y}^2}} & 0 & \frac{\dot{y}}{\sqrt{\dot{x}^2 + \dot{y}^2}} & 0 & 0 & 0 & 0 & 0 \\ 0 & -\frac{\dot{y}}{\dot{x}^2 + \dot{y}^2} & 0 & \frac{\dot{x}}{\dot{x}^2 + \dot{y}^2} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}. \quad (6.7)$$

6.1.3. Noise Model

To correctly calculate the Kalman gain and apply the update equations, the noise in both the system model, and state observations needs to be modeled using a set of covariance matrices. Namely, the process noise covariance \mathbf{Q} , and the measurement noise covariance \mathbf{R} . The process noise will be modelled as an extra unpredicted perturbation in acceleration. This is known as a piecewise white noise model [24]. A common choice is to choose the standard deviation to be around half of the maximum expected acceleration. The variance of the acceleration is then given by

$$\sigma_a^2 = \frac{a_{max}^2}{4}. \quad (6.8)$$

The variance for each direction is given by

$$\mathbf{C} = \begin{bmatrix} \sigma_{a,x}^2 & 0 & 0 \\ 0 & \sigma_{a,y}^2 & 0 \\ 0 & 0 & \sigma_{a,z}^2 \end{bmatrix}. \quad (6.9)$$

The process covariance matrix can then be calculated by defining a matrix

$$\mathbf{G} = \begin{bmatrix} \frac{d_T^2}{2} & 0 & 0 \\ d_T & 0 & 0 \\ 1 & 0 & 0 \\ 0 & \frac{d_T^2}{2} & 0 \\ 0 & d_T & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \frac{d_T^2}{2} \\ 0 & 0 & d_T \\ 0 & 0 & 1 \end{bmatrix}. \quad (6.10)$$

The covariance matrix can then be found through

$$\mathbf{Q} = \mathbf{GCG}^\top. \quad (6.11)$$

Since it is difficult to determine the variances of the measurements in absolute values, a heuristic approach will be used to determine the quality of an observation, and scale the variance with respect to an initial value. For the velocity magnitude estimate, the variance is scaled by the ratio between the peak of the gradient and the average background value. For the course estimate, the variance is scaled by the magnitude of the velocity vector.

It is difficult to determine the quality of the radar height estimate. This is because while the height estimate itself may be accurate, its accuracy in representing the motion of the radar depends on the terrain and how well the elevation is compensated by the DEM. Because of this, the variance of the height estimate will be left as a constant tunable parameter.

The Xsens INS estimates and radar estimates can technically be fused directly. However, since the INS data is the state estimate of a previous Kalman filter, the noise between different states is correlated according to the covariance matrix of its state estimate. As this matrix is unknown, the fusion result might be sub-optimal. Because of this, the radar estimates will also be considered separately and compared to the INS estimates for a qualitative indication of performance. The noise covariance of the INS is set to a constant diagonal matrix containing the variances specified in the datasheet [1].

6.1.4. Initial Results

We first look at estimates from two simulated data sets, to observe the behavior of the system model. Estimates from both a constant acceleration data segment and sinusoidal motion segment are provided as input to the proposed Kalman filter. The estimated position can be seen in Figure 6.1. This figure shows both the strength and weaknesses of a given model. Since sinusoidal motion has a constantly changing velocity and acceleration, it is difficult to track using a constant acceleration model. On the other hand, the model can predict constant accelerations in the right figure quite well.




Figure 6.1: This figure shows a comparison of the constant velocity Kalman filter output for two different simulated maneuvers in the horizontal plane. The Kalman filter uses only the radar velocity estimate to obtain the position. The same parameters were used in both scenarios

The estimates from the INS and the radar are sampled at different frequencies. This problem can be solved by using a variable time step and sorting estimates in the order in which they arrive. This means that the measurement matrix is also updated at each time step to reflect the available measurements. The problem with this approach is that it can cause discontinuities in the filter output. This effect can be alleviated with a better model of the system dynamics but is still present to some extent. Wang et al.[51] propose to fit a Bezier curve to the low-frequency data points and the next prediction, effectively interpolating the measurements. Since the goal is to obtain a smooth estimate of the position able to be used for SAR imaging, the radar estimates will be interpolated to ensure that the trajectory does not contain discontinuities. The radar estimates from two different experimental data segments are used as an input to the Kalman filter. The first segment consists of flight along a straight line, while the second segment contains some maneuvers. The state is initialized using the first entry in the Xsens position estimate. Since the estimates from all fused sensors provided no qualitative insight, we compare the estimated position using only radar estimates first, to get an indication of the positioning accuracy of the radar estimates. This is shown in Figure 6.2. Figure 6.3 shows the horizontal error between the radar position estimate and the Xsens position estimate. The error in the height estimate can be seen in Figure 6.4. In Chapter 7, the results for different fusion configurations will be quantitatively compared using SAR imaging quality metrics.




Figure 6.2: This figure shows a comparison of the constant velocity Kalman filter output for two different experimental data segments in the horizontal plane. The left image shows the estimate for a 45 second data segment consisting of linear flight in the horizontal direction. The right image shows the estimate for a 30 second data segment following an arbitrary flight path. The Kalman filter uses only the radar velocity estimate to obtain the position. The same parameters were used in both scenarios




Figure 6.3: This figure shows the horizontal error in meters between the radar-only output of the Kalman filter and the Xsens position estimate as shown in Figure 6.2




Figure 6.4: This figure shows the horizontal error in meters between the height output of the Kalman filter and the Xsens height estimate.

The initial results are very promising, as the estimates with experimental data are only performed using 8 antenna elements on one side of the antenna array. We have shown in Chapter 5 that the RMS error of the velocity estimate is greatly reduced by incorporating estimates from multiple directions, reducing both the bias caused by non-zero attitude and the error caused by scene contrast.

6.1.5. Tightly Coupled Kalman Filter

Ideally, the radar estimates should be fused with the raw IMU and GNSS data. Unfortunately, there was no access to the Xsens raw data for the current experiment. To still achieve this to some extent, the mission logs from the drone flight computer were downloaded. The following measurements are available:

- The GNSS coordinates in latitude and longitude.
- Barometer readings from 2 separate sensors.
- Magnetometer readings from 3 separate sensors.
- Accelerometer readings from 3 separate sensors.
- Gyroscope readings from 3 separate sensors.

It should be noted that these sensor readings may be of lower quality than the Xsens MTi-G-710 readings, as the Xsens unit is a high-grade "flagship" INS. The estimates for a 30 s data segment were processed by the Intelligent Autonomous Systems department at TNO, using their proprietary Kalman filter implementation. The filter is based on an Unscented Kalman filter using a kinematic model. Aside from position, velocity, attitude, and acceleration, the Kalman filter also estimates the bias of the attitude and inertial sensors. We will present the results from the filter output to more accurately show the effect that the presence and absence of different sensors have on the accuracy of the state estimate.

All sensors other than GNSS will be grouped under the term "inertial sensors". The Kalman filter was applied to the dynamic data segments using 4 different configurations.

- First, the navigation performance of the inertial sensors is analyzed without any external sensors.
- The second configuration incorporates the standard GNSS and inertial sensor fusion implementation.
- The third configuration incorporates estimates from radar along with the inertial sensors.
- The final configuration incorporates data from all sources into the Kalman filter framework.

The fused results can be seen in Figure 6.5. Again, for now, we can only make a qualitative conclusion on the results of radar-only navigation. When using only the inertial sensors, the estimate quickly drifts and becomes unusable within a matter of seconds (Figure 6.6). By incorporating the velocity estimate from the radar, both the position estimate, and IMU bias estimate can be maintained (Figure 6.7). This is an important result as it stresses the value of an accurate velocity estimate in a GNSS-denied scenario.

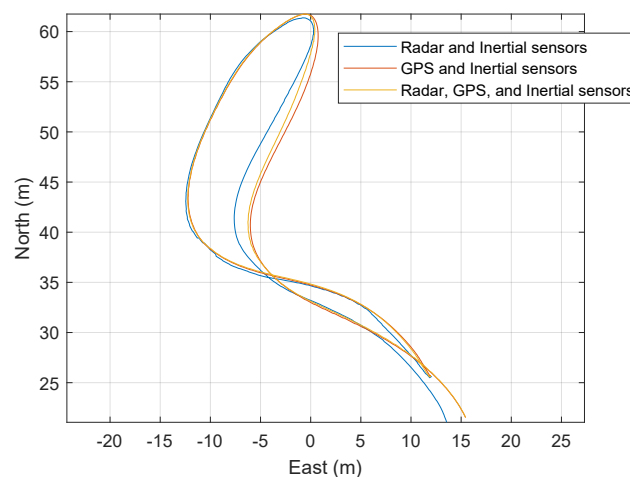


Figure 6.5: Sensor fusion result for different Kalman filtering configurations using a tightly coupled Kalman filtering approach. The total duration of the flight path is 30 seconds.

In practical scenarios, the bias in the inertial sensors may already have been estimated before the GNSS signal is lost. Thus navigation using the inertial sensors may perform slightly better in practice. However, the




Figure 6.6: The position estimate using only inertial sensors. Due to bias in the sensors, the estimate drifts very quickly and is practically unusable.

bias in the sensors is typically not constant as is shown in Figure 6.7. While the estimated bias is relatively constant over a short period, it slowly changes over time.




Figure 6.7: This figure shows the bias in both the accelerometers (left), and gyroscopes (right) as a function of time when using only radar as an additional sensor. The plots (unlabeled) in each figure correspond to both y and x biases. The bias slowly converge to a semi-constant value.

6.2. Incorporating Multi-beam autofocus

While we did not verify the performance of the proposed autofocus algorithm with experimental data, some insight can be provided on how it can be optimally combined with the other proposed methods. When presented, the trajectory corrections obtained were directly applied to the radar trajectory. However, in experimental scenarios, we can expect the phase corrections for each region to be noisy. Because of this a better approach would be to incorporate the phase updates into the EKF framework as a radial velocity in the direction of the focusing regions.

6.3. Conclusion

In this chapter, we combined the course and velocity estimates from the radar and incorporated them into a constant acceleration EKF. The position estimates were qualitatively compared to the output of the INS, and the observed horizontal error was found to be around 2 meters for data segments of 45 and 30 seconds when using only radar data in the EKF. We also analyzed the results from a tightly coupled Kalman filter which uses raw data from the inertial sensors. From this, it was qualitatively shown that the radar velocity estimate is a very valuable addition when a GNSS signal is not available. Sensor fusion between the Xsens INS and

radar was also performed. However, the figures lack enough information to convey the advantage of such a configuration. In the next chapter, we will develop a quantitative approach to determine the accuracy of the position estimate using SAR imaging and compare the results from different sensor fusion configurations.

III

Evaluation and discussion

7

Quantifying Performance

In the first part of this thesis, we introduced our problem and the background behind the tools required to solve it. We also analyzed similar problems in literature to use as a starting point. In the second part, we developed a processing chain to determine both the height of the radar system above ground, and the course and velocity magnitude in the horizontal plane. This was done by exploiting the specifics of the radar and making use of its omnidirectional beamforming capabilities.

The developed methods were quantitatively tested using simulated data, and qualitatively compared to the INS output. The last chapter combined all sub-methods into an EKF framework to estimate the position of the radar from the different sensor inputs. Here, we qualitatively compared the output of the radar-only EKF and found it to be very promising. We also showed that the radar velocity estimate is a very useful addition in GNSS-denied scenarios through the use of a tightly coupled Kalman filter using raw inertial data.

In this chapter, we will quantify the performance of the proposed methods using backprojection SAR imaging. Since an accurate estimate of position is required to obtain a high quality image, we can use image quality metrics to quantify the accuracy of the position estimates. The first section will introduce the metrics used to test image quality. This will be followed by a comparative analysis using the two experimental data segments as detailed in Appendix C.

7.1. SAR Quality Metrics

The chosen image quality metrics need to test both the short term accuracy and the long term accuracy of the position estimates. To do this, we will divide the data segment into a number of sub-apertures. The short term accuracy is then tested by examining the focus of each sub-aperture, and the long term accuracy is tested by comparing the registration between subsequent sub-apertures. Appendix C shows the details of the measurement setup. There are a number of reference objects in the scene, including 2 corner reflectors and a vehicle labeled "Jeep". These reference objects will be used to determine the focus and registration of each SAR image.

The performance tests are split between a number of different cases, each employing different sensor combinations. The cases are defined as follows:

1. The first case consists of the position estimate output by the INS without any post-processing. Since the position output of the INS is a Kalman filtered output, this represents a standard baseline performance for a typical user.
2. In the second case, only estimates from the radar (velocity, height) will be used in the Kalman filter framework described in the previous chapter.
3. The third case will combine both the position output from the INS and its velocity output in the Kalman filter.
4. The fourth case will combine all estimates in the Kalman filter: radar height, radar velocity, INS position and INS velocity.

For each sub-aperture, 814 pulses are used, yielding an integration time of 1 second. The trajectories consist of both constant motion, turns, and decelerations to test the performance in a number of different scenarios. The images are created using a standard backprojection algorithm as outlined in Chapter 2.

7.1.1. Peak To Background Ratio

To quantify the focus of the image, and get an indication of the detectability of targets, the ratio of the peak squared intensity and the average squared background intensity of the image will be used. This is known as the Peak to Background Ratio (PTBR) [36]. The PTBR is given by

$$\text{PTBR} = \frac{\max[I(k_x, k_y)]^2}{\langle I(k_x, k_y) \rangle^2}, \quad (7.1)$$

where I represents the pixel values of the SAR image indexed by k_x and k_y , and $\langle \cdot \rangle$ denotes an averaging operation. As a corner reflector will be used to determine this metric, the PTBR will be highest when all pulses add coherently.

7.1.2. Spatial Resolution

Aside from the detectability of targets, their resolvability is also important. This is determined by the resolution of the SAR image. The resolution can be determined by examining the response from a point target. In this case, the corner reflector is used as an experimental point target. As the radar movement is dynamic, we will average the spatial resolution in both dimensions. The resolution is determined using the half power width of the response from the corner reflector. This is illustrated in Figure 7.1.




Figure 7.1: The image on the left shows the response from the corner reflector for a single sub-aperture. The image on the right shows the image thresholded at -3dB from the peak value. The average between the East and North resolution of this response is defined as the resolution of the image.

7.1.3. Image Registration

While a short integration time may yield a focused image, short integration times do not reflect the long term positioning accuracy. To determine the relative positioning accuracy with respect to a certain initial position, we can use both the known geometrical locations of the reference objects, and the SAR image registration between sub-apertures. Errors in the velocity will be most noticeable for course errors by a rotation of the entire image around the aperture location. The rest of the errors will be seen as absolute shifts in the position of the reference targets. In a near-range case, errors in positioning can cause distortion in the scene as the error is not constant over the entire image. An illustration of the effects caused by errors in course, cross-range, and height, can be seen in Figure 7.2.

Aside from focus and registration we can also examine the geometrical location of the reference objects with respect to the actual known location. This allows us to determine how well the radar estimates can track its position from the last known GNSS location. In addition, it also gives an indication of the noise in the position estimates of the INS.




Figure 7.2: An illustration of the registration errors caused by relative errors between different sub-apertures. The first image shows the effect of an error in course, the second image shows the effect of an error in cross range, while the third image shows that a 3D error can cause distortion in the image due to the range to each target being different.

7.2. Short-term Performance

A visual comparison between the four different cases for a single sub-aperture from a mostly linear trajectory is shown in Figure 7.3. From the defocused image in Figure 7.3a, we can already observe that the INS produces a suboptimal position estimate. Visually, using a radar-only approach (Figure 7.3b) produces a similar result as fusing the position and velocity estimate from the INS (Figure 7.3c). The same can be said for the case when all estimates are combined (Figure 7.3d).

The same experiment was repeated for a different trajectory, where the radar is decelerating into a turn. The results are seen in Figure 7.4. Here it can be seen that the radar estimates perform better than the INS under the influence of an acceleration. We can also see some geometrical distortion around the radar trajectory when using the INS estimates caused by an incorrect height estimate.

To quantify the image focus for the linear trajectory, we zoom in on the corner reflectors, and compute the average PTBR and spatial resolution for a number of sub-apertures. The total duration used for each trajectory is highlighted in Figure 7.5. The PTBR for subsequent 1 second sub-apertures is then calculated for each case, using both the linear and dynamic trajectory. Figure 7.6 shows the normalized average PTBR for the linear trajectory. We can apply the same method to the dynamic trajectory. However, due to the orientation of the active antenna elements, the corner reflectors are not properly illuminated for a large part of the trajectory. Because of this, "Corner1" and "Jeep" will be used to compute the PTBR. For "Corner1", the trajectory from 13 s to 17 s is used. For "Jeep", the entire highlighted segment from 13 s to 26 s is used. While the jeep is an extended target, it consists of strong point-like reflections at each aperture location, making it useful for comparison. The average PTBRs for the dynamic trajectory are shown in Figure 7.8.

The spatial resolution is computed for both corner reflectors in the linear case using the entire highlighted trajectory. The average resolution can be seen in Figure 7.7. In the dynamic case the average spatial resolution is calculated for corner reflector 1 using the trajectory from 13 s to 17 s. This is shown in Figure 7.9.

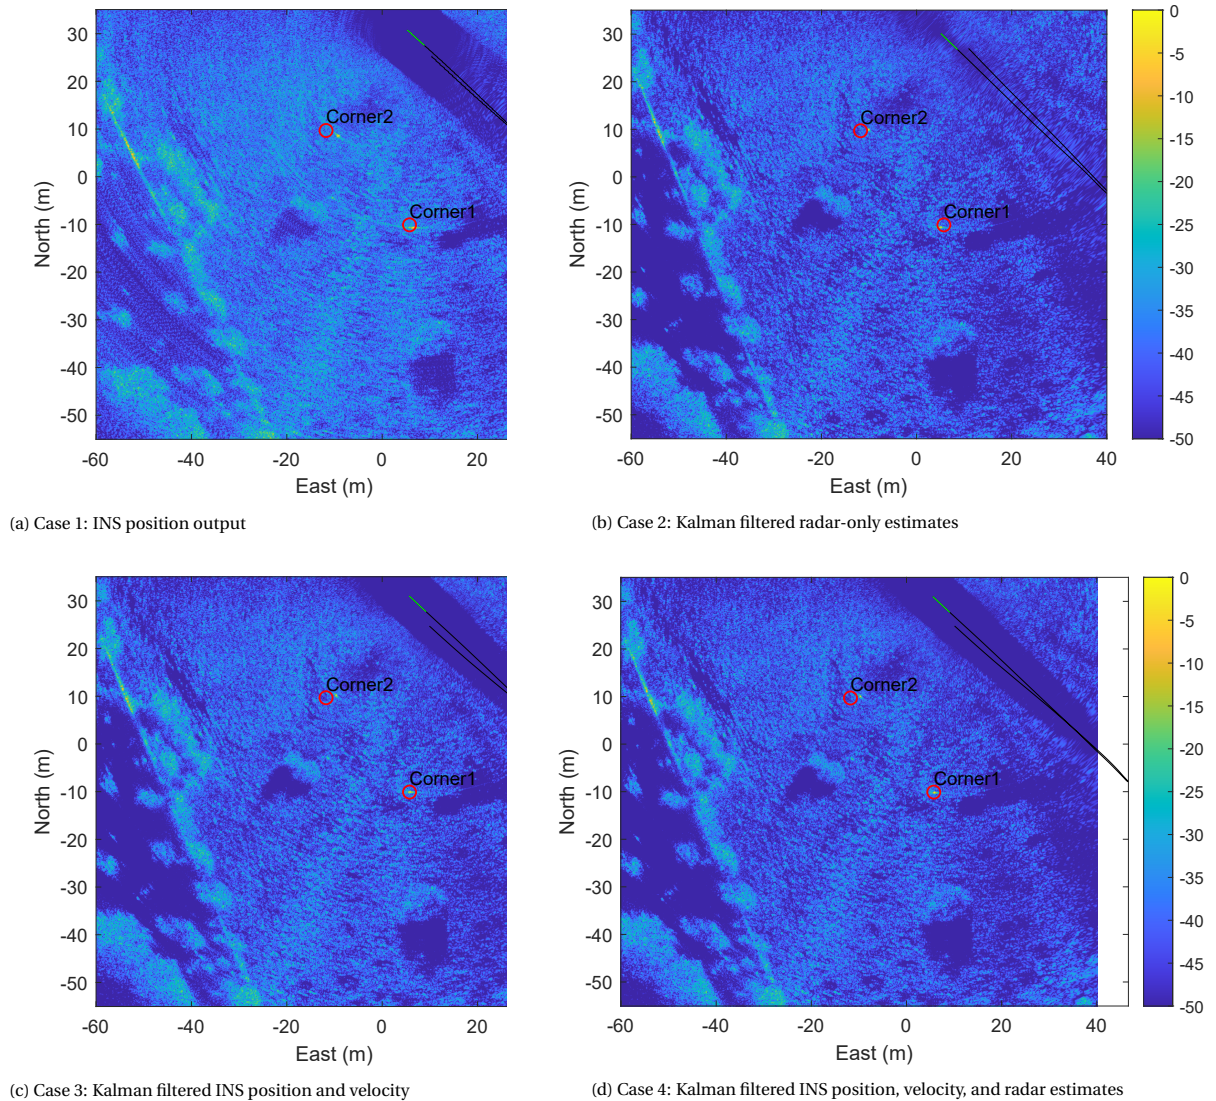


Figure 7.3: This figure shows SAR images for a single sub-aperture from a linear trajectory for 4 different cases. The velocity of the radar during acquisition is relatively constant. The radar trajectory is shown in black and the 1 s integration period is highlighted in green. The ground truth location of the reference objects are highlighted as well. The images are displayed in dB. Details about the measurement setup can be found in Appendix C. The integration time starts at 18 seconds from the start of the recording.

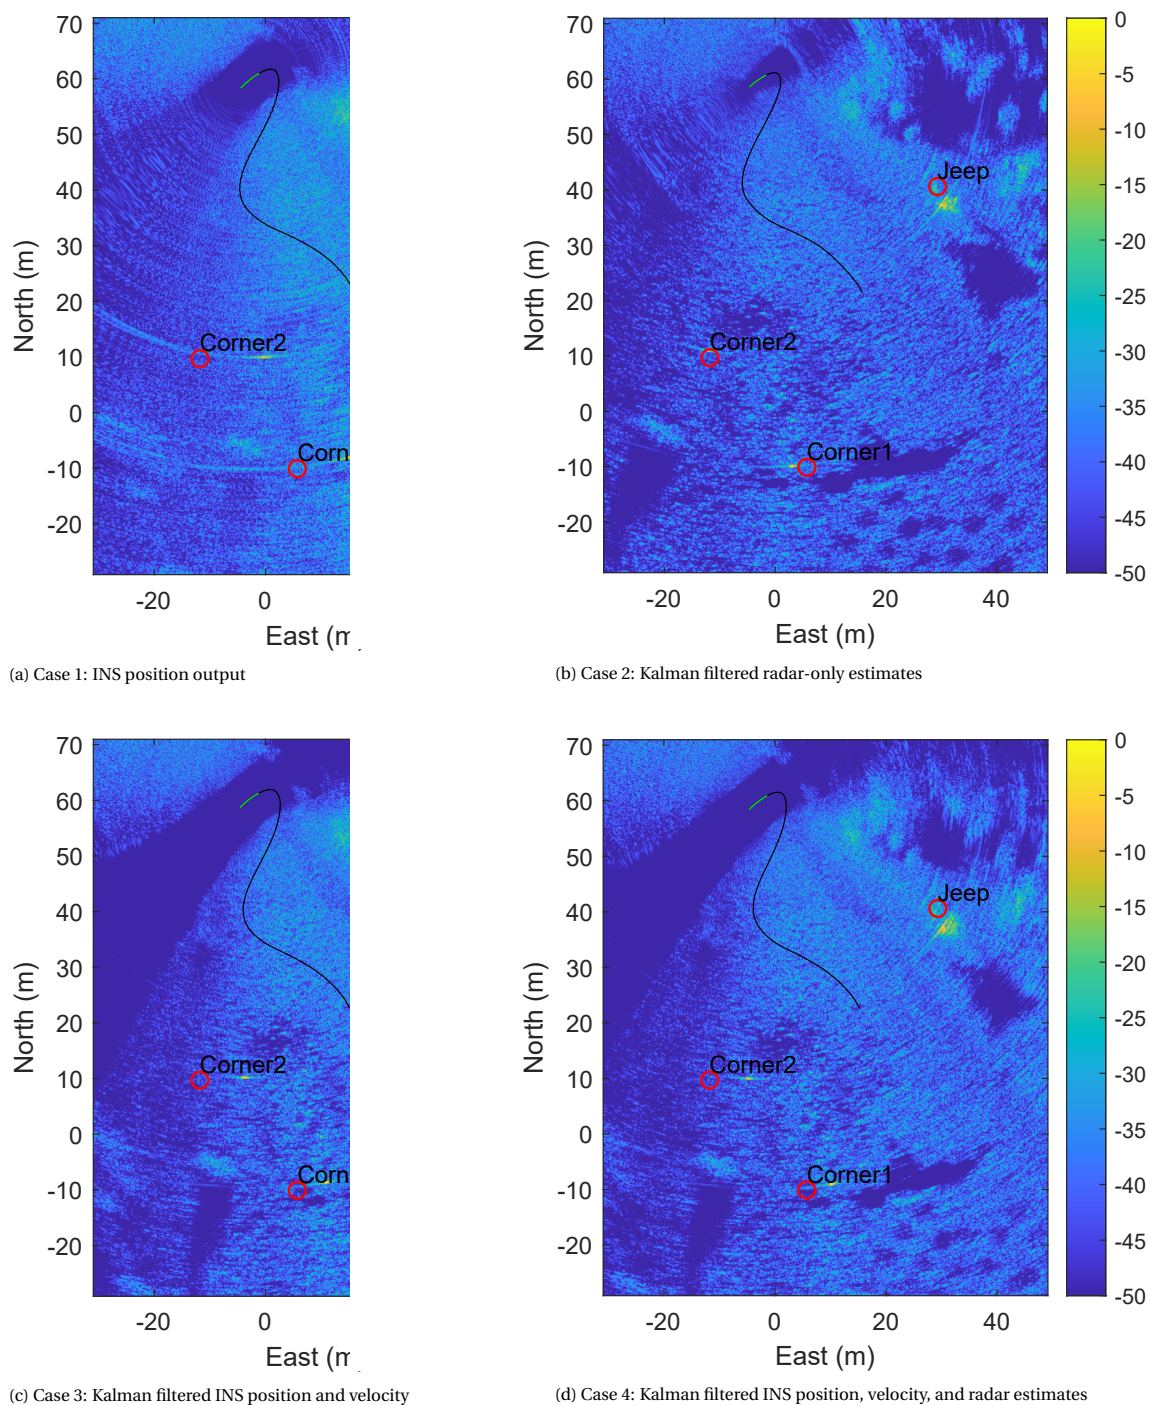


Figure 7.4: This figure shows SAR images for a single sub-aperture from a dynamic trajectory for 4 different cases. The radar is turning and decelerating during the integration time. The radar trajectory is shown in black and the 1 s integration period is highlighted in green. The ground truth location of the reference objects are highlighted as well. The images are displayed in dB. Details about the measurement setup can be found in Appendix C. The integration time starts at 16 seconds from the start of the recording




Figure 7.5: The period used to calculate the average image focus for both the linear (left) and dynamic(right) trajectory. The segment highlighted in green corresponds to the total trajectory used to determine the image focus. Note that this period is divided into 1-second sub-apertures. The total length of the integration time is 13 seconds and 14 seconds for the linear and dynamic case respectively. The time in seconds at each aperture location is also shown.




Figure 7.6: The average PTBR of both corner reflectors for each sub-aperture for the linear trajectory. Each graph value corresponds to a different sensor configuration. The PTBRs for each corner reflector are normalized to the baseline INS value to make it easier to compare the results.

It can be concluded that for the linear trajectory, on average, the position estimates obtained from the radar produce a better focused image than the sensor configurations using only INS position estimates. This result is consistent for both corner reflectors. Similar performance is obtained by fusing INS position and velocity estimates. In the case where estimates from radar and INS are combined in a Kalman filter, a further improvement is seen in both the PTBR and resolution, surpassing both a radar-only and INS-only approach.

The performance metrics are more difficult to evaluate in the dynamic case, as the reference targets were not properly illuminated during the entire trajectory and were in the forward looking regions of the antenna. However, the evaluated PTBR for the first corner reflector and jeep target show consistent result, indicating increased performance for fused data. The images obtained in in Figure 7.4 also provide some empirical evidence to the added value of the radar estimates.

The performance can probably be further improved by tuning the filter and/or applying more advanced adaptive Kalman filtering techniques. However, the result shows that for a 1-second coherent integration time, the radar position estimates significantly improve upon the performance of an off-the-shelf INS.




Figure 7.7: The average spatial resolution of each corner reflector over all sub-apertures for the linear trajectory. Each graph value corresponds to a different sensor configuration. The shown resolution is the average resolution in x and y for the aperture positions highlighted on the left in Figure 7.5.




Figure 7.8: This figure shows the average PTBR for corner reflector 1, and the jeep present in the scene using the dynamic trajectory. The PTBR for the corner reflector is calculated using 5 one-second sub-apertures from 13 s to 17 s as shown in Figure 7.5. For the jeep, the entire highlighted segment from 13 s to 17 s is used. The PTBRs are normalized to the baseline INS value to make it easier to compare the results.

7.3. Long-term Performance

While it is already clear that combining estimates from radar with the INS produces better focused images, the position estimate may drift over time. In this section, we will examine the long term positioning accuracy using each sensor configuration. This will be done by determining the position of the corner reflector in a series of SAR images and calculating the course and range error responsible for the observed offset. The locations of the reference objects were measured using handheld GNSS devices and can be considered as ground truth. The course offset relative to the ground truth can be determined from the angle between the vector from the current aperture to the ground truth location and the vector from the current aperture to the observed location. The calculated course offset is shown in Figure 7.10 and the calculated range error is shown in Figure 7.11.

As the Kalman filter was tuned based on the image quality, relatively large emphasis is placed on the velocity estimates. This is also observed in the relative course error. Overall, the course error of the fused estimate relative to the initial location is low, drifting approximately 2 degrees from the initial location over 13 s for the linear trajectory, and 2 degrees over 6 s for the dynamic trajectory. The added benefit of the INS estimates can be seen in the dynamic trajectory in Figure 7.10, resulting in a reduced drift compared to the radar-only estimate. The Kalman filter is initialized with the INS positions at 0 s. This indicates that the radar




Figure 7.9: The average spatial resolution over the sub-apertures from 13 s to 17 s using the dynamic trajectory. The resolution is calculated using corner reflector 1.

velocity estimate has a relatively low bias, as the position estimate is maintained with relatively low bias up to the observation window.

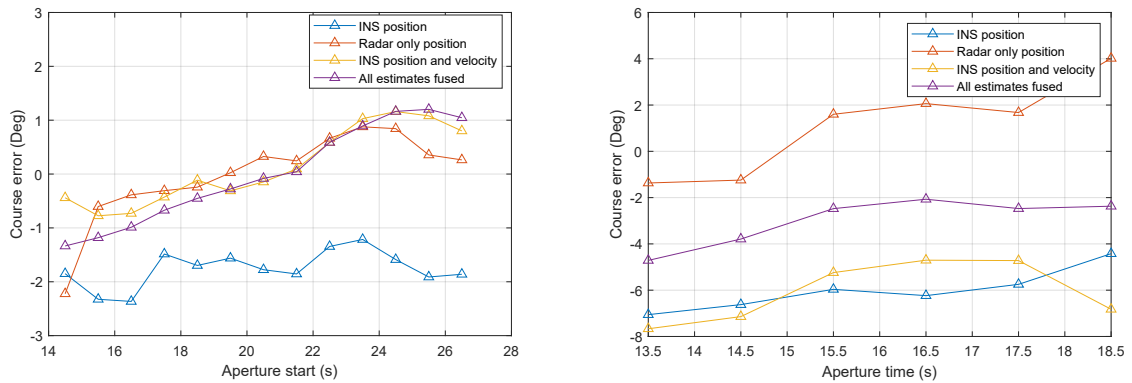


Figure 7.10: The course error along each position with respect to the ground truth course. The image on the left shows the course error for the linear trajectory and the image on the right shows the course error for the dynamic trajectory. The ground truth course is determined using the vector from the actual location of the corner reflector, to the estimated center of the current aperture.




Figure 7.11: The range error along each position with respect to the ground truth range. The image on the left shows the range error for the linear trajectory and the image on the right shows the range error for the dynamic trajectory.

8

Conclusion and Recommendations

The research question formulated at the start of this thesis was framed in the context of coherent processing on an agile radar platform; specifically SAR imaging. As such, we presented the signal processing principles behind the system of interest in Chapter 2, introducing both the concept of FMCW radars and the principles behind digital beamforming and coherent processing. In Chapter 3, we looked towards the literature and presented the state of the art techniques in radar-aided motion estimation/compensation.

Using the field of SAR motion compensation as a starting point, we analyzed the specifics of a drone-radar in Chapter 4 to determine which kinematic states could be observed in radar data, and how they were represented. With this knowledge, we provided extensions to current state-of-the-art techniques in radar altimetry, Doppler centroid estimation, and autofocus in Chapter 5. Here we proposed new, generalized algorithms for dealing with data from an omnidirectional array, and provided adaptations to deal with the effects of attitude and high squint angles. We also introduced a new autofocus technique for omnidirectional SAR systems termed multi-beam autofocus, which is able to correct the position of the radar while focusing an image. This technique was verified using simulated data.

After obtaining the estimates from radar data, we turned to sensor fusion in Chapter 6 and presented an EKF framework for combining estimates from radar data with an INS. We also showed results from a tightly-coupled Kalman filter implementation, stressing the added benefit of radar velocity estimates in GNSS-denied environments.

In Chapter 7, we evaluated the performance of the developed velocity estimator on experimental data using quality metrics for SAR imaging, giving a quantitative result to the developed methods. In this chapter, we reflect on our initial research questions and answer them based on our findings.

8.1. Conclusion

Based on the results presented throughout this thesis, most notably the performance evaluation in Chapter 7, we can now formulate concise answers to our research questions.

- ***What are the current algorithms being employed for positioning and motion compensation in high-resolution radar modes?***

In Chapter 3, we found that there is limited research addressing the problem of radar-aided positioning for multi-copter platforms. The methods presented in the literature deal mostly with antenna arrays looking to one side, typically forward, and are usually based on the velocity of strong targets in the reference pulse. The field of SAR motion compensation provides forms of velocity estimation, and coherent scan-matching by methods such as Doppler centroid estimation and autofocus. However, these methods typically only provide one-dimensional motion compensation for fixed-wing radar platforms. Nevertheless, these methods provided a robust framework for the techniques developed throughout this thesis.

- ***Which relevant parameters can be estimated from radar data, and by what means?***

By analyzing simulated radar data in Chapter 4, we found that the horizontal velocity of the radar can be estimated using the Doppler shift of the background clutter, eliminating the need for strong reference targets. We also found that the height can be estimated from the ground reflection, providing an absolute position in the Up direction. We also showed that relative position from pulse to pulse can be estimated using the autofocus method presented in Chapter 5.

- ***How can the circular radar array be used to improve the accuracy of the parameter estimation techniques?***

In Chapter 5, we developed techniques to estimate the height and horizontal velocity of the radar, and showed that combining estimates from multiple digital beams reduced the variance of the overall estimate. In simulations, incorporating narrower beams reduced the variance further. However, this might not be the case in practical scenarios with high contrast, as the estimate is more easily skewed. We also found through simulated scenarios that a digital Janus configuration is effective in suppressing the offset caused by a non-zero attitude. The proposed velocity estimation technique was tested on experimental data using a subset of 8 antenna elements and was found to be close to the error margin of the INS estimates. Additionally, we also showed how the idea of autofocus can be extended to an omnidirectional array by focusing multiple regions using multiple independent digital beams, potentially providing 3D trajectory corrections from pulse-to-pulse on the order of a single wavelength using only a rough estimate of the initial velocity.

- ***How can the estimates from the circular radar array be used to improve the positioning accuracy of the radar platform?***

In Chapter 6, we transformed our radar estimates into position estimates using an EKF employing a constant acceleration model, providing a sensor-fusion framework for radar and INS observations. Here we showed that position estimates obtained from radar data only are accurate enough to stay within 2.5 m of the INS estimate over a period of 45 s. With this result, we showed the potential for GNSS-denied navigation with our method, opening avenues for further research in autonomous radar-only navigation.

- *To what extent can onboard INS-based positioning accuracy be improved by incorporating data from an omnidirectional radar array?*

In Chapter 7, the performance of the developed methods were quantified using SAR image quality metrics on experimental data segments, providing a definitive answer to the main research question. It was found that in terms of peak-to-background ratio (PTBR) and resolution, using the radar estimates significantly improved SAR imaging quality compared to the baseline INS approach. For a standard SAR acquisition trajectory, the average PTBR increased by a factor of 3.9, while the average resolution is improved by a factor of 1.7. This is a very important result since the general consensus is that backprojection, especially on agile platforms, requires the position to be measured accurately using some inertial sensor. With our results, we have essentially shown that focused images from arbitrary trajectories can be formed by exploiting the beamforming capabilities of the radar array and estimating motion explicitly, practically eliminating the need for an INS.

In the case where both radar and INS data is available, we have shown that fusing these two estimates yields an additional focusing performance, surpassing both a radar-only and INS-only approach. The PTBR increased by a factor of 4.2 compared to the baseline approach, while the resolution is improved by a factor of 2. The improved image focus does not only yield a higher quality SAR image, but the observed increase in PTBR also suggests an increased detectability of weak targets in strong clutter, opening further avenues for target detection through SAR imaging.

8.2. Recommendations

While the methods developed in this thesis provided very promising initial results, there are some aspects that could be improved by further research.

While accurate, the proposed velocity estimator is currently limited by the pulse repetition frequency (PRF) of the radar. The velocity estimation can be made more robust by adaptively changing the estimation method when the PRF threshold is crossed. When the forward velocity is higher than the ambiguity limit, the velocity estimation should be performed on the side looking beams for which this threshold is not yet crossed by estimating the centroid velocity instead of the centroid angle.


The proposed multi-beam autofocus was only verified using simulated data since only 8 elements of the receiver array were available in experimental data. This technique should be further investigated using omnidirectional experimental data. It should also be investigated whether the estimates obtained from autofocusing should be incorporated in the initial sensor fusion step, or applied at a later stage. In addition, there are a number of parameters that can be experimented with. For example, the number of antenna elements used per focused pulse, and the number of pulses which are focused together. Alternatively, the focusing could also be applied between short subsequent sub-apertures to further reduce noise and interference. The autofocus principle could also be used to determine other parameters such as attitude. This could be done by maximizing the image sharpness as a function of the radar orientation.

A non-coherent scan matching approach can also be employed to estimate motion from pulse to pulse in high contrast environments. This has the advantage of being able to better exploit structure in the scene compared to the proposed coherent approach. Using this approach, the drift in course over longer periods could be reduced. For increased correlation, the scan-matching can be performed on short subsequent synthetic apertures.

The proposed sensor fusion algorithm was tested using the filtered output of the INS. A further improvement in positioning performance could be realized by using a more accurate kinematic model. In addition, a tightly coupled Kalman filter using raw sensor estimates could also result in increased performance by incorporating a more accurate estimate of the noise covariance. The sensor fusion step can also be further tuned using more advanced adaptive filtering.

In conclusion, we have shown that 3D radar-aided positioning can be achieved using an omnidirectional radar-array, and provides a significant improvement in positioning accuracy. Its application to practical scenarios provides an interesting multidisciplinary challenge for future research.

Appendix



SAR motion Compensation

A.1. Doppler centroid estimation

A.1.1. Frequency domain Estimation

Frequency domain Doppler centroid algorithms can be categorized by the window functions used in the correlation step. The most notable window functions are energy balancing, matched-correlation, and maximum likelihood estimation.

Energy balancing

The Doppler centroid can be calculated by observing the power spectrum of the SAR data along a constant range azimuth line[34]. Assuming the power spectrum is symmetric, the Doppler centroid can be found by finding the energy centroid of the spectrum. This is also known as energy balancing. It can be described by using the following reference function:

$$R(f) = \begin{cases} 1 & \frac{B_d}{2} \leq f \leq 0 \\ -1 & 0 \leq f \leq \frac{B_d}{2} \end{cases}, \quad (\text{A.1})$$

where B_d is the bandwidth of the power spectrum.

Matched-correlation

The Energy balancing approach can be improved to take into account the two-way antenna power pattern, $W(f) = G^2(f)$. This can be done by setting the Reference function equal to its derivative [4, 12], $W'(f)$.

$$R(f) = W'(f) \quad (\text{A.2})$$

Maximum likelihood estimation

Under the assumption of complex Gaussian distributed, uniform backscatter, a maximum-likelihood estimator that reaches the Cramer-Rao lower bound can be formulated[4]. The estimator can be written in terms of a correlation based estimator with weighting function

$$R(f) = \frac{G'(f)}{G^2(f)}. \quad (\text{A.3})$$

A problem with applying correlation based methods as-is, is that the estimation will be very sensitive to non-homogenities in the scene. That is, the variance of the estimator is proportional to the scene contrast[34]. This can also be seen as reducing the number of effective samples [4]. This can be improved by using longer estimation windows in high contrast scenes. For a multi-copter platform, this is less than ideal since the velocity of the platform may not be constant over a large number of samples. Another approach presented in [27], is to perform the Doppler centroid on multiple sets of samples grouped by their intensity, creating multiple spectra with similar contrast.

A.1.2. Time domain estimation

The Doppler frequency of a signal can also be estimated by measuring the change of phase from pulse to pulse in the time domain.

For a given azimuth sample $s(\eta)$, the phase difference between another sample separated by $\Delta\eta$ is given by the angle of the average cross correlation coefficient at lag 1 [12].

$$C(\Delta\eta) = \sum_{\eta} s^*(\eta) s(\eta + \Delta\eta)$$

The fractional part of the Doppler centroid as given in Appendix 2.3.4 is given By

$$f'_{\eta_c} = \frac{F_a}{2\pi} \angle(C(\eta))$$

This method is referred to as correlation Doppler estimation (CDE). CDE has been shown to be equivalent to a frequency domain Doppler centroid estimator with windowing function

$$R(f) = \sin\left(\frac{2\pi f}{\text{PRF}}\right). \quad (\text{A.4})$$

In order to lessen the bias from bright targets, the signs of the real and imaginary parts of the signals can be used for cross correlation [34]. This is referred to as sign Doppler estimation (SDE). This also lessens the computational cost, which is beneficial for real time systems.

A.1.3. Image processing techniques

Aside from frequency and time domain methods performed on echo data. There are a number of Doppler centroid methods that are based on approaches related to image processing. Namely, methods based on slope estimation [11, 56], morphological edge detection [35], and image entropy [32].

Slope estimation DCE

Slope estimation based DCEs estimate the Doppler centroid by estimating the associated squint angle of the antenna. This is done by measuring the slope of the target azimuth response. In [56], a method based on the radon transform, called geometry-based Doppler estimation (GDE) is proposed.

This method applies the radon transform on the range compressed SAR data, $g(x, y)$.

$$Rf(\hat{g}(\rho, \theta)) = g(x, y)$$

Each line in the range azimuth plot is projected onto a point with coordinate θ, ρ . Here, θ represents the angle of the line and ρ represents the distance from the image center.

The squint angle is then estimated by finding the angle for which the variance of the radon transform along the range direction is approximately maximized.

To obtain an estimate of the Doppler centroid or platform velocity, either one of these parameters has to be known before hand. For a platform with along track velocity V_{st} , the Doppler centroid is approximately given by

$$f_{\eta_c} = 2 \frac{V_{st}}{\lambda_c} \sin(\theta_s)$$

where θ_s is the squint angle. This means that this method cannot be applied as-is for velocity estimation. However, this method could be useful to determine the general velocity direction by finding the squint angle of beams in different directions. This could be used to apply corrections to the direction of the estimated velocity vector.

Some improvements to this algorithm are proposed in [11]. The most notable ones being:

1. Removing the quadratic range cell migration
2. Performing a localized radon transform
3. Adding estimator quality criteria

This method works best when strong point-like scatterers are present in the scene, as they will have a more defined azimuth slope.

Edge detection DCE

A property of the Doppler spectrum for a forward looking radar, is that it is symmetric. In [35], this property is exploited and applied to Doppler estimation. Because of the symmetry of the Doppler spectrum in the forward looking region, the highest frequency is found at the center of the antenna beam, creating an edge at the highest frequency component.

The Doppler frequency is then extracted using morphological edge detection followed by curve fitting. This method has been shown to achieve good estimation accuracy even in the presence of low signal to noise ratio (SNR). The drawback of this method is that it can only be applied in the forward looking region as the highest frequency component does not represent the Doppler centroid for other squint angles. Since pointing the beam directly forward would require knowledge of the velocity vector, this technique actually estimates the velocity magnitude, and not the Doppler centroid.

Entropy minimization DCE

Another image-based DCE is the minimum-entropy Doppler estimator (MEDE). In this method, originally applied to Doppler beam sharpening [32], the quality of an image is quantified through its entropy. The scene is processed with different Doppler centroid values to iteratively minimize the total entropy given By

$$H(f'_{\eta_c}) = - \sum_{m=1}^M \sum_{n=1}^N D_{m,n}(f'_{\eta_c}) \ln D_{m,n}(f'_{\eta_c}),$$

where $D_{m,n}(f'_{\eta_c})$ is given by

$$D_{m,n}(f'_{\eta_c}) = \frac{I_{m,n}(f'_{\eta_c})}{\sum_{m=1}^M \sum_{n=1}^N I_{m,n}(f'_{\eta_c})}.$$

Here, $I_{m,n}$ is the intensity of a given pixel.

A.2. Autofocus Using Row-By-Row coordinate Descent

A popular method for solving SDPs are so called interior point methods (IP). These are also employed by off-the-shelf solvers such as cvx. However, IP methods are not suited for large problems due to their complexity. This can pose a problem when dealing with a large number of pulses. One low-complexity method for solving SDPs, is the row-by-row method presented in [53]. The row-by-row method exploits the structure of matrix \mathbf{X} to compute cheap updates at each iterate. The row-by-row method has been applied to MIMO detection for BPSK signals [50]. This section will show that the same method can also be applied for the autofocus problem without any significant adjustments.

To start, a barrier function is added to problem 5.25

$$\begin{aligned} & \underset{\mathbf{X}}{\text{minimize}} && \text{Tr}(-\mathbf{C}\mathbf{X}) - \sigma \log \det \mathbf{X} \\ & \text{subject to} && \text{diag}(\mathbf{X}) = (1, \dots, 1) \end{aligned} \quad (\text{A.5})$$

where σ is the barrier parameter. The barrier function ensures that the solution stays within the set of positive semidefinite matrices. as \mathbf{X} and \mathbf{C} are both hermitian matrices, they can be separated into blocks.

$$\mathbf{X} = \begin{bmatrix} X_{11} & \xi_1^H \\ \xi_1 & \bar{X}_{11} \end{bmatrix}, \mathbf{C} = \begin{bmatrix} C_{11} & c_1^H \\ c_1 & \bar{C}_{11} \end{bmatrix} \quad (\text{A.6})$$

where ξ_1 represents the 1st column of \mathbf{X} with the 1st row removed and \bar{X}_{11} represents \mathbf{X} with the 1st row and column removed. The same definition holds for c_1 and \bar{C}_{11} . The optimization of ξ_1 can be written as

$$\underset{\xi_1 \in \mathbb{R}^{n-1}}{\text{minimize}} \quad 2\mathbf{c}_1^H \xi_1 - \sigma \log(1 - \xi_1^H \bar{X}_{11}^\dagger \xi_1) \quad (\text{A.7})$$

This problem has a closed form solution which can be obtained by setting its gradient to zero. the optimal value for ξ_1 is then given by

$$\xi_1^* = \begin{cases} -\frac{1}{2\gamma} (\sqrt{\sigma^2 + 4\gamma} - \sigma) \bar{X}_{11} \mathbf{c}_1 & \text{if } \gamma > 0. \\ \mathbf{0}, & \text{if } \gamma = 0. \end{cases} \quad (\text{A.8})$$

where $\gamma = \mathbf{c}_1^H \bar{\mathbf{X}}_{11} \mathbf{c}_1$. In the backprojection algorithm, the image is constructed by coherently adding pulses. This makes it possible to focus pulses as they are added to the image. Under the assumption that the previous pulses are already focussed, (A.8) only needs to be computed once for each pulse.

B

Raw Data simulator

While real experimental data is always preferred when analyzing performance, a simulation can also provide valuable insight into the system. A simulation enables testing of algorithms in ideal and non-ideal scenarios, while having access to the ground truth of each parameter.

For the purpose of Testing and analysis, A raw signal simulator was developed in Matlab. The simulator simulates the signal received in each receiver channel from a number of point targets. The simulator incorporates the antenna pattern of the receiver in both azimuth and elevation. The simulator can be used to model both strong stationary targets, and background clutter.

B.1. Signal model and Simulation geometry

The signal in each receive channel is simulated on a sweep by sweep basis, where each sweep corresponds to a different location. The location of the receive elements is computed using the parameters of the receive array and a yaw pitch and roll angle.

The antenna pattern is incorporated by calculating the orientation of the receiver with respect to the inertial reference frame. The appropriate antenna gain is then calculated from the relative angle of arrival of a given target signal. For a given half-power beamwidth in azimuth and elevation, The antenna gain is approximated by the following sinc squared function:

$$W_a(\theta_{az}, \theta_{el}) = \text{sinc}^2\left(0.886 \frac{\theta_{el}}{BW_{el}}\right) \cdot \text{sinc}^2\left(0.886 \frac{\theta_{az}}{BW_{az}}\right) \quad (\text{B.1})$$

Where θ_{el} and θ_{az} are the azimuth and elevation angles of the target with respect to the antenna element orientation. The orientation of an antenna element is defined as the normal vector from the center of the antenna ring to the position of the antenna element. BW is denotes the half power beamwidth in the azimuth and elevation dimensions.

B.2. Target Distribution

In addition to specific point target configurations, the simulator can also be used to model background clutter. This can be done by considering the statistics of a typical SAR scene. Typically, there are multiple scatterers in a single resolution cell. The phase of a scattering cell will depend on the position of the radar with respect to each target. As the position of the radar changes, the intrinsic phase of each scatterer is added randomly, resulting in a circular symmetric Gaussian process [34].

The scattering cells can be simulated using a collection of spatially uniform distributed point targets with a circular symmetric Gaussian distribution. For real valued signals, this is equivalent to a Rayleigh distributed magnitude, and a uniformly distributed phase. While this creates the correct magnitude distribution, the intrinsic phase of each scattering cell is stationary. However, If a sufficient number of scatterers are simulated, multiple targets will fall within the same resolution cell, creating the desired statistical distribution as the position of the radar changes over time.

Figure B.2 shows the range-Doppler image of a clutter simulation. For this simulation, approximately 2.5×10^5 point targets are simulated at random locations in a radius of 100m around the radar. The radar has a




Figure B.1: Simulated antenna geometry

constant velocity of 5 m/s, and the beam is steered perpendicularly to the velocity vector using 8 antenna elements. The antenna pattern is clearly visible, and the image also exhibits a speckle pattern which is expected from random background clutter




Figure B.2: Range Doppler image for simulated clutter. For this simulation, 2.5×10^5 point targets were simulated in a radius of 100m. The beam is steered perpendicularly to the velocity vector using 8 antenna elements

B.2.1. Simulated Trajectories

Each simulated trajectory simulates 5 s of data and incorporates around 2.5×10^5 point targets to simulate background clutter. The simulations consist of the following scenarios:

- Linear motion with constant velocity
- Motion with cross range acceleration
- Sinusoidal motion

For each scenario, a data segment without pitch, and with a 15 degree down pitch was simulated.

C

Experimental Setup

C.1. Radar System Description

The experimental radar platform used throughout this thesis is a recent development in drone-mounted radar systems presented by Otten et al. [39]. Here, a novel, omnidirectional radar system is proposed for agile drone platforms. The radar system has been developed, and tested in a variety of experimental scenarios at TNO, the Netherlands. The radar system consists of two circular antenna arrays. The transmitter array consists of 16 elements which can be switched on and off to shape the transmit beam, while the receive array consists of 32 digitized receive channels, allowing for digital beamforming on receive. The system is operated on X-band frequencies, and the transmit signal is modulated with a Frequency Modulated Continuous Wave (FMCW) waveform.

An image of the radar system can be seen in Figure C.1. Figure C.2 shows the radar system mounted to a multi-copter platform. The radar system, including battery, weighs approximately 800 grams and is well-suited for an agile platform. Table C.1 shows a list of system specifications.




Figure C.1: The radar system designed by TNO. The image on the left shows the two antenna rings. The larger ring being the receiver ring and the smaller ring being the transmitter ring. The image on the right shows the sensor in the upright position with the INS mounted on top of the. Also seen on the image are the metal reflectors which direct the beam approximately 30 degrees from horizontal.



Figure C.2: The radar system mounted to a multi-copter platform hovering above the ground.

Table C.1: Overview of the system specifications for the radar system developed by TNO. The half-power beamwidths for the antenna elements are also shown. Note that these correspond to both receive and transmit elements as they are approximately the same.

Data acquisition	
Operating frequency	9 - 10 GHz (X-band)
Waveform	FMCW
Bandwidth	1 GHz
PRF	814 Hz
Antenna	
TX Elements	16
TX Diameter	12 cm
RX Elements	32
RX Diameter	22 cm
Element beam properties	
Azimuth _{-3dB}	100°
Elevation _{-3dB}	50°
Depression angle	30°

Alongside conventional radar modes, the system also allows for high-resolution SAR imaging. Currently, the position of the radar is tracked using the Xsens MTi-G-710; an INS consisting of a gyroscope, accelerometer, magnetometer, barometer, and GNSS module. Each output is combined to provide estimates of velocity, position, and attitude. Its specifications are listed in Table C.2. To establish some consistency throughout the rest of this thesis, the system of interest will be referred to as the 'drone-radar'.

Table C.2: Overview of measurement accuracy for the Xsens MTi-G-710[1].

Parameter	Measurement	Value
Sensor data		
Roll/pitch	Static	0.2
	Dynamic	0.3
Yaw		0.8
Position and Velocity		
Horizontal position	1 σ	1.0 m
Vertical position	1 σ	2.0 m
Velocity accuracy (@30 m/s)	RMS	0.05m/s

C.2. Measurement Setup

To test the developed algorithms in realistic scenarios, and give credibility to the simulated results, several experimental data segments were recorded. All data were recorded in the same area for several different flight paths and receiver configurations.

C.2.1. Recording area

The data was recorded over grassy terrain, consisting of a relatively large open area. The area also contains some vegetation and a wall (figure C.3). In terms of elevation, the terrain is relatively flat, but has a slight slope, resulting in a peak to peak elevation deviation of around 1 m. For reference, several stationary targets were placed in the scene. Namely, 3 corner reflectors differing in size, and a metal box. The Figure below shows a map of the area with the locations of the reference objects highlighted.




Figure C.3: The image on the left shows a photograph of the measurement area from the radar mounted camera. The image was taken with the drone hovering above the ground at an altitude of approximately 10 meters. The legs of the drone are also visible on the image. The image on the right shows a map of the area, with some highlighted reference objects.

C.2.2. Radar setup

The radar was mounted to the drone using a rigid mounting system, which can be seen in Figure C.4. This means that the attitude of the radar is tied to the attitude of the drone. The signal parameters used during recording can be seen in Table C.3.

Unfortunately, due to technical problems in the radar software, only 8 receive channels were able to be recorded simultaneously. To make the most out of the available experimental data, different maneuvers were performed during the recordings to emulate as many different scenarios as possible.

Table C.3: Signal parameters used for all recordings.


Parameter	Value
PRF	814 Hz
Sampling Rate	2 MHz
Center frequency	9.5 Ghz
Bandwidth	1 Ghz

C.2.3. Flight trajectories

This section will give an overview of the flight trajectories for the experimental data segments that are referenced throughout this thesis.



Figure C.4: The image on the left shows a photograph of the drone and radar system with the GNSS antennas visible. The radar mounted camera is also visible. The image on the right shows the radar system mounted to the drone. In the image, the Xsens G710 INS is also visible as the red module on top of the radar.



(a) Linear flight path

(b) Non-linear flight path

Figure C.5: The two trajectories used for the experimental data segments. The trajectory on the left consist of a straight pass in a side looking setup. The trajectory on the right consists of an arbitrary maneuver, more representative of actual flight. The recording lengths are 45 seconds and 30 seconds respectively. The shown trajectory is the output of the INS.

To test the velocity estimator, two data segments are used. These are shown in Figure C.5. Both trajectories use a setup with an array of 8 active receive elements. For trajectory C.5a, the squint is close to zero, while for trajectory C.5b the squint changes throughout the maneuver.

D

Geometry And Pre-processing

This chapter will give an overview of the geometry of the radar, and give some insight into the steps needed to compute the location of each antenna element. The first section will present the definitions used for the global reference frame, and the convention used when defining rotations. The second section will outline some pre-processing steps which precede many of the algorithms presented in this thesis.

D.1. Global Reference Frame and Euler Angles

Throughout this thesis, the global reference frame is defined as East, North, Up. This means that the local axis is x,y,z respectively. Rotations are applied in the positive direction using a right handed convention. The rotations are applied in the following order.

- The first rotation is around the global Z axis (Up) with an angle α (Yaw)
- The second rotation is around the rotated local y axis (y') with an angle β (Pitch)
- The final rotation is around the rotated local x axis (x'') with an angle γ (Roll)

The rotations can also be described by the following multiplication of rotation matrices:

$$\mathbf{R}(\alpha, \beta, \gamma) = \mathbf{R}(\alpha)\mathbf{R}(\beta)\mathbf{R}(\gamma), \quad (\text{D.1})$$

where

$$\mathbf{R}(\alpha) = \begin{bmatrix} \cos \alpha & -\sin \alpha & 0 \\ \sin \alpha & \cos \alpha & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (\text{D.2})$$

$$\mathbf{R}(\beta) = \begin{bmatrix} \cos \beta & 0 & \sin \beta \\ 0 & 1 & 0 \\ -\sin \beta & 0 & \cos \beta \end{bmatrix}, \quad (\text{D.3})$$

$$\mathbf{R}(\gamma) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \gamma & \sin \gamma \\ 0 & -\sin \gamma & \cos \gamma \end{bmatrix}. \quad (\text{D.4})$$

A visual illustration of the order of rotations can be seen in Figure D.1. Figure D.2 gives an illustration of how the reference frame would look for the drone hovering above the ground.




Figure D.1: The order of rotations as used throughout this thesis. the accents above each axis represent how many times that axis has already been rotated.




Figure D.2: The global and local reference frame visualized for the drone hovering above the ground.

D.1.1. Computing Antenna Positions

Before any data processing, the position of each antenna element at each pulse location needs to be computed using the known antenna geometry, and the attitude information provided by the INS. When computing the position of each antenna element, the offset between the antenna array and the GNSS antenna also needs to be accounted for, as there is a significant lever arm when the drone pitches or rolls. To compute the position of each antenna element, the following information is required:

- The attitude offset between the radar and the INS. This applies mostly to the yaw direction and is defined by using a single antenna element as a reference. In this case, the first antenna element is used.
- The position offset between the GNSS antenna and the radar.
- The diameter of each antenna array.
- The offset between the receive and transmit array.

For the receiver array, if it is assumed that the offset in roll and pitch is zero, the position of the n_{th} antenna element in the inertial reference frame is given by

$${}^{\mathcal{N}}P_{RX,n} = {}^{\mathcal{N}}P_{INS} + {}^{\mathcal{N}}\mathbf{R}(\alpha, \beta, \gamma)({}^{\mathcal{G}}P_{RX,c} + {}^{\mathcal{C}}P_{RX,n}), \quad (\text{D.5})$$

where the caligraphic notation determines the current reference frame. \mathcal{C} denotes the reference frame centered at the middle of the antenna array, \mathcal{G} is the reference frame centered at the GNSS antenna, and \mathcal{N} is the reference frame centered at the global origin. the subscript c indicates the center of the antenna array. The rotation matrix applies the correct yaw, pitch and roll represented by α, β and γ respectively. The position of

a single antenna element with respect to the center of the receive array is described by,

$$\phi_{P_{RX,n}} = -\frac{d_{RX}}{2} \frac{n-1}{32} 2\pi. \quad (\text{D.6})$$

Note that there is a minus sign due to the numbering of antenna elements being defined as clockwise in an ENU coordinate system. The same steps can be applied to the transmitter array by adding the offset between the center of the two arrays. An image of the calculated antenna positions for an arbitrary pulse location can be seen in Figure D.3.




Figure D.3: The antenna setup used for the tracks in Figure C.5

D.2. Data Processing

Before the data is processed in the estimation algorithms, several pre-processing steps are applied. The pre-processing steps prepare the data for further processing and improve some aspects of performance. This section will discuss downsampling, range loss correction, windowing and Doppler processing.

D.2.1. Downsampling

At first, it may seem strange to downsample the data, as this removes information. While this is true to some extent, it can also have a positive impact on performance. Recall from Chapter 2 that the maximum observable range depends on the fast-time sampling frequency. In some cases, however, the area of interest may only correspond to a narrow range interval. The processing can be optimized by only computing the Fourier transform for the frequency of interest. This is also known as the zoom- Fast Fourier Transform (FFT) [25]. By downsampling in fast-time, the maximum range is limited. However, the FFT now requires fewer samples to obtain the same range resolution in the area of interest. This can be seen from the equation for the FFT resolution

$$\delta f = \frac{f_s}{\text{NFFT}} \quad (\text{D.7})$$

Where f_s is the sampling rate and NFFT is the number of samples used to compute the FFT. It should be noted that the data is first low-pass filtered before downsampling to remove aliasing.

D.2.2. Range Compression And Loss Compensation

As explained in Chapter 2, Range compression can be performed through a Fourier transform. For experimental data, there is a loss caused by the propagation of waves in a spherical manner. The one-way loss is proportional to $\frac{1}{R^2}$. The reflection also experiences the same loss causing the total effect for a single target to be $\frac{1}{R^4}$. However, since the background reflects most of the energy back, the loss for the entire scene is only proportional to $\frac{1}{2R^2}$, as the scene background can be seen as a wall that doesn't alter the shape of the transmitting wave.




Figure D.4: The average range profile for an experimental data segment 1 second in length. It can be seen that the range loss is approximately quadratic.

D.2.3. Doppler Processing

The Doppler frequency is determined from pulse to pulse using a slow-time Fourier transform. This is applied after the data has been range-compressed. The number of slow-time samples used per pulse is chosen based on an empirical estimate of the acceleration of the radar platform. As the Doppler spectrum provides an average velocity over a certain time interval, it is important to ensure that the total change in velocity during this interval is not too great. On the other hand, enough pulses need to be used such that there is sufficient frequency resolution. Based on experiments with simulated data, it was found that the noise in the velocity estimate was acceptable when at least 150 samples were used. In cases where more digital beams are incorporated, from the perspective of Doppler Centroid Estimation, the effective resolution increases proportionally, allowing for shorter time spans. However, this is only true when the beams are uncorrelated.

Bibliography

- [1] Mti_usermanual.pdf. https://www.xsens.com/hubfs/Downloads/usermanual/MTi_usermanual.pdf. (Accessed on 09/19/2020).
- [2] Ahn. Home, Aug 2020. URL <https://www.ahn.nl/>.
- [3] Joshua N. Ash. An Autofocus Method for Backprojection Imagery in Synthetic Aperture Radar. *IEEE Geoscience and Remote Sensing Letters*, 9(1):104–108, January 2012. ISSN 1558-0571. doi: 10.1109/LGRS.2011.2161456.
- [4] R. Bamler. Doppler Frequency Estimation and the Cramer-Rao Bound. *IEEE Transactions on Geoscience and Remote Sensing*, 29(3):385–390, 1991.
- [5] Dan Barnes and Ingmar Posner. Under the Radar: Learning to Predict Robust Keypoints for Odometry Estimation and Metric Localisation in Radar. *arXiv:2001.10789 [cs]*, February 2020.
- [6] G. Brown. The average impulse response of a rough surface and its applications. *IEEE Transactions on Antennas and Propagation*, 25(1):67–74, 1977. doi: 10.1109/TAP.1977.1141536.
- [7] T.M. Calloway and G.W. Donohoe. Subaperture autofocus for synthetic aperture radar. *IEEE Transactions on Aerospace and Electronic Systems*, 30(2):617–621, April 1994. ISSN 1557-9603. doi: 10.1109/7.272285.
- [8] Sarah H. Cen and Paul Newman. Radar-Only Ego-Motion Estimation in Difficult Settings via Graph Matching. *arXiv:1904.11476 [cs]*, April 2019.
- [9] S. Clark and G. Dissanayake. Simultaneous localisation and map building using millimetre wave radar to extract natural features. In *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No.99CH36288C)*, volume 2, pages 1316–1321, Detroit, MI, USA, 1999. IEEE. ISBN 978-0-7803-5180-6. doi: 10.1109/ROBOT.1999.772543.
- [10] S. Clark and H. Durrant-Whyte. Autonomous land vehicle navigation using millimeter wave radar. In *Proceedings. 1998 IEEE International Conference on Robotics and Automation (Cat. No.98CH36146)*, volume 4, pages 3697–3702 vol.4, 1998.
- [11] I. G. Cumming and Shu Li. Improved Slope Estimation for SAR Doppler Ambiguity Resolution. *IEEE Transactions on Geoscience and Remote Sensing*, 44(3):707–718, 2006.
- [12] Ian G. Cumming and Frank Haychee Wong. *Digital Processing of Synthetic Aperture Radar Data: Algorithms and Implementation*. Artech House, 2005.
- [13] S. Dill, E. Schreiber, M. Engel, A. Heinzl, and M. Peichl. A Drone Carried Multichannel Synthetic Aperture Radar for Advanced Buried Object Detection. In *2019 IEEE Radar Conference (RadarConf)*, pages 1–6, April 2019. doi: 10.1109/RADAR.2019.8835814.
- [14] Michael I. Duersch and David G. Long. *Backprojection Autofocus for Synthetic Aperture Radar*. 2013.
- [15] Michael Israel Duersch. *Backprojection for synthetic aperture radar*. PhD thesis, Brigham Young University, 2013.
- [16] P. Eichel, D. Ghiglia, and C. V. Jakowatz. Speckle processing method for synthetic-aperture-radar phase correction. *Optics letters*, 14 1:1–3, 1989.
- [17] A. Evers, N. Zimmerman, and J. A. Jackson. Semidefinite relaxation autofocus for bistatic backprojection sar. In *2017 IEEE Radar Conference (RadarConf)*, pages 1332–1337, 2017. doi: 10.1109/RADAR.2017.7944412.

- [18] R. Faragher. Understanding the basis of the kalman filter via a simple and intuitive derivation [lecture notes]. *IEEE Signal Processing Magazine*, 29(5):128–132, 2012.
- [19] Walter R. Fried. Principles and Performance Analysis of Doppler Navigation Systems. *IRE Transactions on Aeronautical and Navigational Electronics*, ANE-4(4):176–196, December 1957. ISSN 2331-0804. doi: 10.1109/TANE3.1957.4201552.
- [20] Franck Gerossier, Paul Checchin, Christophe Blanc, Roland Chapuis, and Laurent Trassoudaine. Trajectory-oriented EKF-SLAM using the Fourier-Mellin Transform applied to Microwave Radar Images. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4925–4930, St. Louis, MO, October 2009. IEEE. ISBN 978-1-4244-3803-7. doi: 10.1109/IROS.2009.5354548.
- [21] LeRoy A. Gorham and Linda J. Moore. SAR Image Formation Toolbox for MATLAB. In *Algorithms for Synthetic Aperture Radar Imagery XVII*, volume 7699, page 769906. International Society for Optics and Photonics, April 2010. doi: 10.1117/12.855375.
- [22] Zhenyu Guo, Hongbo Zhang, and Shaohua Ye. Simplified and approximation autofocus back-projection algorithm for SAR. *The Journal of Engineering*, 2019(20):6408–6412, 2019. ISSN 2051-3305. doi: 10.1049/joe.2019.0406.
- [23] Andreas Hantsch and Wolfgang Menzel. A 76 GHz Folded Reflector Antenna for True Ground Speed Measurement. 2006.
- [24] L. Hong. Discrete constant-velocity-equivalent multirate models for target tracking. *Mathematical and Computer Modelling*, 28(11):7 – 18, 1998. ISSN 0895-7177. doi: [https://doi.org/10.1016/S0895-7177\(98\)00161-7](https://doi.org/10.1016/S0895-7177(98)00161-7). URL <http://www.sciencedirect.com/science/article/pii/S0895717798001617>.
- [25] E. Hoyer and R. Stork. The zoom fft using complex modulation. In *ICASSP '77. IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 2, pages 78–81, 1977. doi: 10.1109/ICASSP.1977.1170177.
- [26] M. Jankiraman. *FMCW Radar Design*. 2018.
- [27] M. Y. Jin. Optimal Doppler Centroid Estimation for SAR Data from a Quasi-Homogeneous Source. *IEEE Transactions on Geoscience and Remote Sensing*, GE-24(6):1022–1025, 1986.
- [28] M. S. Kay. *Fundamentals of Statistical Signal Processing. Estimation Theory*. Prentice-Hall, 1993.
- [29] D. V. Khablov. Autonomous Navigation System of Ground Transport Based on Doppler Sensors for Measuring Vector Velocity. *Measurement Techniques*, 61(4):384–389, July 2018. ISSN 1573-8906. doi: 10.1007/s11018-018-1438-x.
- [30] Marc Lavielle. Using penalized contrasts for the change-point problem. *Signal Processing*, 85:1501–1510, 08 2005. doi: 10.1016/j.sigpro.2005.01.012.
- [31] Kuang-Hung Liu, Ami Wiesel, and David C. Munson. Synthetic Aperture Radar Autofocus via Semidefinite Relaxation. *IEEE Transactions on Image Processing*, 22(6):2317–2326, June 2013. ISSN 1941-0042. doi: 10.1109/TIP.2013.2249084.
- [32] T. Long, Z. Lu, Z. Ding, and L. Liu. A DBS Doppler Centroid Estimation Algorithm Based on Entropy Minimization. *IEEE Transactions on Geoscience and Remote Sensing*, 49(10):3703–3712, 2011.
- [33] Zhi-quan Luo, Wing-kin Ma, Anthony Man-cho So, Yinyu Ye, and Shuzhong Zhang. Semidefinite Relaxation of Quadratic Optimization Problems. *IEEE Signal Processing Magazine*, 27(3):20–34, May 2010. ISSN 1558-0792. doi: 10.1109/MSP.2010.936019.
- [34] S. N. Madsen. Estimating the Doppler Centroid of SAR Data. *IEEE Transactions on Aerospace and Electronic Systems*, 25(2):134–140, 1989.
- [35] D. Mao, Y. Zhang, Y. Zhang, Y. Huang, and J. Yang. Doppler Beam Sharpening Using Estimated Doppler Centroid Based on Edge Detection and Fitting. *IEEE Access*, 7:123604–123615, 2019.
- [36] A. Martínez and J. Marchand. Sar image quality assessment. 1993.

- [37] Mostafa Mostafa, Shady Zahran, Adel Moussa, Naser El-Sheimy, and Abu Sesay. Radar and Visual Odometry Integrated System Aided Navigation for UAVS in GNSS Denied Environment. *Sensors*, 18(9):2776, August 2018. ISSN 1424-8220. doi: 10.3390/s18092776.
- [38] C. Oliver and S. Quegan. *Understanding Synthetic Aperture Radar Images*. January 2004.
- [39] M. Otten, N. Maas, R. Bolt, M. Caro-Cuenca, and H. Medenblik. Circular Micro-Sar for Mini-UAV. In *2018 15th European Radar Conference (EuRAD)*, pages 321–324, 2018.
- [40] R. P. Perry, R. C. DiPietro, and R. L. Fante. Coherent integration with range migration using keystone formatting. In *2007 IEEE Radar Conference*, pages 863–868, 2007.
- [41] Lei Ran, Zheng Liu, Tao Zhang, and Tao Li. Autofocus for correcting three dimensional trajectory deviations in synthetic aperture radar. In *2016 CIE International Conference on Radar (RADAR)*, pages 1–4, October 2016. doi: 10.1109/RADAR.2016.8059312.
- [42] R. K. Raney. The delay/doppler radar altimeter. *IEEE Transactions on Geoscience and Remote Sensing*, 36(5):1578–1588, 1998.
- [43] Antonio Scannapieco, Alfredo Renga, Giancarmine Fasano, and A. Moccia. Ultralight radar for small and micro-uav navigation. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-2/W6:333–338, 08 2017. doi: 10.5194/isprs-archives-XLII-2-W6-333-2017.
- [44] M. Schartel, R. Burr, W. Mayer, N. Docci, and C. Waldschmidt. UAV-Based Ground Penetrating Synthetic Aperture Radar. In *2018 IEEE MTT-S International Conference on Microwaves for Intelligent Mobility (ICMIM)*, pages 1–4, April 2018. doi: 10.1109/ICMIM.2018.8443503.
- [45] M. Schartel, R. Burr, W. Mayer, and C. Waldschmidt. Airborne Tripwire Detection Using a Synthetic Aperture Radar. *IEEE Geoscience and Remote Sensing Letters*, 17(2):262–266, February 2020. ISSN 1558-0571. doi: 10.1109/LGRS.2019.2917917.
- [46] Timothy J. Schulz. Optimal Sharpness Function for SAR Autofocus. *IEEE Signal Processing Letters*, 14(1): 27–30, January 2007. ISSN 1558-2361. doi: 10.1109/LSP.2006.881525.
- [47] Xiaoyi Shen, Markku Similä, Wolfgang Dierking, Xi Zhang, Changqing Ke, Meijie Liu, and Manman Wang. A new retracking algorithm for retrieving sea ice freeboard from cryosat-2 radar altimeter data during winter–spring transition. *Remote Sensing*, 11(10):1194, May 2019. ISSN 2072-4292. doi: 10.3390/rs11101194. URL <http://dx.doi.org/10.3390/rs11101194>.
- [48] Aron Sommer and Jörn Ostermann. Image Reconstruction without an Inertial Navigation System using Backprojection Autofocus for Synthetic Aperture Radar. In *2019 20th International Radar Symposium (IRS)*, pages 1–10, June 2019. doi: 10.23919/IRS.2019.8767455.
- [49] D. E. Wahl, P. H. Eichel, D. C. Ghiglia, and C. V. Jakowatz. Phase Gradient Autofocus—a Robust Tool for High Resolution SAR Phase Correction. *IEEE Transactions on Aerospace and Electronic Systems*, 30(3): 827–835, 1994.
- [50] Hoi-To Wai, Wing-Kin Ma, and Anthony Man-Cho So. Cheap semidefinite relaxation MIMO detection using row-by-row block coordinate descent. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3256–3259, May 2011. doi: 10.1109/ICASSP.2011.5946716.
- [51] Y. Wang, D. Kostić, S. T. H. Jansen, and H. Nijmeijer. Filling the gap between low frequency measurements with their estimates. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 175–180, 2014. doi: 10.1109/ICRA.2014.6906606.
- [52] Weidong Yu and Zhaoda Zhu. Comparison of Doppler Centroid Estimation Methods in SAR. In *Proceedings of the IEEE 1997 National Aerospace and Electronics Conference. NAECON 1997*, volume 2, pages 1015–1018 vol.2, 1997.
- [53] Zaiwen Wen, Donald Goldfarb, Shiqian Ma, and Katya Scheinberg. ROW BY ROW METHODS FOR SEMIDEFINITE PROGRAMMING. page 21.

-
- [54] C. Xu, L. Daniel, E. Hoare, V. Sizov, and M. Cherniakov. Comparison of Speed over Ground Estimation Using Acoustic and Radar Doppler Sensors. In *2014 11th European Radar Conference*, pages 189–192, 2014.
- [55] A. F. Yegulalp. Fast backprojection algorithm for synthetic aperture radar. In *Proceedings of the 1999 IEEE Radar Conference. Radar into the Next Millennium (Cat. No.99CH36249)*, pages 60–65, 1999.
- [56] Young-Kyun Kong, Byung-Lae Cho, and Young-Soo Kim. Ambiguity-Free Doppler Centroid Estimation Technique for Airborne SAR Using the Radon Transform. *IEEE Transactions on Geoscience and Remote Sensing*, 43(4):715–721, April 2005. ISSN 0196-2892. doi: 10.1109/TGRS.2005.843955.