Radar Based Navigation in Unknown Terrain

DISSERTATION

Kyle J. Kauffman, Civilian, USAF

AFIT/ENG/DS/12-03

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

AFIT/ENG/DS/12-03

Radar Based Navigation in Unknown Terrain

DISSERTATION

Presented to the Faculty

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

In Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

Kyle J. Kauffman, B.S.C.E., M.S.C.S.

Civilian, USAF

December, 2012

AFIT/ENG/DS/12-03

Radar Based Navigation in Unknown Terrain

Kyle J. Kauffman, B.S.C.E., M.S.C.S.
Civilian, USAF

Approved:

_____   _____
Dr. J. Raquet, PhD (Chairman)   30 NOV 2012
                                                                        Date

_____   _____
Dr. Y. Morton, PhD (Member)   Nov. 21, 2012
                                                                        Date

_____   _____
Dr. D. Garmatyuk, PhD (Member)   Nov. 21, 2012
                                                                        Date

_____   _____
Lt Col K. Fisher, PhD (Member)   21 Nov 2012
                                                                        Date

Accepted:

_____   _____
M. U. Thomas   12 Dec 2012
Dean, Graduate School of Engineering   Date
   and Management

AFIT/ENG/DS/12-03

Abstract

There is a great need to develop non-GPS based methods for positioning and navigation in situations where GPS is not available. This research focuses on the development of an Ultra-Wideband Orthogonal Frequency Division Multiplexed (UWB-OFDM) radar as a navigation sensor in GPS-denied environments. A side-looking vehicle-fixed UWB-OFDM radar is mounted to a ground or aerial vehicle continuously collecting data. A set of signal processing algorithms and methods are developed which use the raw radar data to aide in calculating the vehicle position and velocity via a simultaneous localization and mapping (SLAM) approach. The radar processing algorithms detect strong, persistent, and stationary reflectors embedded in the environment and extract range/Doppler measurements to them. If the radar is the only sensor available, the measurements are used to directly compute the vehicle position via an extended Kalman filter (EKF) or Levenberg-Marquardt solver. If an existing navigation platform is available, the measurements are combined with the other sensors in an EKF. The developed algorithms are tested via both a series of airborne simulations and a ground-based experiment. The airborne simulations are performed with simulated commercial-grade, tactical-grade, and navigation-grade INS systems available. The experiment is performed with an indoor mobile platform containing an HG1700 tactical-grade INS and an X-band 500MHz UWB-OFDM radar. For all configurations, the computed navigation solution performance is analyzed with the following sensor availability: radar-only, INS-only, and combined radar/INS. In both simulation and experimental scenarios, the integrated INS/UWB-OFDM system shows significant improvements over an INS-only navigation solution. The radar-only system performs well assuming high availability of reflectors to track, with decreased performance when reflector-less environments are encountered.

Table of Contents

v

List of Figures

List of Tables

List of Abbreviations

Radar Based Navigation in Unknown Terrain

I. Introduction

The problem of providing a robust, reliable navigation solution to aerial and ground vehicles is of great interest. The safety and success of commercial transportation systems, remotely operated vehicles, guided projectiles, maritime operations, and many other applications relies on real-time knowledge of the vehicle's position, velocity, and orientation. The development of navigation sensors to provide these real-time estimates is the focus of a large field of research.

Rudimentary navigation has been performed since the beginning of recorded history. One of the earliest known navigation methods is celestial navigation, whereby careful observations of the stars gives the observer information about their position. By positioning fixed stars with a sextant and using a chronometer to estimate the time of day early sailors could calculate their latitude and longitude [1]. Alternatively, when sailors were near land they could use knowledge of known landmarks on the shore to estimate their position. If neither of these sources of absolute position information were available they could still estimate their location by carefully recording their heading and speed, which told them how far they had moved and in what direction. By continuing to record their movement over time they could create a track of their travels since their last absolute position update, a process known as dead-reckoning navigation.

While there are many dead-reckoning measurement devices (chip logs, airspeed indicators, etc.), modern navigation solutions typically rely on inertial navigation systems (INS). INSs use a combination of accelerometers and gyroscopes to estimate the linear and angular accelerations of the navigation platform, which can be used to calculate the change in position of a platform or vehicle much like heading and speed. Gyro-based navigation was initially explored in the early 1900s [2], with increasing

interest due to the development of the V1 and V2 guided rockets during World War II [1]. However, like all dead-reckoning approaches, INS-based navigation only works well over short periods of time; the utilization of relative position information results in position drift, a type of error which builds up over time and eventually causes the navigation solution to be useless.

The launch of the first Global Positioning System (GPS) satellite in 1978 [3] provided a highly available source of absolute position information, allowing vehicles to calculate accurate non-drifting navigation solutions continuously. Unfortunately, the long distances between terrestrial vehicles and the GPS satellites combined with power constraints on the satellites leaves GPS vulnerable to both intentional and unintentional jamming [4]. GPS also cannot be used in indoor or underground environments, where the satellite signals are too weak to be received. When GPS is denied or unavailable, the vehicle must find other means of improving the INS's position drift.

One approach is to integrate INS with other sensors, resulting in a combined sensor navigation platform. A number of alternative sensors have been used for navigation, including cameras [5,6], AM radio [7], sonar [8], television [9], GSM [10], and LIDAR [11]. Each of these sensors has advantages as well as phenomena which limit their operation. For example, optical sensors are easily obstructed by cloud cover or smoke. The addition of alternative sensors allows the navigation platform to operate under more diverse environmental conditions. Each additional sensor also provides a source of independent information about the vehicles' navigation solution, resulting in a lower error bound. Thus the combination of a wide variety of sensors results in a better and more robust navigation solution.

1.1 Radar-based navigation

This dissertation proposes the use of radar as a navigation sensor. Radars are a form of echo-location technology originally developed during World War II to detect and locate objects. Radars operate by transmitting radio signals anywhere between 2MHz and 300GHz [12] and recording returns of the signal from reflecting

objects. The information gathered about a remote object includes its range from the radar, its radial velocity, its angular direction with respect to the radar, and its size and shape. There are many modern applications of radar, including air traffic control, surveillance, and military operations; one application of particular interest is the development of imaging radar. Imaging radars, also known as synthetic aperture radars (SAR), use a series of transmissions to construct high resolution images of an environment [12].

In theory, the ability to use radars to generate images of a scene allows usage of a radar as a navigation sensor on-board an integrated navigation sensor array. One simple integration method would be to use SAR to construct imagery and then utilize existing image-based navigation approaches [13]. However, this approach has several issues:

1. The construction of SAR imagery in real-time is computationally expensive, often being performed on parallelized clusters of computers. This prohibits direct SAR image-based navigation on platforms which have limited resources, such as small unmanned aerial vehicles (UAVs) or unmanned ground vehicles (UGVs).

2. Since high-precision navigation is desired, an ultra-wideband (UWB) radar is desired, because wide bandwidth signals have higher range resolution. UWB data processing requires even more computation to process than narrow-band signals, increasing the need to develop high speed processing algorithms. [14,15]

3. The phenomenology encountered by optical sensors is different from those operating in radar bandwidths. For example, cloud cover will obscure land features from a camera mounted on a UAV, but a low-frequency radar will penetrate the cloud cover and the earth's surface [15].

4. The features which can be utilized by SAR will have very different characteristics than those of a camera, and necessitate a different approach to feature detection, extraction, tracking, and positioning. A radar collects down-range and Doppler

3

data from targets, whereas a camera is an angular sensor which (inherently) has no knowledge of the depth of a scene.

The use of radar as an integrated navigation sensor will require first overcoming these obstacles by developing a radar signal processing approach capable of real-time SAR feature extracting and tracking.

1.2 Problem Statement

The problem considered is that of navigating a UAV/UGV through terrain/buildings under the following conditions:

- The terrain/building is unknown, such that no *a priori* map of the environment is available.

- GPS is not available, either due to insufficient signal strength or intentional/unintentional jamming. The potential presence of a jamming source necessitates the use of RF waveforms with resilience to interference.

- Reflectors are available in the environment which are strong, persistent (over short durations of time), stationary, and isolated from each other.

- The radar clutter is reasonably modeled using statistical clutter models [12].

- An INS and a UWB SAR system are available on the vehicle.

The goal of this research is to use a UWB SAR sensor and an INS as an integrated navigation platform. As the vehicle moves through its environment, it combines information obtained from the UWB radar and INS to compute the navigation solution (position, orientation, velocity) of the vehicle.

1.3 Research Contributions

The problem of integrating data collected from an imaging radar into a navigation solution requires careful consideration of the requirements of a navigation sensor

and the use of computationally inexpensive approximations to optimal processing methods. The primary research contributions of this dissertation are as follows:

1. The development of a set of SAR data processing algorithms enabling the real-time extraction of SAR observations for the purpose of navigation.

2. The development of a navigation filter capable of using the information contained in the SAR observations to correct INS errors.

3. The development of an experimental UGV system prototype which contains both an INS and an experimental jamming-resistant UWB radar system. The radar waveform transmitted is an orthogonal frequency division multiplexed (OFDM) symbol, which allows for resilience to jamming [16].

4. Experimental validation of the algorithms/navigation filter using the system prototype.

1.4 *Dissertation Outline*

The remainder of this dissertation is organized as follows. Section II provides a mathematical background for general navigation methods, radar processing techniques, and inertial systems. Section III describes the novel data processing/estimation algorithms for integrating SAR data into an INS navigation solution, as well as the developed experimental system prototype. Section IV presents simulated results for UAV-based navigation. Section V provides experimental results for UGV-based navigation using the experimental system prototype. Finally, Section VI details the conclusions of the study and potential future work.

II. Mathematical Background

In this chapter, we review previously developed concepts and techniques which aid and inform our algorithm and model design in the next chapter.

2.1 Conventions

2.1.1 Notation.   The following notational conventions are used in this dissertation:

- A hat $(\hat{\cdot})$ indicates an estimated quantity.

- A single dot $(\dot{\cdot})$ and double dots $(\ddot{\cdot})$ on top of a symbol indicate a first-order and second-order derivative with respect to time.

- A ∂ is used to represent the derivative of a quantity with respect to another quantity, e.g. $\frac{\partial^2 y(x)}{\partial^2 f(x)}$

- When indexing into a vector/matrix, $[\cdot]$ brackets are used. For example, $[\mathbf{a}_k]_l$ is the lth row of the vector $\mathbf{a}_k$.

- A δ represents the error of a quantity (i.e. $\delta(\cdot)$ is the error of $(\cdot)$).

- The reference frame of a quantity (if given) is denoted by a superscript. For example, $(\cdot)^b$ represents the parenthesized quantity in the b frame.

- Bold lower case variables are row or column vectors.

- Bold upper case variables are matrices.

2.1.2 Reference Frames.   Geometric vector quantities must be expressed with respect to a frame of reference. A frame of reference consists of two things: an origin point and an ordered basis of orthogonal unit-length vectors. For an arbitrary 3-dimensional frame f, we will denote the basis vectors $\mathbf{x}^f$, $\mathbf{y}^f$, and $\mathbf{z}^f$. In this section we will outline some common reference frames, as detailed in [1, 13].

In general, an inertial frame is defined as a non-rotating frame without non-gravitational acceleration acting upon it. The earth-centered inertial (ECI) frame i

is defined by an origin at the center of mass of the earth and the $\mathbf{x}^i,\mathbf{y}^i,\mathbf{z}^i$ vectors pointing towards observable stars. The assumption is that the relative motion of the stars with respect to the earth has a negligible effect on the basis set when using the frame to position local objects. Thus the i frame does not rotate with respect to the earth, sun, or solar system, but instead instead maintains reference with the galaxy. With this in mind, it may not be a truely "inertial" frame, as the motion of the fixed stars will cause the ECI frame to have an extremely small rotation.

The earth-centered earth-fixed (ECEF) frame e is defined by an origin at the center of mass of the earth and the $\mathbf{x}^e,\mathbf{y}^e,\mathbf{z}^e$ vectors pointing towards the Greenwich meridian, the equatorial plane at 90 degrees east longitude, and the north pole respectively. This frame's basis vectors are aligned with fixed points on the earth, and thus it is a rotating frame. However, it has the advantage that objects on the surface of the earth (or otherwise synchronized with the earth's rotation, such as geosynchronous orbits) do not change position in the ECEF frame.

The local level tangential navigation frame n' is defined by an origin at the center of mass of the navigation vehicle and $\mathbf{x}^{n'},\mathbf{y}^{n'},\mathbf{z}^{n'}$ vectors pointing North, East, and down (towards the earth's center of gravity) respectively (NED). This frame does not change with the rotation of the vehicle, but its origin is continually changing with respect to the earth and the ECEF frame. This frame is convenient when quantities are easily expressed in relation to the vehicle position, such as ranges to targets.

The fixed local level navigation frame n is defined by an origin at the initial navigation vehicle's position and $\mathbf{x}^n,\mathbf{y}^n,\mathbf{z}^n$ vectors pointing North, East, and down respectively. The n frame differs from the n' frame in that the origin does not track the vehicle position over time. This frame is useful for modeling navigation problems where the earth's surface can be approximated as flat, such as indoor environments.

The body frame b is defined by an origin fixed to the vehicle and $\mathbf{x}^b,\mathbf{y}^b,\mathbf{z}^b$ vectors pointing out the nose, right wing, and bottom of the vehicle respectively. This frame is similar to the n' frame except that it rotates with the vehicle. It is useful when

measurements or other quantities are only known with respect to the body of the vehicle, such as a strapdown inertial device.

2.2 *Asymptotic Complexity*

Estimation error is not the only metric of optimality applicable to estimators. In many contexts, the speed at which the algorithm calculates its estimate is also important. To analyze this problem, we must introduce metrics which measure the computational burden of the algorithms generating estimates.

Computational burden is a particularly difficult optimality criteria to quantify. Modern superscalar processors feature instruction level parallelism, out of order instruction execution, data pipelining, multi-layer caching structions, data dependency elimination, and many more optimizations. The existence of such complex computation engines makes simple algorithm metrics such as counting the number of floating point adds and multiplies virtually useless, as the number of floating operations per second will vary from algorithm to algorithm [17]. In order to optimize real-world algorithms, brute-force approaches which try all available execution strategies are needed, such as the FFTW implementation of the fast Fourier transform (FFT) in [17]. The results of such optimization will vary wildly depending on the specific processor architecture and platform on which the algorithm is run.

One alternative metric for computational burden is to consider how an algorithm will scale as data inputs get larger. Because computer speed is ever increasing, this approach allows the future applicability of the algorithm. Donald Knuth popularized the "Big-O" notation [18] for this metric, which is defined here. Let n be the size of the input data to an algorithm, $f(n)$ be the number of operations performed by the algorithm for an input of size n, and $g(n)$ be a function describing the "order" of the algorithm. Then the notation

$$f(n) = O(g(n)) \tag{1}$$

implies that $f(n)$ is bounded above by the function $g(n)$ as n gets large. Formally, we say there exists positive integers n_0, C such that

$$|f(n)| \leq C|g(n)|, \forall n > n_0 \tag{2}$$

Similarly, the notation

$$f(n) = \Omega(g(n)) \tag{3}$$

implies that the function $f(n)$ is bounded below by the function $g(n)$ as n gets large. Formally, we say there exists positive integers n_0, C such that

$$|f(n)| \geq C|g(n)|, \forall n > n_0 \tag{4}$$

Finally, the notation

$$f(n) = \Theta(g(n)) \tag{5}$$

implies that the function f is bounded above and below by the function g as n gets large, and is thus both $O(g(n))$ and $\Omega(g(n))$. This metric can be used to classify algorithms according to their asymptotic complexity. For example, Fourier transform implementation which is $\Theta(n^2)$ will run much slower than an implementation which is $\Theta(n \log n)$ for large n. Indeed, the ability to run in $\Theta(n \log n)$ time is the distinguishing factor that separates the FFT from other discrete Fourier transform (DFT) implementations.

2.3 Geometric Dilution of Precision

When using ranges between a set of sensors and a target to estimate the target's position, it is important to consider the orientation of the sensors with respect to the target. If the ranges contain errors, then the geometry of the sensor configuration will

result in geometric dilution of precision (GDOP). Fig. 1 illustrates this phenomenon. In A), the two small filled circles are sensors which have each collected a range to the target from themselves. Assuming the ranges to be error-less, we can state that the target must lie somewhere along a circle with radius equal to the range, as shown by the circular lines. Since the target must lie on *both* of the lines, there is only one possible target location: the intersection of the two lines. There is another intersection between the lines (not shown), but we can ignore this possibility if we know the target lies to the right of the sensors. In B) we consider the same scenario but let the ranges have errors. The target is no longer known to lie on the lines, but instead in an area plus or minus a small distance from the lines. Similarly the target does not have to lie at the intersection point, but instead could be anywhere in the green shaded area. In C), we consider the same scenario as B) but with bad geometry (both sensors give us ranges along the same axis). We see that the green shaded area where the target could reside is much larger. Thus the geometry in B) is superior to that in C) if we want to know the position of the target in both the x and y axes.

When we consider multiple sensors and/or multiple ranges collected over time, the GDOP becomes harder to visualize. However, in general having sensors geometrically diverse (i.e., distributed and not clumped) yields higher accuracy estimates.

2.4 Estimation

2.4.1 Introduction. The goal of estimation is to predict the value of one or more hidden quantities given a set of observations which contain information about them. There are two primary models used in estimation: classical estimation, where the hidden quantities are assumed to be deterministic but unknown, and Bayesian estimation, where the hidden quantities are assumed to be random. In both models, one must first decide on an optimality criterion which defines what an "optimal" estimator is. The second step is to derive an estimator that is optimal with respect to the chosen criterion. In this section, some common classical and Bayesian estimators and metrics of optimality are discussed.
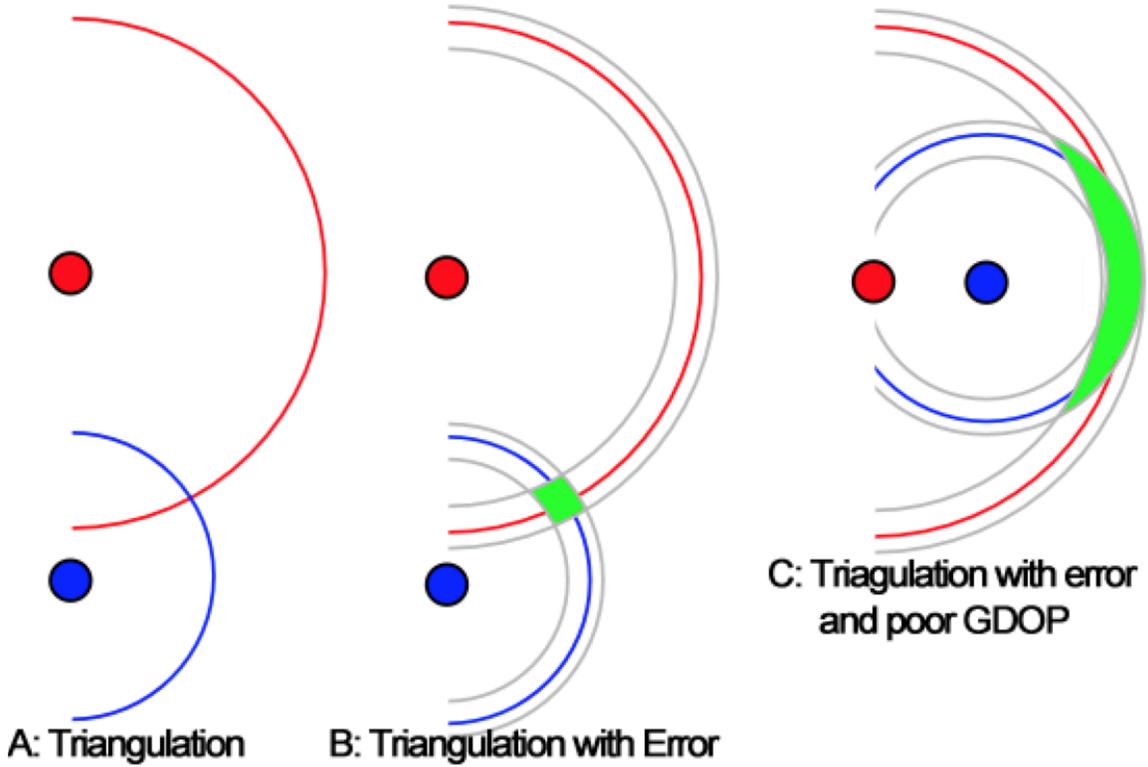
Figure 1: Illustration of GDOP for range-based localization, taken from Wikimedia Commons (attribution: Colin O'Flynn). A) two sensors receiving ranges to a target, with one solution. B) same as A but showing errors on ranges, yielding an area of possible solutions. C) same as B but with poor geometric distribution of sensors, yielding a larger area of possible solution (higher error).

2.4.2 Classical Estimation. Classic estimation calculates an estimate of a hidden deterministic set of quantities given a set of observations. Let $\mathbf{x}$ be a vector of unknown quantities (possibly of size 1) we wish to estimate and $\mathbf{z}$ be a set of observations which contains information about $\mathbf{x}$. Then we want to come up with an estimator $\hat{\mathbf{x}}(\mathbf{z})$ such that $\hat{\mathbf{x}}(\mathbf{z}) \approx \mathbf{x}$. $\hat{\mathbf{x}}$ is technically a function mapping a random variable (the observations) to a random variable (the estimator). That is,

$$\hat{\mathbf{x}} : (\mathbf{\Omega} \to \mathbb{R}) \to (\mathbf{\Omega} \to \mathbb{R}) \tag{6}$$

where $\mathbf{\Omega}$ is the sample space. Note that in practical use, the estimator is applied to a realization of the observations (our sensor output) to produce a realization of

11

the estimator. For brevity, we will omit the function notation from now on and refer simply to $\hat{\mathbf{x}}$ when we talk about the estimator. Thus you can think of $\hat{\mathbf{x}}$ as "an estimate of the quantity $\mathbf{x}$ which uses the observations". This estimate is in fact a RV, and so it has moments such as mean and variance.

There are many metrics for evaluating the performance of an estimator $\hat{\mathbf{x}}$. One popular optimality criterion is to find an estimator which is unbiased (i.e. $E[\hat{\mathbf{x}}] = \mathbf{x}$) and has minimal variance (i.e., $Var(\hat{\mathbf{x}}) < Var(\hat{\mathbf{x}}')$ for all other possible estimators $\hat{\mathbf{x}}'$). Such an estimator is called the minimum variance unbiased (MVU) estimator. MVU estimators are, in general, difficult to find. However, they do exist for particular problems, such as the linear estimation problem considered in the next section.

2.4.3 Linear MVU Estimator. Consider the problem of estimating $\mathbf{x}$ given that the observations $\mathbf{z}$ are related to $\mathbf{x}$ via

$$\mathbf{z} = \mathbf{H}\mathbf{x} + \mathbf{w} \tag{7}$$

where $\mathbf{H}$ is a matrix and $\mathbf{w}$ is a vector of random variables acting as a noise source for the observations. If the noise sources in $\mathbf{w}$ are assumed to be Gaussian with zero mean and covariance matrix $\sigma^2\mathbf{I}$, then the MVU estimator for $\mathbf{x}$ is [19]

$$\hat{\mathbf{x}} = (\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T\mathbf{z} \tag{8}$$

The covariance of the estimate $\hat{\mathbf{x}}$ is

$$E[(\hat{\mathbf{x}} - \mu_{\hat{\mathbf{x}}})(\hat{\mathbf{x}} - \mu_{\hat{\mathbf{x}}})^T]] = \sigma^2(\mathbf{H}^T\mathbf{H})^{-1} \tag{9}$$

This result relies on the system being linear and the noise sources being uncorrelated Gaussian with equal noise strength. While the MVU may be derivable for other scenarios on a case-by-case basis, no generalized form exists for an MVU estimator.

2.4.4 Alternative Metrics and Sub-optimal Estimators. MVU estimators do not exist or are unsolvable for many problems of interest. For these problems we must either use sub-optimal estimators or use a different optimality criterion.

One sub-optimal MVU estimator which is easy to find and can be applied to many problems is the maximum likelihood estimator (MLE). Although no optimality can be claimed for the MLE, it can be proven that the MLE converges to the MVU estimator as the number of observations goes to infinity (i.e. the dimensionality of $\mathbf{z} \to \infty$) [19]. Thus the MLE is considered approximately optimal (in the MVU sense), especially if the number of observations is large.

One popular alternative optimality condition is least squares error (LSE). Unlike MVU or MLE approaches, LSE estimators do not consider the probabilistic properties of the observations, and thus can make no claims about optimality or variance. Instead, only the deterministic model of how the observations relate to the unknown quantities is used. LSE estimators attempt to estimate the hidden quantities such that the observations predicted by the estimate (the prediction is made by passing the estimate through the model) closely match the actual received observations. Formally, we wish to find an $\hat{\mathbf{x}}$ such that the summation

$$\sum_{i=1}^{N}([\mathbf{h}(\hat{\mathbf{x}})]_i - [\mathbf{z}]_i)^2 \tag{10}$$

is minimized. N is the length of $\mathbf{z}$, $\mathbf{z}_i$ is the ith element of $\mathbf{z}$, $[\mathbf{h}(\hat{\mathbf{x}})]_i$ is the ith element of $\mathbf{h}(\hat{\mathbf{x}})$, and $\mathbf{h}$ is a deterministic function which relates the unknown quantity $\mathbf{x}$ to the observation $\mathbf{z}$.

2.4.5 Linear Least Squares. If we constrain the function $\mathbf{h}$ from the last section to be linear

$$\mathbf{z} = \mathbf{Hx} \tag{11}$$

then the LSE criterion is to find a $\hat{\mathbf{x}}$ such that we minimize the quantity

$$(\mathbf{z} - \mathbf{H}\hat{\mathbf{x}})^T(\mathbf{z} - \mathbf{H}\hat{\mathbf{x}}) \tag{12}$$

The estimator which minimizes this quantity is given by [19]

$$\hat{\mathbf{x}} = (\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T\mathbf{z} \tag{13}$$

Note that this is in fact the same result as the MVU estimator for additive Gaussian noise sources. If we made the same assumptions about the noise sources here, the LSE estimator is in fact the optimal MVU estimator. This result demonstrates that, although LSE is not generally optimal in any sense, it tends to yield estimators which are reasonable. If the noise sources were in fact not zero-mean additive Gaussian, then we cannot make any guarantees about the optimality of the LSE estimate; as a rule of thumb, we can say the results it produces tend to be sensible, and apply to a wide range of problems for which the MVU estimator is not available, such as the non-linear problems considered in the next section.

2.4.6 Non-linear Least Squares. Non-linear problems where

$$\mathbf{z} = \mathbf{h}(\mathbf{x}) \tag{14}$$

are in general not able to be estimated by the MVU method. In this case, we can use a non-linear form of the LSE estimator. In particular, we will try to find an estimate $\hat{\mathbf{x}}$ such that

$$(\mathbf{z} - \mathbf{h}(\hat{\mathbf{x}}))^T(\mathbf{z} - \mathbf{h}(\hat{\mathbf{x}})) \tag{15}$$

14

Although no optimality can be claimed, we expect a reasonable estimate of the unknown quantities, since the produced estimate combined with the observation model will closely predict the observations we collected.

There are many approaches to non-linear LSE minimization. Two standard approaches are the Gauss-Newton (GN) method and the gradient descent method. GN starts with an intial guess and then uses a Taylor series expansion to approximate the non-linear function. The minimum of the approximated function is then found via standard root-finding methods. This process is repeated iteratively until GN converges to a stable point [19]. Gradient descent also starts with an initial guess, and steps in the direction of the steepest downward descent (i.e., the gradient of the function at the current location). The gradient descent method keeps stepping in the direction of the gradient vector until it reaches a stable point.

Both of these methods have trade-offs. Gradient descent is highly susceptible to finding local minimums, and GN is more susceptible to convergence issues. The widely used Levenberg-Marquardt algorithm (LMA) is an alternative non-linear least squares estimator which combines GN and gradient descent [19]. LMA has a tunable parameter λ which allows it to act as GN (slow convergence) or gradient descent (fast convergence), or anywhere in-between. LMA allows the user to tune the convergence rate to best suit a particular problem and is a generalization of both GN and gradient descent.

It should be noted that none of these methods can guarantee to find the point which minimizes the LSE. All methods are susceptible to finding local minimums and convergence issues. Even if they could guarantee to find the LSE point, as previously discussed the LSE estimator is not guaranteed to be optimal. Thus non-linear least squares estimates must be considered to be potentially incorrect with no guarantees. The particular problem must be studied and tested to see if non-linear least squares is a suitable approximation.

2.4.7 Bayesian Estimation. Bayesian estimators make the assumption that the probability distribution of the quantity to be estimated is known. This is in contrast to classical estimation, where the quantity is deterministic but unknown. This allows us to compute the probability distributions of estimators. Let R be a random variable, and $\hat{R}$ be the estimate produced by our estimation algorithm. One popular metric is the quantity

$$M_e = E[(\hat{R} - R)^2] \tag{16}$$

where E is the expected value operator. This metric is called the mean squared error (MSE). An estimator which minimizes this error is called the minimum MSE (MMSE) estimator. In general, the MMSE estimator is difficult to find, but it can be found in certain contexts, such as linear problems.

2.4.8 Bayesian Estimation of Linear Dynamic Systems. A linear dynamic system assumes that each state is a linear function of the previous state with additive Gaussian noise:

$$\mathbf{x}_{k+1} = \mathbf{\Phi}_k \mathbf{x}_k + \mathbf{w}_k \tag{17}$$

where $\mathbf{x}_k$ is the state vector at time k, $\mathbf{\Phi}_k$ is the a matrix called the discrete-time dynamics matrix, and $\mathbf{w}_k$ is a vector of jointly-Gaussian random variables at time k with covariance matrix $\mathbf{Q}_k$. Note that this model implies that each state is only a function of the previous state, making it a Markov process. We further assume that the observations we have of this system are linearly related to the states:

$$\mathbf{z}_k = \mathbf{H}_k \mathbf{x}_k + \mathbf{v}_k \tag{18}$$

where $\mathbf{H}_k$ is a matrix relating the observations to the states at time k, and $\mathbf{v}_k$ is a vector of jointly-Gaussian random variables at time k with covariance matrix $\mathbf{R}_k$.

We wish to calculate estimates $\{\hat{\mathbf{x}}_k, \hat{\mathbf{x}}_{k+1}, \ldots\}$ of the state matrix over time given the observations.

The solution to this problem is called the Kalman filter (KF). The KF assumes we have an initial estimate of the state at some time k, and then iteratively calculates the state estimate from that time forward. The first step is to calculate the uninformed estimate of the state at time $k + 1$ (i.e. we haven't used the information from the observation at $k + 1$ to inform the filter yet). Then we have [20]

$$\hat{\mathbf{x}}_{k+1}^- = E[\boldsymbol{\Phi}_k \mathbf{x}_k + \mathbf{w}_k] = \boldsymbol{\Phi}_\mathbf{k} \hat{\mathbf{x}}_k \tag{19}$$

This step is relatively straightforward as we simply use the dynamics equation to propagate our estimate forward. Because our noise sources are zero mean they do not affect our mean estimate. We then inform the estimate at $k + 1$ with the observations we've gathered:

$$\hat{\mathbf{x}}_{k+1}^+ = \hat{\mathbf{x}}_{k+1}^- + \mathbf{K}_{k+1}(\mathbf{z}_{k+1} - \mathbf{H}_k \hat{\mathbf{x}}_{k+1}^-) \tag{20}$$

That is, our informed estimate is equal to our uninformed estimate plus the difference between our measurements and the measurements predicted by our uninformed estimate. This difference is known as the residual of the measurements. The $\mathbf{K}$ matrix is the Kalman gain, which scales the influence of the residual such that the KF is the optimal MMSE estimator. The gain is given by [20]

$$\mathbf{K}_k = \mathbf{P}_k^- \mathbf{H}_k^T (\mathbf{H}_k \mathbf{P}_k^- \mathbf{H}_k^T + \mathbf{R}_k)^{-1} \tag{21}$$

where $\mathbf{R}_k$ is the covariance of the measurements

$$\mathbf{R}_k = E[\mathbf{z}_k \mathbf{z}_k^T] \tag{22}$$

and $\mathbf{P}_k^-$ is the covariance matrix of the state estimate $\hat{\mathbf{x}}_{k+1}^-$ given by

$$\mathbf{P}_{k+1}^- = \mathbf{\Phi}_k \mathbf{P}_k^+ \mathbf{\Phi}_k^T + \mathbf{Q}_k \tag{23}$$

where $\mathbf{Q}_k$ is the covariance of the measurements

$$\mathbf{Q}_k = E[\mathbf{x}_k \mathbf{x}_k^T] \tag{24}$$

and

$$\mathbf{P}_k^+ = (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k) \mathbf{P}_k^- \tag{25}$$

These equations can be used iteratively to estimate the state of the system over time. As measurements come in, the state estimate (and its associated covariance) is propagated to the correct time and then updated with the observations. If no observations are made at that time, then $\hat{\mathbf{x}}_{k+1}^+ = \hat{\mathbf{x}}_{k+1}^-$. This process can proceed indefinitely.

2.4.9 Non-linear Bayesian MMSE Estimator.   The extended KF (EKF) is an extension to the KF to allow it to apply to non-linear problems. The EKF approximates the non-linear system as a linear system in order to apply the standard Kalman gain. This approximation is done by taking a first-order Taylor series expansion on the non-linear equations in the system. Depending on the source of the non-linearity, this may be either the measurement model or the dynamics model. If the measurement model is defined by a matrix H in the KF, the EKF would define a matrix H which is the Jacobian of the non-linear observation function h. Similarly, a non-linear dynamics model F would be replaced by the Jacobian of the non-linear dynamics function f. Note that in calculating the residual, the non-linear function h is still used in lieu of the Jacobian, since it is not necessary to be linear and the usage of the true observation function is more accurate.

18

The EKF model is only optimal in the MMSE sense if the errors introduced by linearization errors are zero. In practice, the linearization error can be significant, especially for highly non-linear systems and when initial state estimates are inaccurate. Additionally, the larger the state error grows in an EKF, the more significant the linearization errors are [20]. This is due to the linearization only being valid in the neighborhood of the truth, as the function is really non-linear. Thus linearization error must be considered when applying the EKF to a real problem. Although no guarantees can be made about the optimality of an EKF with non-zero linearization errors, in practice the solution is very close to the optimal solution for small errors.

2.4.10 Particle Filter.   The particle filter (PF) is a brute-force attempt to provide a better estimator than the EKF for non-linear recursive Bayesian problems. A PF initializes a large number of weighted particles which "sample" the pdf of a state at a given time. The particles represent a discrete approximation to the pdf, and thus allow highly non-Gaussian and arbitrary distributions of the state estimate. The particles are propagated through the non-linear dynamics by individually passing each particle through the dynamics function. There are many methods of updating particle weights when an observation is available. The most common approach is to lower the weights of the particles which are unlikely given the observation realization and the observation model. Since every state and probability is represented as a discrete approximation, the PF allows highly non-linear models and highly non-Gaussian noise distributions. The drawback to the PF is the sheer computational burden required to minimize the discrete approximation error. This computational burdern makes PFs undesireable for applications unless the problem has shown itself to be too non-linear or non-Gaussian for the EKF to produce reasonable results.

2.5 Detection

Detection is the problem of deciding whether or not an object is present given a set of observations (i.e., noisy data). There are two possible types of error: the first

type is known as a "false alarm", which is when we decide that the object is present when in reality it isn't. The second type is known as a "missed detection", which is when we decide the object is not present but in reality it is. There are thus two optimality criterion for detectors. The first is the detector's probability of false alarm P_{FA}, which is the probability that the detector will decide the object is present given that it is not present. The second is the probability of missed detection P_{MD}. It is also common to discuss the complementary probability for P_{MD}, namely the probability of detection P_D defined as

$$P_D = 1 - P_{MD} \tag{26}$$

Minimizing one type of error is easy—for example, we could always decide the object is present, in which case our P_{MD} is zero. In general, the task of choosing a detector is to find one which minimizes both errors simultaneously. The importance of minimizing one type of error over the other will depend greatly on the particular problem.

2.5.1 Neyman-Pearson Detector. The Neyman-Pearson (NP) detector minimizes P_{MD} for a desired P_{FA}. Thus when using NP detectors it is crucial to know what an acceptable P_{FA} is for the problem considered. Let $\mathbf{x}$ be the set of observations collected, O be the event that the object is present, O' be the event that the object is not present, and the notation $\Pr(A; B)$ read as "the probability of receiving data set A given that B is true". Then the NP detector decides the object is present if

$$\frac{\Pr(\mathbf{x}; O)}{\Pr(\mathbf{x}; O')} > T \tag{27}$$

where T is a threshold found by integrating the denominator [19]. The NP detector maximizes the P_D for a chosen P_{FA} as long as the threshold is chosen correctly. Note that the ratio in (27) is also called the likelihood ratio (LR), as it is the likelihood of

getting the received data given O over the same quantity given O'. The NP detector is therefore also referred to as the likelihood ratio test (LRT).

  2.5.2 Matched Filter.   Consider the problem where the received data is a known signal with additive white Gaussian noise (AWGN):

$$\mathbf{x} = \mathbf{s} + \mathbf{w} \tag{28}$$

where $\mathbf{s}$ is a known deterministic signal and $\mathbf{w}$ is a vector of white Gaussian noise. The NP detector for this problem decides the signal is present if [19]

$$\sum_{n=0}^{N-1} x_n s_n > T \tag{29}$$

where T depends on the desired P_{FA}. This summation is known as the matched filter (MF), which is widely used in signal processing. This is the NP detector only for signals in white noise; however, many real-world noise sources can be approximated as band-limited AWGN. The MF is also popular due to the availability of a fast implementation via the FFT:

$$\sum_{n=0}^{N-1} x_n s_n = [IFFT\{FFT\{\mathbf{x}\}^c FFT\{\mathbf{s}\}\}]_0 \tag{30}$$

where $()^c$ refers to the element-wise complex conjugate and $[]_0$ is the 0th element of the resulting vector. The other elements of the right-hand side are the matched filter summations for varying time-lags. Thus the FFT implementation allows the fast computation ($\Theta(n \log n)$) of the MF for arbitrary time lags, which is useful when $\mathbf{x}$ may contain $\mathbf{s}$ with an unknown time-delay.

2.6   Radar Imaging and Navigation

2.6.1   Introduction.   Radars have been used to generate imagery as early as the 1950s [21]. The original idea of generating images from sequential radar data collections came from the study of microscopes [22], which outlined wavefront reconstruction theory. The first known implementation of an imaging radar was by a Goodyear Aircraft Corp. engineer [14]. The lack of high-speed computers initially required the use of analog approximation, via the Fresnel approximation [23]. In the 1970's, a digital form of this approximation was developed [21], as well as a new digital processing method called polar format processing [24]. While this development was occuring in the radar community, a simultaneous development was occuring in the field of medical imaging [25]. The development of computerized axial tomography (CAT), which revolutionized medical diagnostics, was developed in the 1970s based on the same principles as SAR wavefront reconstruction. In particular, CAT imagery is generated by transmitting an X-ray beam through an object to be imaged repeatedly, while rotating the X-ray transmitter in a circle around the object. The resulting set of data is then processed to form a digital image of the object. This setup is identical to spotlight SAR where an aerial vehicle (AV) is flown in a circle around an imaging plane and data is collected periodically. The only difference between the two techniques is the usage of X-ray vs. radio waves [14].

Modern SAR imaging study has developed a wide range of new digital processing techniques, but each technique is still fundamentally based on the principles developed in the 1970s. The rest of this section will discuss two digital construction algorithms which are popular today; namely, backprojection and polar format processing [21].

2.6.2   Polar Format Algorithm.   The critical theory which allows SAR imagery to be constructed is the projection-slice theorem [26]. This theorem states that "the one-dimensional Fourier transform of a projection function $p_\theta(u)$ is equal to the two-dimensional Fourier transform $G(X, Y)$ of the image to be reconstructed when the two dimensional Fourier transform is evaluated along a line in the Fourier plane

which lies at the same angle θ measured from the X axis" [14]. That is, let $g(\cdot, \cdot)$ be a 2D reflectivity map and $p_\theta(\cdot)$ be a projection function running through the 2D space at angle θ, such that $p_\theta(\zeta)$ samples the 2D reflectivity map at the polar coordinates with angle θ and magnitude ζ. Then the projection-slice theorem states that

$$G(U \cos \theta, U \sin \theta) = P_\theta(U) \tag{31}$$

where U is a dummy variable and G, P are the Fourier transforms of g and p respectively. We see then that collecting a set of time-domain samples along lines at different angles $p_{\theta_1}, p_{\theta_2}, \ldots$ and taking the Fourier transform of each collection will allow us to populate the space G via the projection-slice theorem. The inverse Fourier transform of G is g, the reflectivity map we desire. Thus this process will allow us to construct a reflectivity map (i.e. image) of a 2D scene. The basic methodology is:

1. Circle a target, collect linear samples

2. Compute the Fourier transform of each data collection, and then use the projection-slice theorem to map these onto the 2D Fourier transform of the 2D reflectivity map of the target.

3. Compute the inverse Fourier transform of the resulting data set, which yields the 2D reflectivity map.

One issue in the above description is that it doesn't consider discrete systems. In a realistic system, we will only be able to collect samples of the projection function p. The DFT will then only contain samples of P. If the image to be constructed is also digital, then G will be discrete. The issue then is that the discrete samples of P will not be located at the same points as the discrete samples of G, and we cannot then directly use the projection-slice theorem. Fig. 2 from [14] illustrates this mismatch between G and P. The collected samples are located at the circular points, and the reconstructed image samples are at the grid intersections. The solution is to interpolate the polar samples to estimate the grid intersections, allowing image

23
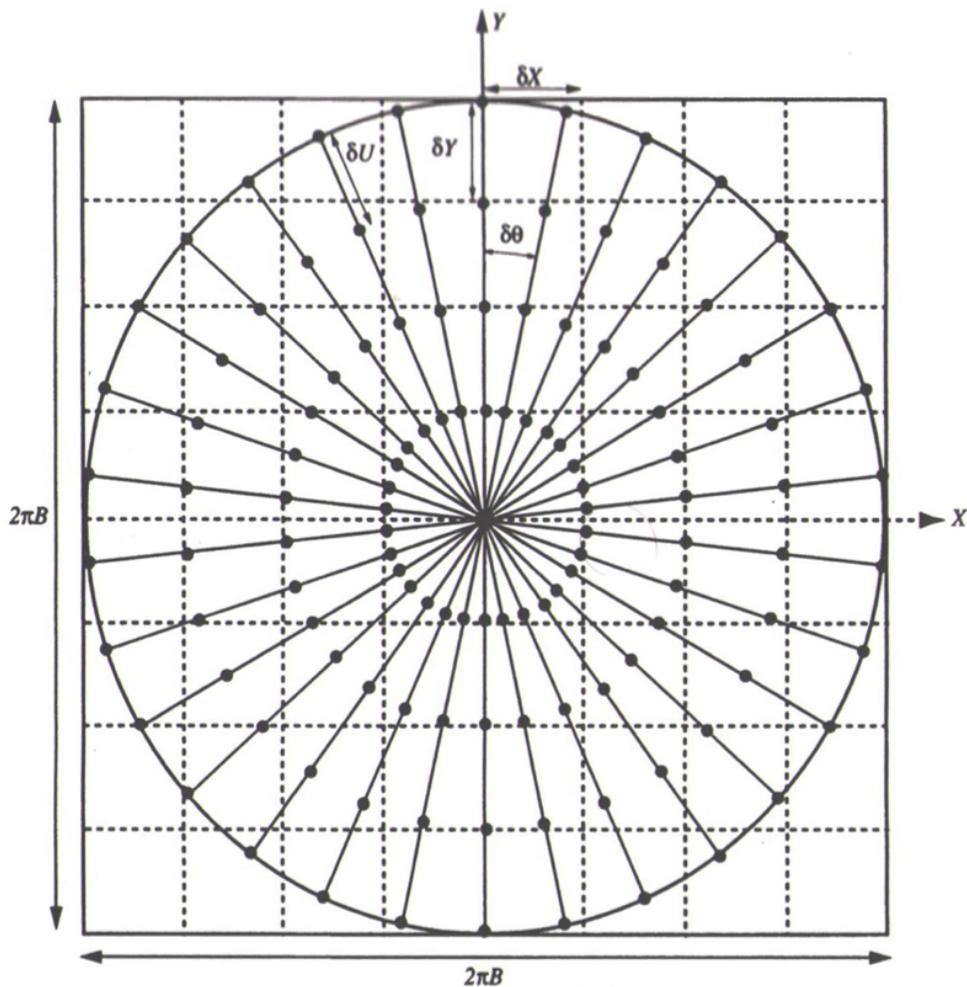
Figure 2:   Polar format mismatch between collected data and the reconstructed image. (Taken directly from [14])

formulation to proceed as normal. This approach is to image formulation is called the polar format algorithm, as the primary difficulty is in mapping the polar data onto a rectangular grid. Note that the principles described here also extend to 3D scenarios such as navigation by simply formulating the 3D version of the theorem [14].

2.6.3 Backprojection.   Backprojection is also based on the projection-slice theorem and is theoretically equivalent to polar format processing. It uses the convo-

lution property of the Fourier transform to rewrite the projection-slice theorem into the time domain. The resulting construction method is [14]

$$g(\zeta_1, \zeta_2) = \frac{1}{4\pi^2} \int_{-\pi/2}^{\pi/2} [p_\theta * h](\zeta_1 \cos\theta + \zeta_2 \sin\theta) d\theta \qquad (32)$$

where $*$ is the convolution operator and h is a filtering kernel which is a function of the geometry. This method yields the same results theoretically as the polar format algorithm, but may yield better results practically due to avoiding the frequency domain interpolation issues previously discussed. Instead, the calculation of the convolution $p_\theta * h$ becomes the critical factor in the accuracy of the produced image. This may be difficult, as p is sampled and only available at discrete θ points, which results in the use of approximations to form the image.

2.6.4   Radar Image-based Navigation.   Radar based navigation sensors have been actively developed in the form of terrain aided navigation systems (TANS) [27–30]. In these studies, the presence of digital terrain elevation data (DTED) is assumed. The data collected from the radar is correlated to DTED maps which were collected *a priori*. This necessarily precludes navigating over unknown terrain, as no DTED will be available.

To navigate in unknown terrain, it is necessary to extract features from the radar image and position them while simultaneously estimating the vehicle position. This technique is called simultaneous localization and mapping (SLAM). SLAM has been extensively studied in image-aided navigation [13], however only cursory work has been done exploring the use of radar using such an approach [31].

2.6.5   Radar Doppler-Aided Navigation.   There has been previous interest in using on-board range-Doppler radars for dead-reckoning navigation [32–34]. These systems use a method called Doppler-aiding, where reflections from an assumed surface or object are processed to estimate the Doppler frequency of the reflection. This

Doppler can then be used as a velocity update in the navigation filter, aiding and correcting the velocity state of the navigation solution.

Although this approach will reduce the error of a stand-alone INS navigation solution, it does not provide nearly as much information to the filter as a tracked target state-based SLAM approach, such as the one developed in Chapter III. Suppose we have an AV with an onboard INS and radar system. A Doppler-aiding system will measure the Doppler to the ground or one or more of the radar features in the environment. This Doppler estimate—containing errors—will be used to estimate the velocity of the AV, which will be given to the navigation filter as a noisy measurement. Thus the filter has access to a noisy velocity measurement (from the radar) and a noise acceleration measurement (from the INS). Since both of these measurements must be integrated to obtain position information, the position estimate generated by the EKF will have drift due to the errors in the measurements. In addition, Doppler aiding only measures velocity in the direction of plane travel (i.e., 1-dimension), and cannot be used to correct or estimate AV bearing or flight path.

In contrast, the approach developed in Chapter III measures ranges to positioned (via SLAM) reflectors. If enough reflectors are available with good geometry (low GDOP), a full 3-D position solution is possible using a similar solution to that of GPS. The SLAM solution can estimate the location of the AV in any direction that it has observability in. Thus the solution accuracy in 3-D depends on the GDOP induced by the environment and reflector locations. In combination with an INS, the calculated navigation solution will be much higher than that of a Doppler-aiding solution.

In addition, even if bad geometry or limited reflector availability is encountered, the SLAM-based solution will not drift as much as an INS. To observe this, we consider a hypothetical scenario where the AV flies in a circle around a terrain with at least one always visible feature. The Doppler-aiding configuration will drift as a function of its Doppler estimate error, as stated before. The SLAM approach would continually

observe a range from at least one reflector, constraining the possible AV location to within sight (detection range) of that target. Thus as long as the AV flew in sight of that reflector, the AV solution error would be constrained (i.e., not drift away). Although this scenario is not realistic, it highlights the increased information provided by the SLAM approach. The drift in a SLAM feature-tracking approach will increase linearly as a function of the time a particular feature is in view (feature detection range), and the error in ranges are not integrated to form a position solution (as is the case for INS and Doppler-aiding).

2.6.6   Clutter Models.   One of the primary sources of interference for imaging radars is the existence of unwanted nuisance scatterers called clutter. Clutter models are used to describe background noise generated from scatterers which are not targeted by the radar. What constitutes clutter depends widely on the specific application. For example, a radar attempting to detect low flying vehicles will consider all stationary objects on the surface as clutter; whereas an imaging radar attempting to generate a picture of the surface of the Earth will consider low flying vehicles clutter. In our case, we wish to extract the location of strong stationary persistent scatterers from a target scene. Thus things like foliage, grass, tree canopies, and moving objects are considered clutter. Our ability to navigate will be based on our ability to distinguish useful targets for navigation from background noise.

There are many studies done on airborne clutter in particular environments such as forest canopies [35–37] and sea/water bodies [38]. The standard approach in unknown terrain is to use a statistical clutter model, such as chi-squared or log-normal clutter distributions [39–41]. In [40] we see that UWB clutter at low grazing angles is well approximated by independent log-normal noise on each receiver channel. Since our airborne simulations in Chapter IV are at high grazing angles, the long tails of log-normal statistical clutter models aren't necessary. We therefore use independent Gaussian noise sources on each receiver channel.

2.7 Orthogonal Frequency Division Multiplexing

2.7.1 Introduction.   Frequency division multiplexing (FDM) is a technique that allows the transmission of multiple sub-signals inside a single signal by encoding them in different frequency bands. FDM is extremely common in every day life—for example, the radio in your car actually receives one signal containing many radio station signals transmitted at different frequencies. When you turn the dial to a specific station, you are specifying which frequency range (sub-signal) you wish to receive. FDM requires each transmitted signal to be non-overlapping. For example, suppose you have two audio signals you want to transmit, and each requires 100 kHz of bandwidth. You might decide to transmit the first audio signal in the frequency band $0 - 100$ kHz. Then the next audio signal would have to be transmitted at $100 - 200$ kHz, as it cannot overlap with the first signal.

One issue with this scheme is that it is very difficult to immediately cut a signal off at a particular frequency, as with our example above which required exactly 100 kHz of spectrum. Typically, a signal will slowly decrease energy at its high and low frequencies. A packing scheme such as our audio signal above will cause the signals to interfere, as both signals bleed into each other's reserved frequency bands. A solution to this problem is to allow guard bands to exist between transmitted signals. In our audio signal example we could encode the first signal at $0 - 100$ kHz and the second at $150 - 250$ kHz, leaving a 50 kHz space between the sub-signals to reduce interference. This scheme "wastes" bandwidth, as the guard band is essentially unused.

A more optimal solution is to keep the channels closely packed and encode the sub-signals such that the energy they bleed into adjacent channels will not affect the recovery of information. This can be done by carefully construction each sub-signal to be orthogonal to all other sub-signals, a scheme known as orthogonal FDM (OFDM). Fig. 3 illustrates an example OFDM waveform spectrum. The red, blue, and green lines represent the energy of channels encoded at $5, 6$, and 6.2 kHz respectively. We see that each channel bleeds energy into every other channel. However, at the exact

channel center, the energy contributed by every other channel is zero. For example, at 6 kHz, both the green and blue lines are zero. Thus the amplitude/phase of the 6 kHz sub-signal can be recovered without interference from adjacent signals. This allows for each sub-signal to be amplitude modulated (AM) and phase modulated (PM), and data encoded to be recovered via measuring the AM and PM of each sub-signal.

OFDM approaches the optimal data transmission rate given by the Shannon-Hartley theorem, and thus is widely used in communications. However, it also has advantages for radar. One advantage is the arbitrary encoding of the spectrum as seen in Fig. 3 and 4. OFDM also has implementation issues with high peak-to-average power ratio (PAPR) requirements and signal processing requirements for tracking and detection due to the necessity of pulse compression. In particular, high-powered amplifiers have a fixed voltage swing which makes it difficult to transmit large power spikes without scaling the transmitted waveform down and thus transmitting at lower power. However, modern amplifiers, OFDM encoding methods, and high speed computational capabilities are removing the limitations typically encountered in OFDM systems.

2.7.2 Signal Model. The reference signal to be transmitted at each point $\mathbf{p}_k$ is an UWB-OFDM pulse defined by

$$s_r(t) = \Re\left\{ \sum_{k=0}^{N_c-1} \xi_k \exp[j2\pi t(f_0 + k\Delta f)] \right\} \tag{33}$$

Each exponential in the summation corresponds to a particular OFDM channel. N_c is the number of channels transmitted, f_0 is the fundamental frequency, Δf is the channel spacing, and ξ_k is the complex-valued modulation for channel k. Typically the OFDM waveform is transmitted as a time-limited pulse. In order for channel orthogonality to hold, the window length must be a multiple of the period of each channel. Assuming the fundamental frequency is zero (i.e., the OFDM symbol is constructed at baseband), the OFDM waveform is then
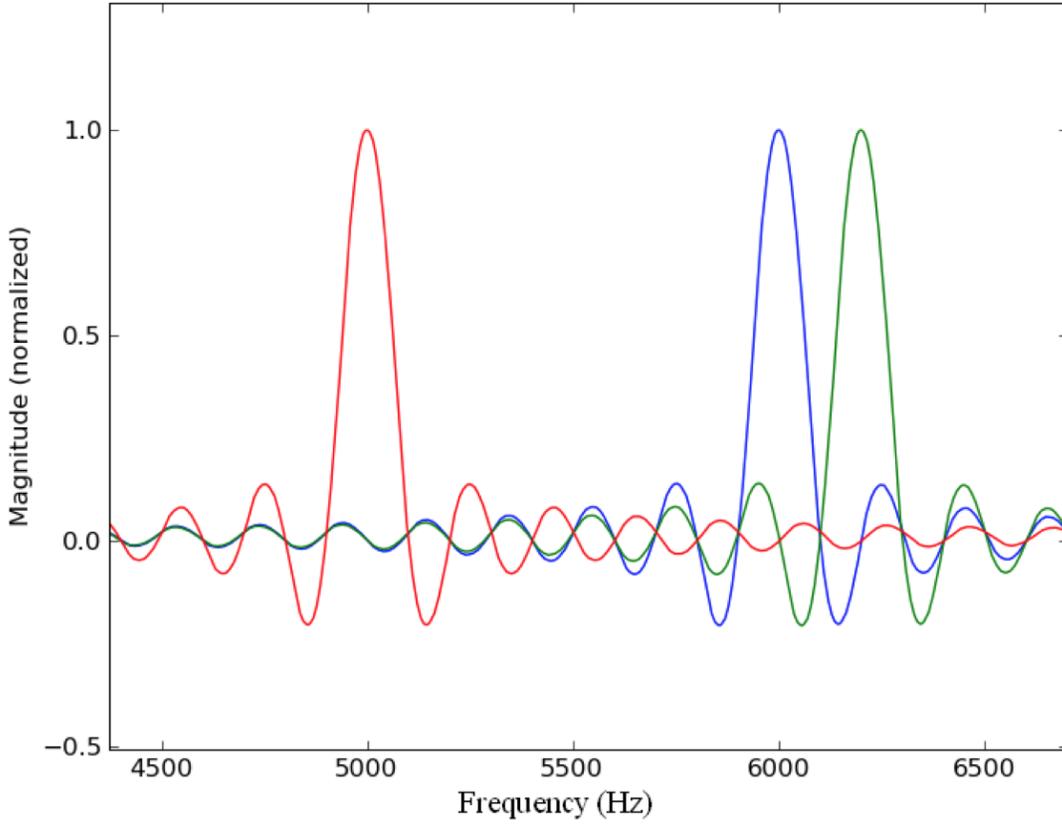
Figure 3: Example OFDM waveform and spectrum packing.

$$s_{ru}(t) = s_r(t) \left(\mathcal{U}(t) - \mathcal{U}\left(t - \frac{1}{\Delta f}\right) \right) \tag{34}$$

where $\mathcal{U}(\cdot)$ is the unit step function. Let the discrete OFDM pulse with sampling frequency f_s be denoted

$$S_{ref}(k) = s_{ru}\left(k\frac{1}{f_s}\right), \quad k = 0, 1, \ldots, \frac{2D_{av}}{c}f_s \tag{35}$$

where D_{av} is the maximum distance to be illuminated by the transmitted beam. Fig. 4 illustrates a transmitted 128ns OFDM symbol with random sub-carrier modulation (64 sub-carriers). The green dots represent the discrete sub-carriers in the continuous complex-valued spectrum (represented here as phase/magnitude). Since the modula-
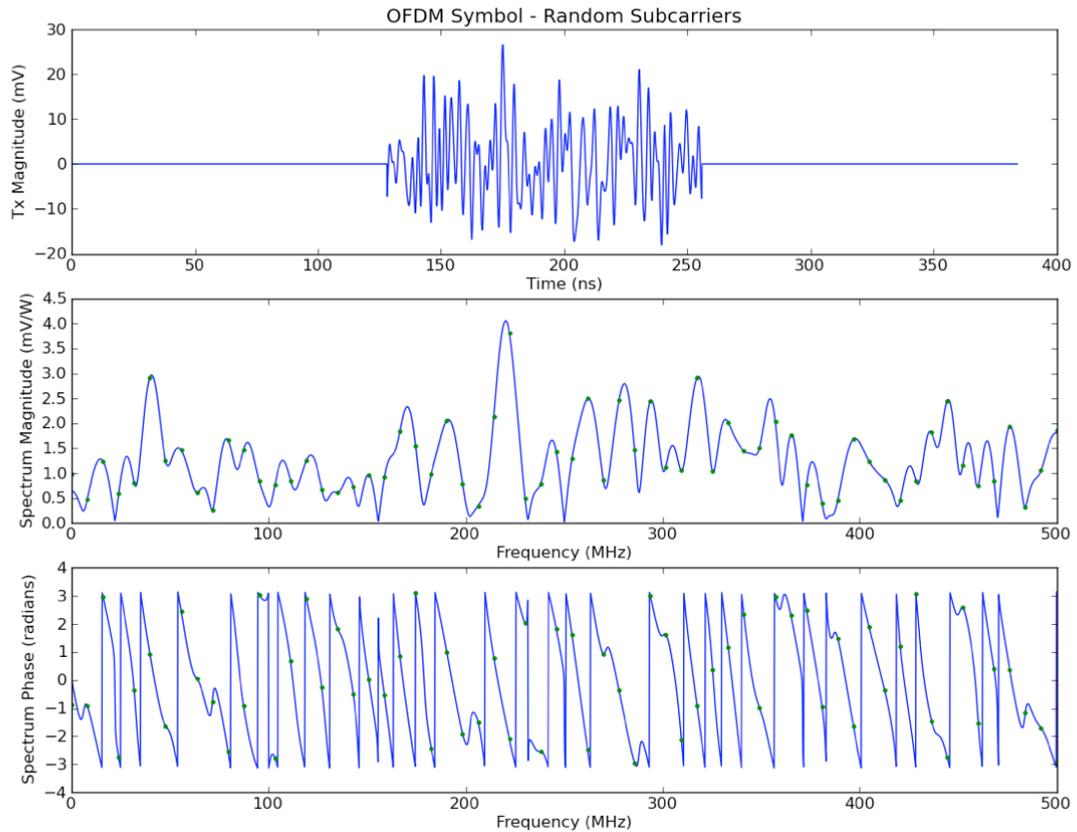
Figure 4:   Example transmitted OFDM symbol with random modulation.

tion is random, we see the discrete points have a random distribution, whereas the continuous spectrum between the discrete points is highly correlated. Fig. 5 shows another OFDM symbol which is modulated with an example communications signal. The first half of the sub-carriers are modulated with zero phase and zero magnitude except for the 2nd, 3rd, 5th and 9th sub-carrier which have positive magnitude. The second half of the sub-carriers are modulated with positive magnitude and one of four possible phases, dictated by the data stream being encoded. The encoding of a symbol with four possible phases to represent two bits of data is called 4-Quadrature Amplitude Modulation (4-QAM). Thus each of the sub-carriers in the second half of the symbol's spectrum are encoded as 4-QAM.

31
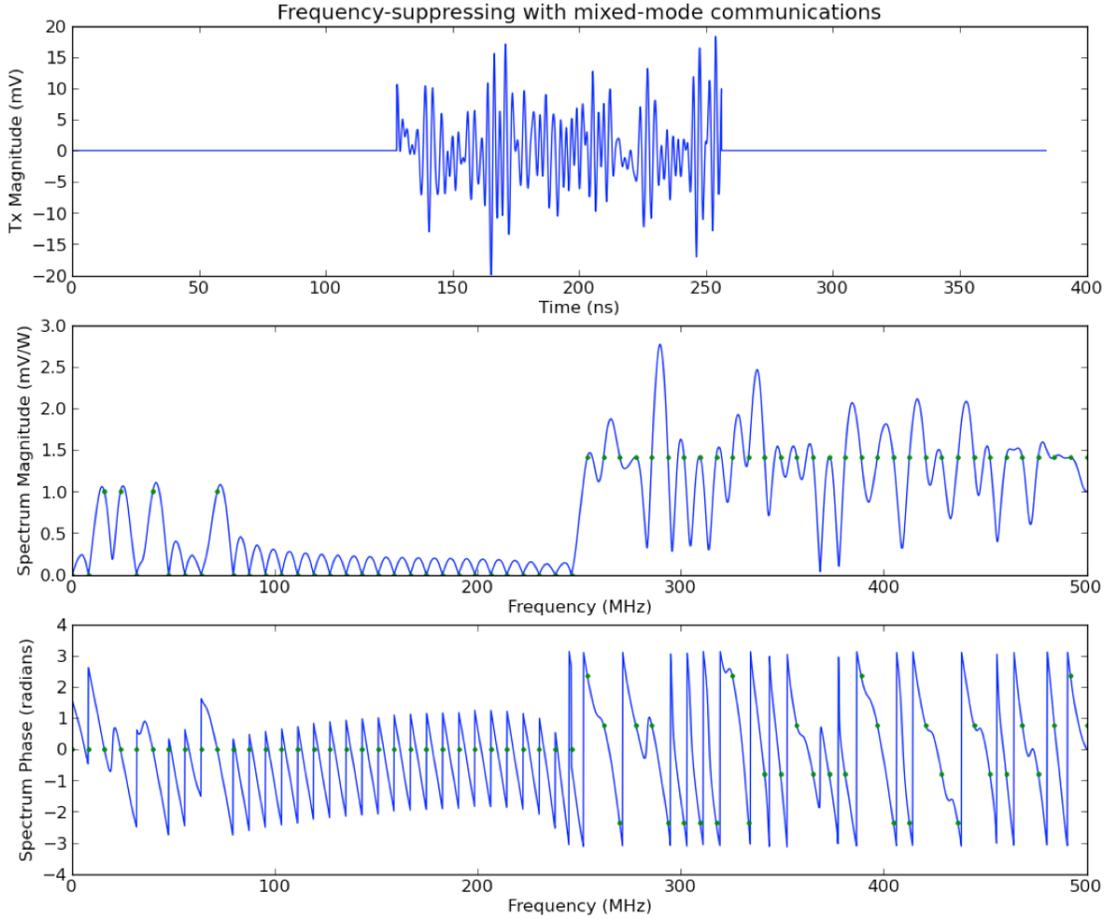
Figure 5: Example transmitted OFDM symbol with preset modulation on the first half of the sub-carriers and 4-QAM on the second half.

2.7.3 Ultra-Wideband. OFDM itself does not necessitate a bandwidth usage—OFDM symbols can be narrow-band or wide-band. However, a fundamental property of radars is that the range resolution (i.e., the minimum distance between two targets that we could still resolve them as distinct and separate from each other) is a function of the bandwidth of the system [14]. Specifically, the range resolution $\Delta d_{\min}$ is given by

$$\Delta d_{\min} = \frac{c_p}{2B_e} \tag{36}$$

32

where c_p is the propagation speed (nominally the speed of light) and B_e is the effective bandwidth of the waveform. For OFDM the effective bandwidth is calculated via

$$B_e = N_c \Delta f \tag{37}$$

We see that the range resolution is therefore directly dependent on the bandwidth used by the full OFDM symbol. For precision navigation, ultra-wideband OFDM symbols are highly desired due to this property.

2.8 Coherent Demodulation

One issue in the design of a radar system is the need for coherent modulation. Suppose we have an arbitrary baseband signal $I(t)$ we want to transmit. We first modulate the signal to the carrier frequency

$$s_T(t) = I(t) \cos(\omega_c t + \Phi_{LO,T}) \tag{38}$$

where $\Phi_{LO,T}$ is the phase of the local oscillator at transmission. The received signal is then

$$s_R(t) = s_T(t - t_d) = I(t - t_d) \cos(\omega_c(t - t_d) + \Phi_{LO,T}) \tag{39}$$

At baseband we have

$$s_{RB}(t) = s_R(t) \cos(\omega_c t + \Phi_{LO,R}) \tag{40}$$

where $\Phi_{LO,R}$ is the phase of the local oscillator when received. Using the half angle trig identity and rewriting we have

$$s_{RB}(t) = \frac{1}{2} I(t - t_d)[HFC + \cos(\Phi_{LO,T} - \omega_c t_d - \Phi_{LO,R})] \tag{41}$$

33

where HFC is a high frequency component which will be filtered out by a low pass filter, yielding:

$$s_{RB-LPF}(t) = \frac{1}{2}I(t - t_d)\cos(\Phi_{LO,T} - \omega_c t_d - \Phi_{LO,R}) \tag{42}$$

If we allow the oscillator phases to be incoherent (i.e., randomly wander), then there is the possibility that a target return will not be detected at baseband, since the term $\cos(\Phi_{LO,T} - \omega_c t_d - \Phi_{LO,R})$ will attenuate the power of the received signal. In communications, we could simply use a phase lock loop (PLL) to set

$$\Phi_{LO,R} = \Phi_{LO,T} - t_d \tag{43}$$

However, in radar, this is infeasible, since we may be tracking multiple targets at once resulting in a set of t_d values. We would need a set of local oscillators and PLLs, one for each target. A more practical solution is to use an I/Q receiver to detect the signal with an arbitrary phase difference. We therefore downconvert twice, generating an in-phase and quadrature-phase baseband signal:

$$s_{RB-I}(t) = s_R(t)\cos(\omega_c t + \Phi_{LO,R}) \tag{44}$$

$$s_{RB-Q}(t) = s_R(t)\sin(\omega_c t + \Phi_{LO,R}) \tag{45}$$

Using trigonometric identities and applying a low-pass filter we have

$$s_{I-LPF}(t) = \Re e\{I(t - t_d)\exp(j(\Phi_{LO,T} - \omega_c t_d - \Phi_{LO,R}))\} \tag{46}$$

$$s_{I-LPF}(t) = \Im m\{I(t - t_d)\exp(j(\Phi_{LO,T} - \omega_c t_d - \Phi_{LO,R}))\} \tag{47}$$

Adding the two returns, we now have a complex exponential which is amplitude modulated by our received signal, regardless of the phase difference between

the transmit LO and receive LO. Thus an I/Q demodulator is required to perform coherent down conversion of radar returns.

III. Algorithm Design and Mathematical Model

This chapter describes the algorithms and mathematical models which have been developed and used in order to study radar-based navigation in airborne, indoor, simulated, and experimental contexts.

3.1 Radar Waveform

There are many waveforms used for radar. Chirps are used due to their easy processing properties; continuous wave transmissions are used for precision Doppler estimation; RF pulses are used due to their easy hardware implementation and unambiguous range measurements [15]. One recent waveform of interest is the Ultra Wideband (UWB) orthogonal frequency division multiplexed (OFDM) pulse. As seen in Chapter II. UWB-OFDM has many properties which make it suitable for precision navigation, including jamming resistance, high range resolution, and potential for use as a low probability of intercept (LPI) device.

We have previously studied the UWB-OFDM waveform for radar extensively [42–44]. In addition, we have access to a UWB-OFDM experimental prototype system that will be described in Section 5.1.1, allowing for experimental validation of the algorithms designed. We therefore chose UWB-OFDM as the waveform transmitted for all simulation and experimental results in this study. Please see Section 2.7.2 for details of the signal model and notation.

3.2 Aerial Vehicle Model

The primary purpose of navigation is to discover the AV position using information collected from our sensor array. For a direct-state EKF implementation, we need to make assumptions about the dynamics (i.e. the range of possible motion we might realistically expect to encounter) of the AV. This is useful for tuning the filter dynamics to match the specifics of the chosen AV (rocket, plane, etc.). This section describes the AV related notation and the dynamics model used.

The position of the AV at time t is denoted

36

$$\mathbf{p}(t) = \begin{bmatrix} p_{\mathbf{x}}(t) \\ p_{\mathbf{y}}(t) \\ p_{\mathbf{z}}(t) \end{bmatrix} \tag{48}$$

The AV collects INS/radar data at a fixed time interval Δt_p with the first transmission at time $t = 0$. Thus the pulses are transmitted at positions

$$\mathbf{p}_k = \mathbf{p}(k \Delta t_p), \quad k = 0, 1, \ldots \tag{49}$$

The transmitted beam is side-looking and directed downward towards the terrain as illustrated in Fig. 6, assuming the AV has an initial heading due south:

$$\frac{\dot{\mathbf{p}}(0)}{|\dot{\mathbf{p}}(0)|} = \mathbf{x} \tag{50}$$

The AV dynamics are modeled using a first-order Gauss-Markov (FOGM) acceleration model [20, 45]. A FOGM process is a stationary Gaussian process defined by its autocorrelation function. Let $\sigma_{\ddot{\mathbf{p}}}^2$ be the process variance and $\tau_{\ddot{\mathbf{p}}}$ be the time constant for the AV acceleration (equal for all axes). Then the autocorrelation of the AV acceleration along the $\mathbf{x}, \mathbf{y}$ and $\mathbf{z}$ axes is given by

$$R_{\ddot{p}_{\mathbf{x}}}(\zeta) = R_{\ddot{p}_{\mathbf{y}}}(\zeta) = R_{\ddot{p}_{\mathbf{z}}}(\zeta) = \sigma_{\ddot{\mathbf{p}}}^2 e^{-\tau_{\ddot{\mathbf{p}}}|\zeta|} \tag{51}$$

The continuous-time AV acceleration dynamics equation is [45]:

$$\frac{d}{dt}\ddot{\mathbf{p}}(t) = -\tau_{\ddot{\mathbf{p}}}\ddot{\mathbf{p}}(t) + \mathbf{w}_{\ddot{\mathbf{p}}}(t) \tag{52}$$

where $\mathbf{w}_{\ddot{\mathbf{p}}}(\cdot)$ is a column vector of three zero mean $i.i.d.$ white Gaussian noise processes defined by

$$E[(\mathbf{w}_{\ddot{\mathbf{p}}}(t))(\mathbf{w}_{\ddot{\mathbf{p}}}(t+\zeta))^T] = \mathbf{I}Q_{\ddot{\mathbf{p}}}\delta(\zeta) \tag{53}$$

and $\delta(\cdot)$ is the Dirac delta function. The noise strength $Q_{\ddot{\mathbf{p}}}$ is chosen to yield the desired auto-correlation function in (51). From [45] we have

$$Q_{\ddot{\mathbf{p}}} = \left(\frac{2}{\tau_{\ddot{\mathbf{p}}}}\right)\sigma_{\ddot{\mathbf{p}}}^2 \tag{54}$$

In order to calculate the AV positions at transmission times $k\Delta t_p$, the continuous dynamics model in (52) must be discretized into a difference equation. We follow the formulation in [45] for our discretization process. The state transition matrix is

$$\mathbf{\Phi}_{\ddot{\mathbf{p}}}(t_1 - t_2) = \mathbf{\Phi}_{\ddot{\mathbf{p}}}(\Delta t_p) = e^{-\tau_{\ddot{\mathbf{p}}}\Delta t_p} \tag{55}$$

Letting $t_1 - t_2 = \Delta t_p$, the discrete time acceleration dynamics equation is

$$\ddot{\mathbf{p}}_{k+1} = \ddot{\mathbf{p}}_k e^{-\tau_{\ddot{\mathbf{p}}}\Delta t_p} + \mathbf{w}_{\ddot{\mathbf{p}}d}(k\Delta t_p), \quad k = 0, 1, \dots \tag{56}$$

where $\mathbf{w}_{\ddot{\mathbf{p}}d}(\cdot)$ is the equivalent discrete time noise process to $\mathbf{w}_{\ddot{\mathbf{a}}}(\cdot)$. $\mathbf{w}_{\ddot{\mathbf{p}}d}(\cdot)$ is a zero mean white Gaussian process defined by

$$E[\mathbf{w}_{\ddot{\mathbf{p}}d}(t_k)\mathbf{w}_{\ddot{\mathbf{p}}d}(t_k)^T] = \int_{t_k}^{t_k+\Delta t_p} \mathbf{\Phi}(t_k - \zeta)Q_{\ddot{\mathbf{p}}}\mathbf{\Phi}^T(t_k - \zeta)\,\mathrm{d}\zeta \tag{57}$$

This integral is evaluated numerically via the Van Loan method described in [20] to generate noise realizations. The noise realizations can then be used with (56) to generate sample AV locations during simulations.

The FOGM acceleration model defined here allows the AV to fly freely, while still constraining its trajectory to realistic flight dynamics. In particular, an AV should avoid discontinuities in acceleration. The time constant $\tau_{\ddot{\mathbf{a}}}$ must be tuned to match the expected flight dynamics. For example, during aerial maneuvers acceleration may
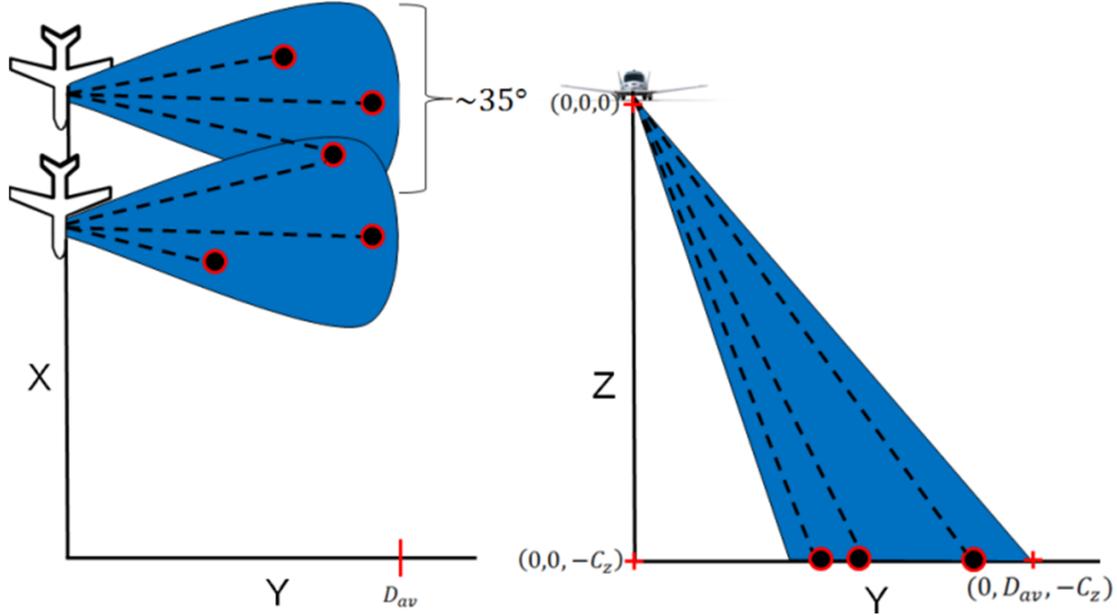
Figure 6:   Illustration of the AV side-looking stripmap SAR configuration.

be completely uncorrelated within fractions of a second, whereas during level flight it may remain correlated for several hours. For the simulations in this dissertation, a time constant $\tau_{\ddot{\mathbf{a}}} = 3$ seconds is used. In (56) we see that the possible acceleration values are defined by the noise strength of $\mathbf{w}_{\ddot{\mathbf{p}}d}$, which is related to $\sigma_{\ddot{\mathbf{a}}}^2$ via (57) and (54). Therefore we must choose the process noise $\sigma_{\ddot{\mathbf{a}}}^2$ proportional to the maximum acceleration expected during flight.

3.3   Environment Model

During simulations, we will need to model the environment and reflector characteristics. This section describes the environment and reflector models used.

The modeling of interaction between the environment and the signal is easiest when performed in the frequency domain. We will consider a single OFDM channel reflecting off a single reflector in this section, as each channel in $s_{ru}(\cdot)$ is a single modulated sinusoid with constant frequency.

Let $s_{ch}(\cdot)$ be a single channel of the form presented in (33). Let ξ_{ch} and f_{ch} be the modulation and frequency of the chosen channel. The single channel OFDM signal is transmitted via the radar, travels through the environment, reflects off the target (possibly more than one), travels back to the radar, and finally is received by the radar. Fig. 7 gives an overview of the environment effects on the signal. The channel is first upconverted to the transmit frequency:

$$s_{tx}(t) = s_{ch}(t)\sin(2\pi f_c t) \tag{58}$$

where f_c is the carrier frequency. The transmitted signal is then scaled to the transmitting power r_P and the antenna gain r_G. During propagation the signal attenuates as the inverse square of the distance and has a time delay proportional to the distance from the AV to the reflector. The signal incident on the reflector is:

$$s_{inc}(t) = \frac{r_P r_G}{4\pi d(\mathbf{p}(T_t), \mathbf{r})^2} s_{tx}\left(t - \frac{d(\mathbf{p}(T_t), \mathbf{r})}{c}\right) \tag{59}$$

where $d(\cdot, \cdot)$ is the distance between two points, $\mathbf{r}$ is the position of the reflector, and T_t is the time of pulse transmission. Note the absence of an explicit Doppler term, even though the AV and reflector have non-zero relative velocity. This is due to the Doppler being approximated as zero over the time interval of a single pulse, i.e. we assume no significant acceleration is experienced during collection. For a fixed Doppler over the time interval of a single pulse transmission, the reference signal used in the matched filter during signal detection and tracking can have matching Doppler added to it. Therefore the Doppler shift for a single pulse will not affect the results of the matched filter other than to add computational burden, and so the Doppler term is omitted here. The Doppler frequency is still estimated for use in initializing the navigation filter; however, it is collected on a pulse-to-pulse basis by examining the complex phase of the I/Q demodulator output in (63). This approach is standard practice in pulse-Doppler radar processing [12].
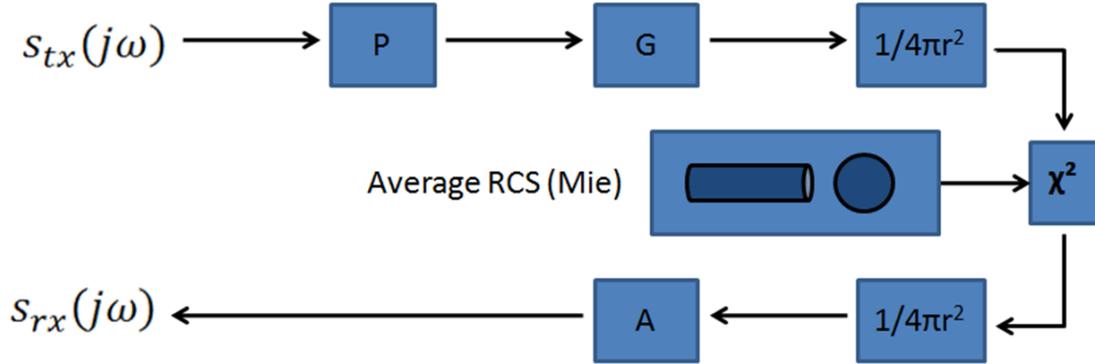
Figure 7: Environment model for channel loss and reflector backscatter.

The incident signal is then reflected off the scatterer as a function of the scattering characteristics of the reflecting surface. The reflectors are modeled as Lorenz-Mie spherical scatterers with Swerling [46] radar cross section (RCS) noise. The Mie scattering model generates a frequency dependent RCS of the scatterer, which will be different for each channel. This model allows us to account for frequency dependent attenuation of the received signal, which is significant for UWB. The Swerling RCS noise allows us to model random variation in RCS as a function of the angle of incidence. Swerling type I scattering was used, giving a density function

$$f_\chi(\zeta) = \frac{1}{\sigma_M} \exp\left(-\frac{\zeta}{\sigma_M}\right) \tag{60}$$

where σ_M is the average RCS provided by the Mie scattering model. In our simulations, σ_M is calculated using an iterative numerical approximation, with a sphere radius of 0.5m. The reflected signal is then

$$s_{return}(t) = s_{inc}(t)\chi \tag{61}$$

where χ is a random variable with the distribution given by (60). The returned signal at the radar experiences an additional inverse square power loss, time delay, and a scale factor accounting for the receiver antenna affective aperture. Equivalent thermal

41

noise is also added to the received signal to model the effects of ambient and receiver noise. Assuming one reflector is present, the received signal is then

$$
\begin{aligned}
s_{rx}(t) = {} & \frac{r_A}{4\pi d(\mathbf{p}(T_t), \mathbf{r})^2} s_{ret}\left(t - \frac{d(\mathbf{p}(T_t), \mathbf{r})}{c}\right) \\
& + w_{th}(t)
\end{aligned}
\tag{62}
$$

where $w_{th}(t)$ is an additive white Gaussian noise (AWGN) process with strength σ_{th}^2 and r_A is the receiver effective aperture.

The $d(\mathbf{p}(T_t), \mathbf{r})$ terms in (59) and (62) introduce an unknown phase modulation in the received signal $s_{rx}(\cdot)$. This unknown phase shift can be estimated through the use of an I/Q demodulator. This is explained in more detail in the next section. The received baseband signal is

$$
\begin{aligned}
s_{rxB}(t) = {} & s_{rx}(t) \sin(2\pi f_c t - \phi_c) \\
& + j s_{rx}(t) \cos(2\pi f_c t - \phi_c)
\end{aligned}
\tag{63}
$$

where ϕ_c is the unknown phase drift of the oscillator between the time of up-conversion and down-conversion (nominally zero).

Eqn. (63) is the received signal for a single channel reflecting off a single reflector. The actual received baseband signal is the sum of all received channels over all targets in the illumination beam:

$$
s_{rxS}(t) = \sum_{\text{targets}} \sum_{\text{channels}} s_{rxB}(t)
\tag{64}
$$

where s_{rxB} is the return of the form shown in (63) for that particular channel and target.

The discrete received baseband signal with sampling frequency f_s is

$$S_{rxS}(k) = s_{rxS}\left(k\frac{1}{f_s}\right), \quad k = 0, 1, \ldots, \frac{2D_{av}}{c}f_s \tag{65}$$

3.4 Radar Processing Algorithms

Due to the computational burden involved in conventional SAR data processing, we process the data on a per-collection basis using conventional radar tracking methodology. As each pulse is transmitted, the collected return samples are processed immediately, and the contained information is integrated directly into the navigation filter. This approach allows real-time navigation updates, and significantly reduces the computational burdern by only requiring 1D processing.

Even with this approach the amount of data may be too large for some applications, such as unmanned AVs (UAVs) and other embedded devices. We therefore present both deterministic and randomized algorithm implementations, where the randomized algorithm only processes selected parts of the data to improve speed.

3.4.1 *Feature Extraction.*   Let $S_{rx}(k, l)$ be the lth sample collected at the position $\mathbf{p}_k$. The matched filter (MF) of the kth collection is then

$$\mathbf{m}_k = \text{IFFT}(\text{FFT}(|S_{rx}(k, \cdot)|)^*\text{FFT}(|S_{ref}(\cdot)|)) \tag{66}$$

where $(\cdot)^*$ denotes the complex conjugate. In practice the reference pulse used in (66) will be oversampled to allow for sub-sample alignment; however, the sampling frequencies of both signals are set to f_s in this dissertation for clarity.

Fig. 8 shows the ideal matched filter response of an OFDM pulse with $N_c = 256$, $f_s = 1\text{GHz}$, $\Delta f = 3.906\text{MHz}$, and random normally distributed channel modulations. Fig. 9 shows the output of (66) for three reflectors with the noise strength of $w_{th}(\cdot)$ set to generate an SNR of 0dB. The ovals correspond to the peaks generated from the three reflectors. We see that the effects of the environment model combined with the sidelobes of the autocorrelation cause ambiguity in the location of targets and

may result in false alarms. This necessitates the use of adaptive thresholding which is normalized to the noise level. This is performed by thresholding the SNR of the MF, instead of the MF directly. We calculate the MF SNR as

$$SNR_{MF}(k) = \frac{P(\max \mathbf{m}_k) - P(\mathbf{m}_k) + \frac{P(\max \mathbf{m}_k)}{N_s}}{P(\mathbf{m}_k) - \frac{P(\max \mathbf{m}_k)}{N_s}} \tag{67}$$

where N_s is the number of samples and $P(\cdot)$ is the average power of the samples. Fig. 10-13 illustrate the calculated MF SNR for different strengths of $w_{th}(\cdot)$. We see that choosing a threshold for the SNR_{MF} is invariant to noise if a constant false alarm rate (CFAR) is desired. This is due to the noise strength normalization term $P(\mathbf{m}_k)$ in the denominator of (67). Note that the MF SNR is not true signal SNR, as it is calculated via (67) as the ratio of the peak power to the average power of a MF output. This quantity has a baseline of 9dB even when no signal is present. If the MF SNR is above the chosen threshold, the peak at $\arg \max \mathbf{m}_k$ is removed from $\mathbf{m}_k$ and (67) is calculated again. This process is repeated iteratively until the MF SNR chosen is below the threshold. This iterative process allows the detection of multiple targets in a single OFDM data collection.

Table 1 shows the calculated probability of false alarm and missed detection *per sample collected* using the MF SNR approach versus a simple threshold on the MF peak. The simple threshold method was tuned to operate at a SNR of 14dB, and the red entries denote a critical failure region. For navigation, missed detections are not of critical importance, as 3-6 tracked reflectors are sufficient. However, P_{fa} is extremely important, as bad measurements introduced into the EKF with low associated covariance matrices may corrupt the entire navigation solution, causing growing linearization errors. Therefore, the CFAR achieved by the MF SNR approach is desired for navigation.

3.4.2 Feature Tracking. The feature extraction algorithm produces a set of observations $\mathbf{z}_k$ collected at position $\mathbf{p}_k$. In this section, we first define a data

Table 1:   Probability of Missed Detection and False Alarm For Simple Thresholding and MF SNR Thresholding

MF SNR	MF P_{fa}	MF P_{md}	Simple P_{fa}	Simple P_{md}
9	< 1%	> 99%	> 99%	< 1%
10	< 1%	45%	32%	< 1%
13	< 1%	3%	2%	< 1%
18	< 1%	< 1%	< 1%	< 1%
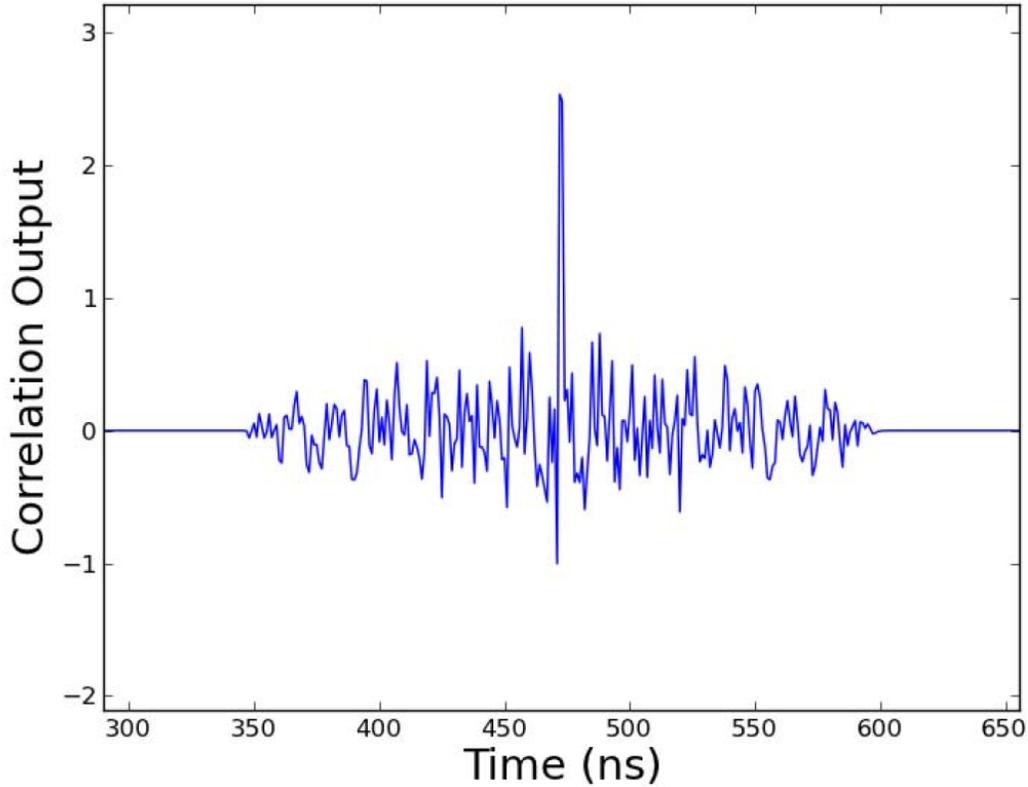
Figure 8:   Matched filter output of an OFDM pulse reflecting off a perfect reflector at a range of 140 meters.

association method using a global nearest neighbor (GNN) approach, and then use the association to implement an M/N detector. Both methods are widely used in conventional radar processing [47, 48].

3.4.2.1   GNN Data Association.   Let $\mathbf{t}_k$ be the set of previously tracked targets at position $\mathbf{x}_k$, initially empty. Once calculated, the position estiamtes of the
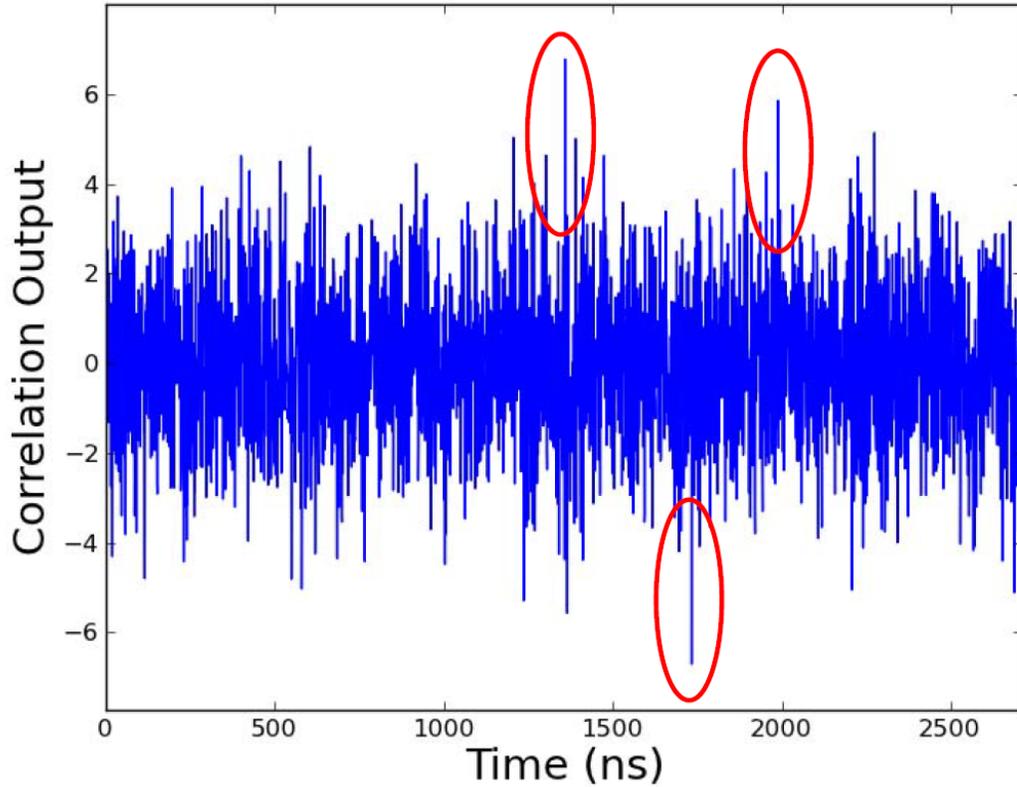
Figure 9:   Matched filter output of an OFDM pulse reflecting off 3 reflectors modeled as Mie scatterers with Swerling noise.

targets in $\mathbf{t}_k$ are stored as part of the EKF state. Our goal is to pair observations in $\mathbf{z}_k$ to a subset of the tracks in $\mathbf{t}_k$. Let N_p be the maximum number of pairings possible and $C(\zeta, \Gamma)$ be the cost function for pairing track ζ with observation Γ. Then we need to find a set of tracks

$$\mathcal{T}_k = \{\mathcal{T}_k(1), \mathcal{T}_k(2), \ldots, \mathcal{T}_k(N_p)\} \subseteq \mathbf{t}_k \tag{68}$$

and a mapping

$$g_k : \mathcal{T}_k \to \mathbf{z}_k \tag{69}$$
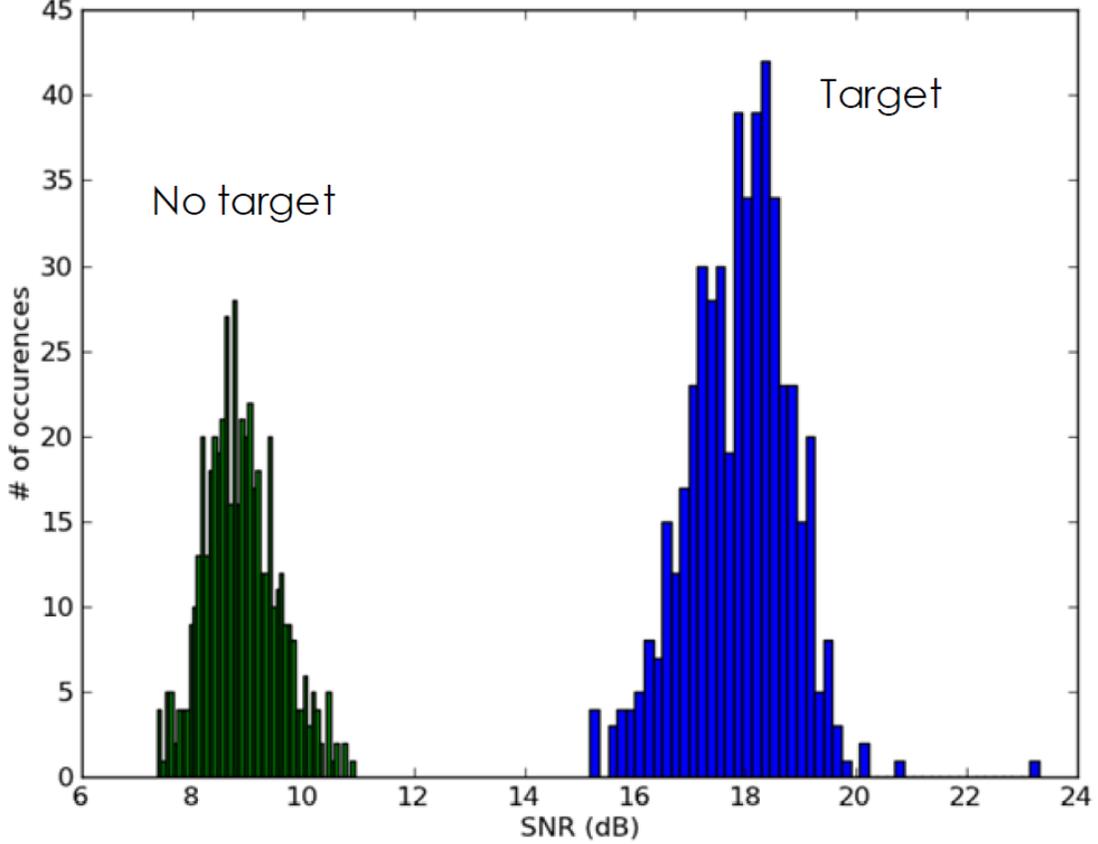
Figure 10: MF SNR histogram for target and no target scenarios with true MF SNR of 18dB.

such that our choices of $g_k(\cdot)$ and $\mathcal{T}_k$ are optimal in the sense of minimizing the direct sum cost function:

$$C'(\mathcal{T}_k, g_k) = \sum_{i=1}^{N_p} C(\mathcal{T}_k(i), g_k(\mathcal{T}_k(i))) \tag{70}$$

Using the EKF computed covariance of previous tracks, we can use the Mahalanobis distance as the pairing cost function. For a diagonal covariance matrix, this would be

$$C(\mathcal{T}_k(i), g_k(i)) = \frac{(\tilde{d}(\mathcal{T}_k(i)) - \tilde{d}'(g_k(i)))^2}{\sigma_{\mathcal{T}_k}^2} \tag{71}$$
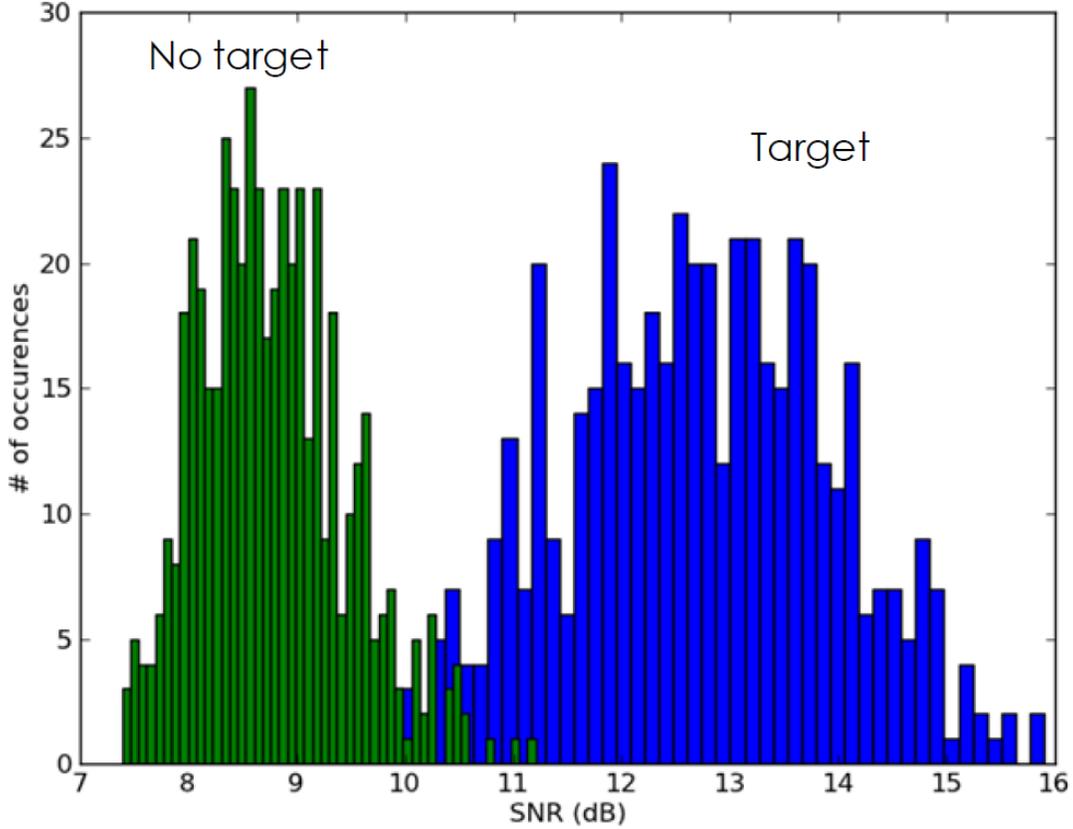
Figure 11: MF SNR histogram for target and no target scenarios with true MF SNR of 13dB.

where $\tilde{d}(\cdot)$ and $\tilde{d}'(\cdot)$ are the estimated range between the AV and the track and observation respectively. For observations, this range is calculated directly from the MF SNR peak location. For tracks, this is calculated via the Euclidean distance between the track position and the best available estimate of the AV position. This calculation can be done by evaluating the function $\hat{d}(\cdot)$ from Section 3.8. Let T_r be the time the observation was received. Then

$$\tilde{d}(\varsigma) = \hat{d}(m_t(\varsigma), T_r) \tag{72}$$

where $m_t(\cdot)$ is a mapping from a track to the estimated track position inside the filter. For confirmed tracks (described in the next section), there is a one to one correspondence between the tracks in $\mathbf{t}_k$ and the EKF states $\mathbf{T}_1(T_r), \mathbf{T}_2(T_r), \ldots$, as

48
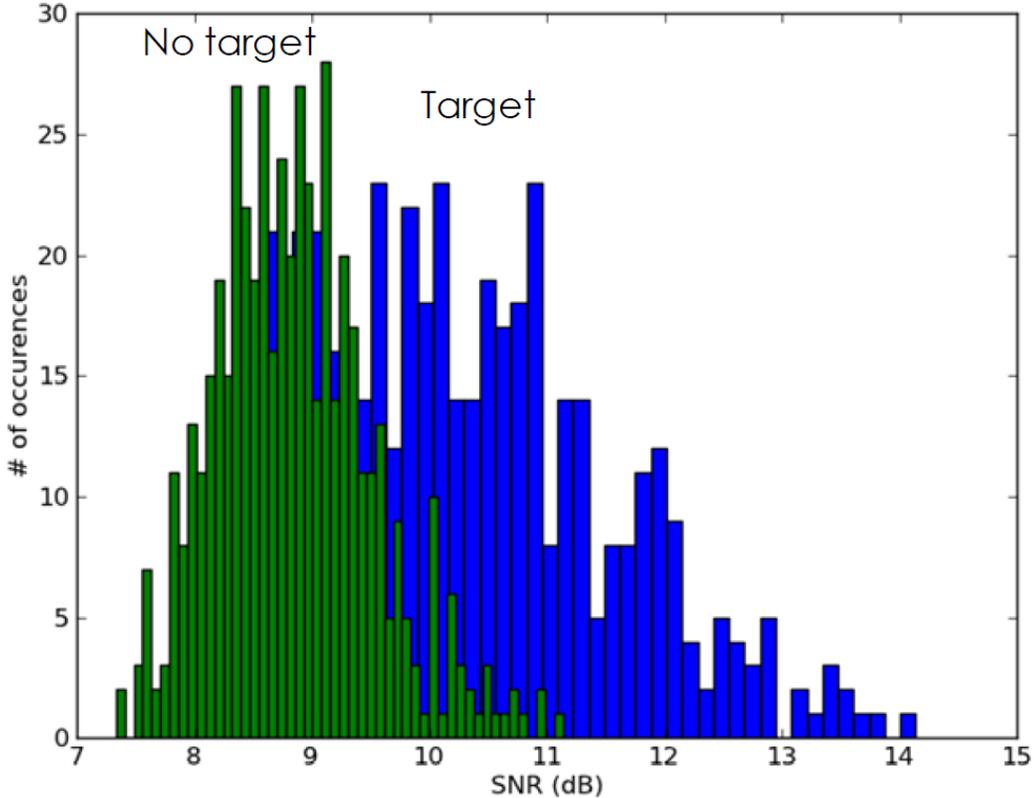
Figure 12:   MF SNR histogram for target and no target scenarios with true MF SNR of 10dB.

described in Section 3.7.1. For unconfirmed tracks, the track position estimate is stored in a separate state vector, which is recorded until the unconfirmed track's deletion.

Note that the algorithm as described above would have the undesired conse-quence of every track receiving an observation if enough are available, due to N_p being fixed. In practice, gating of pairings allowed in (69) is necessary to prevent pairings significantly distant in terms of (71).

3.4.2.2   M/N Detector.   The set of unpaired tracks at position k is denoted

$$\mathcal{T}'_k = \mathbf{t}_k - \mathcal{T}_k \tag{73}$$
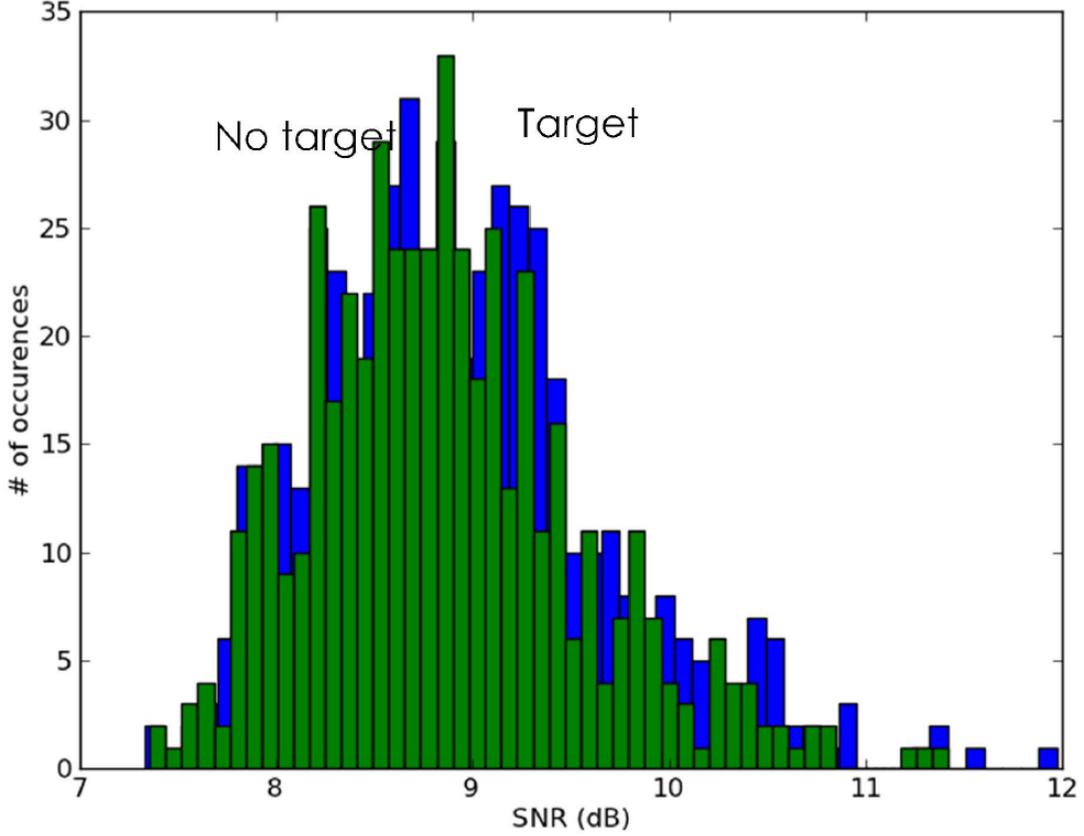
Figure 13:   MF SNR histogram for target and no target scenarios with true MF SNR of 9dB.

The set of unpaired observations at position k is denoted

$$\mathbf{z}'_k = \mathbf{z}_k - \{g_k(\mathcal{T}_k(i))\}, \quad \forall \mathcal{T}_k(i) \in \mathcal{T}_k \tag{74}$$

The set of tracks available at position k is given by

$$\mathbf{t}_{k+1} = \mathbf{t}_k + \mathbf{t}_{k,\text{add}} - \mathbf{t}_{k,\text{remove}} \tag{75}$$

where $\mathbf{t}_{k,\text{add}}$ is the set of new reflectors to track and $\mathbf{t}_{k,\text{remove}}$ is the set of reflector tracks to delete at time k. A reflector is deleted at time position k if it has appeared in $\mathcal{T}'_k$ M out of the last N positions, where the parameters M, N are tunable to achieve a desired CFAR. All unconfirmed observations $\mathbf{z}'_k$ are added to the next set of tracks:

50

$$\mathbf{t}_{k,\text{add}} = \mathbf{z}'_k \tag{76}$$

These tracks are added in order to make them available to the association algorithm. However, they are marked as being unconfirmed tracks and not added to the set of EKF tracks. If the data association algorithm pairs new observations to an unconfirmed track in M out of N sequential collections, the track is confirmed and initialized into the EKF solution, as described in Section 3.9. Unconfirmed tracks are subject to deletion via (75) just as confirmed tracks are.

3.4.3   Stochastic Extensions.   The deterministic algorithms discussed in this section require 1D spectral processing of each data collection. We found the computational requirements for this processing to be feasible for real-time simulation on a desktop PC with GPU parallelization. However, in certain scenarios such as embedded platforms with limited computational resources, the computational burden may need to be further reduced. In this case, a subset of the raw radar data may be used to provide real-time navigation updates.

The matched filter in (66) requires $O(n \log n)$ computation for a data collection of n samples. If the processing platform is unable to compute the MF within the pulse repetition interval (PRI), then the MF can instead be computed every N_{skip} data collections. Initial detection can occur during any data collection when the reflector is illuminated by the beam. Once detected, the path of the reflector can be tracked backwards in time using the tracking method detailed next.

The tracking of targets also relies on the MF output. However, given the estimated position of a track we can perform a local search in the expected range bins for an observation which can be associated with the track. This modification allows us to forgo the FFT-based MF described in (66), and instead perform a direct computation of the needed bins:

$$\mathbf{m}'_k(l) = \sum_{i=0}^{2D_{av}f_s/c} |S_{rx}(k,0)||S_{ref}(l)| \tag{77}$$

where $|S_{rx}(k,0)|$ is zero padded as necessary. This method has an asymptotic complexity of $O(n)$ for a single point computation, and only outperforms the FFT based computation when the number of points to be evaluated is less than $\log n$. Fig. 14 illustrates this process. The MF outputs calculated are lightly shaded. We see that a full correlation is performed in order to find a peak (this occurs ever N_{skip} collections). From this point on, a local search is performed in the local area around the predicted target location. In this way, the vast majority of the data points can remain uncalculated (the white rectangular areas), which reduces the computational burden. This method is ideal for usage in this study, as the navigation problem only requires a sparse set of reflectors in order to achieve a position solution. The stochastic methods described here are therefore used for all simulations in this study.

3.5   INS Error Model

Strapdown INS systems are composed of two primary sensors types, accelerometers and gyros. The accelerometers produce a measurement of specific force, which is a measurement of the forces acting upon the INS computed in the accelerating frame b'. There are many sophisticated INS models which have been developed to accurately model certain types and implementations of INS [1, 49, 50].

In order for the methods developed here to be applicable to a wide range of applications, a simple model based on the model given in [13] is used. The gyroscope and accelerometer both have an additive noise source and a bias. The gyroscope and accelerometer biases are modelled as FOGM processes which are i.i.d. with respect to all axes. Let the gyroscope bias time constant and noise strength be denoted τ_a and σ_a respectively. Let the accelerometer bias time constant and noise strength be τ_b and σ_b respectively. The gyroscope and accelerometer noise sources are modelled as a random walk. Let the gyroscope and accelerometer random walk noise strength
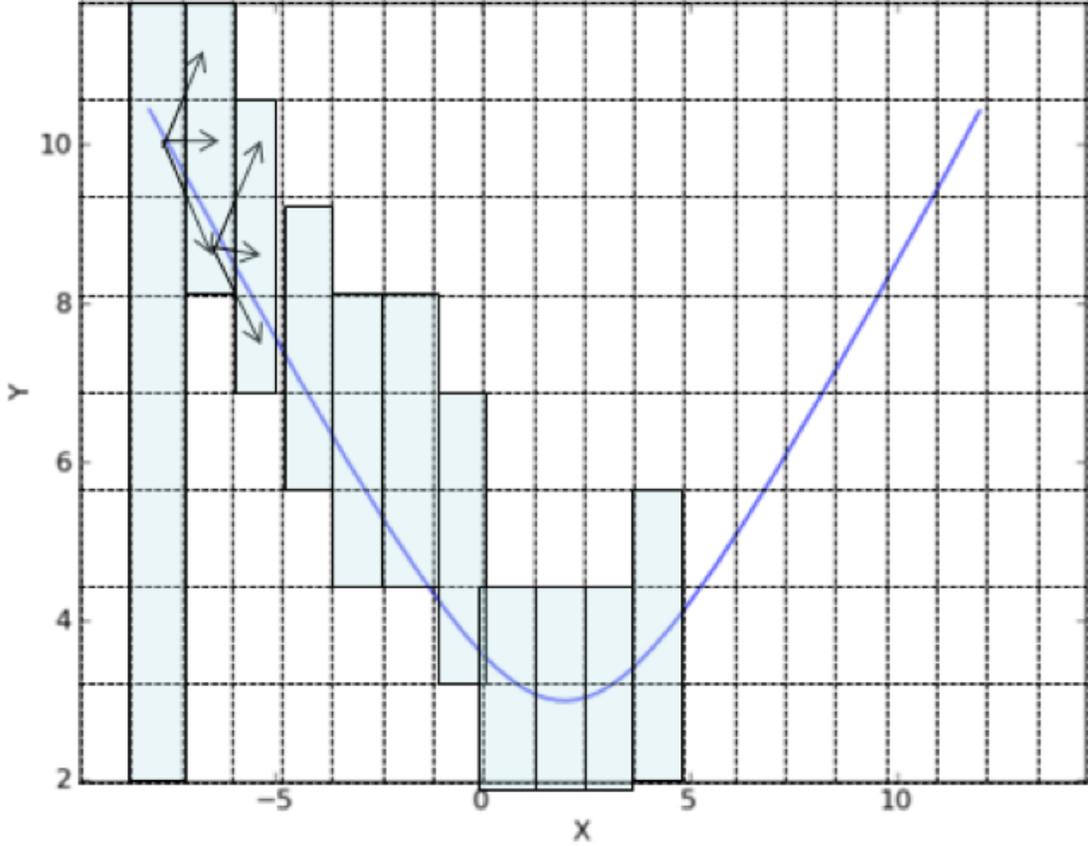
Figure 14:   Stochastic exploration of large SAR data sets.

be denoted $\sigma_{a_{rw}}$ and $\sigma_{b_{rw}}$ respectively. Table 2 shows the actual parameters used for simulations later in this dissertation, for each INS grade simulated.

Table 2:   Parameters used in simulation for different INS Grades. All INS Grades use a time constant of $\tau_a = \tau_b = 3600$ seconds.

INS Grade	σ_a (rad/s)	σ_b (m/s^2)	$\sigma_{a_{rw}}$ (rad/s$^{1/2}$)	$\sigma_{b_{rw}}$ (m/s$^{3/2}$)
Commercial (Cloudcap Crista)	8.7e-3	1.96e-1	6.5e-4	4.3e-3
Tactical (HG1700)	4.8e-6	9.8e-3	8.7e-5	9.5e-3
Nagivation (HG 9900 - H764G)	7.2e-9	2.45e-4	5.8e-7	2.3e-4

3.6   Direct-State Navigation Filter Dynamics Model

In this section we pose the problem of navigation with a stand-alone radar system in such a way that we can apply the EKF equations of 2.4.9. We assume tracked radar reflectors with associated range/Doppler observations are available (as described in Section 3.4) to be used as measurements.

3.6.1   State Model.   This navigation filter uses a direct-state model where we directly model the quantities we are interested in knowing as states, along with auxiliary states necessary for the determination of the desired quantities. The state vector $\mathbf{a}$ contains the AV position $\mathbf{p}$, velocity $\dot{\mathbf{p}}$, and acceleration $\ddot{\mathbf{p}}$, each in three dimensions. We are have states for each radar target $\mathbf{T}_1, \mathbf{T}_2, \cdots$, each of which is a tracked and confirmed target from the radar processing algorithms.

$$
\mathbf{a}(t) = \begin{bmatrix} \mathbf{p}(t) \\ \dot{\mathbf{p}}(t) \\ \ddot{\mathbf{p}}(t) \\ \mathbf{T}_1(t) \\ \mathbf{T}_2(t) \\ \vdots \end{bmatrix} \tag{78}
$$

3.6.2   Dynamics Model.   Our system dynamics are given by

$$
\dot{\mathbf{a}}(t) = \mathbf{F}\mathbf{a}(t) + \mathbf{w}_{\mathbf{a}}(t) \tag{79}
$$

Our dynamics matrix $\mathbf{F}$ attempts to capture the relationship between the states as well as how they propagate over time. Using our knowledge of the physics we can compute the nominal $\dot{\mathbf{a}}$:

$$\dot{\mathbf{a}}(t) = \frac{\partial \mathbf{a}(t)}{\partial t} = \frac{\partial}{\partial t} \begin{bmatrix} \mathbf{p}(t) \\ \dot{\mathbf{p}}(t) \\ \ddot{\mathbf{p}}(t) \\ \mathbf{T}_1(t) \\ \mathbf{T}_2(t) \\ \vdots \end{bmatrix} = \begin{bmatrix} \frac{\partial}{\partial t}\mathbf{p}(t) \\ \frac{\partial}{\partial t}\dot{\mathbf{p}}(t) \\ \frac{\partial}{\partial t}\ddot{\mathbf{p}}(t) \\ \frac{\partial}{\partial t}\mathbf{T}_1(t) \\ \frac{\partial}{\partial t}\mathbf{T}_2(t) \\ \vdots \end{bmatrix} = \begin{bmatrix} \dot{\mathbf{p}}(t) \\ \ddot{\mathbf{p}}(t) \\ \dddot{\mathbf{p}}(t) \\ \frac{\partial}{\partial t}\mathbf{T}_1(t) \\ \frac{\partial}{\partial t}\mathbf{T}_2(t) \\ \vdots \end{bmatrix} \tag{80}$$

Expanding Eq. (79) we have

$$\begin{bmatrix} \dot{\mathbf{p}}(t) \\ \ddot{\mathbf{p}}(t) \\ \dddot{\mathbf{p}}(t) \\ \frac{\partial}{\partial t}\mathbf{T}_1(t) \\ \frac{\partial}{\partial t}\mathbf{T}_2(t) \\ \vdots \end{bmatrix} = \mathbf{F} \begin{bmatrix} \mathbf{p}(t) \\ \dot{\mathbf{p}}(t) \\ \ddot{\mathbf{p}}(t) \\ \mathbf{T}_1(t) \\ \mathbf{T}_2(t) \\ \vdots \end{bmatrix} + \mathbf{w_a}(t) \tag{81}$$

From this representation we can determine the terms in $\mathbf{F}$. For example, obviously $\dot{\mathbf{p}}(t) = \dot{\mathbf{p}}(t)$, we can set the corresponding terms $(\mathbf{F}_{14}, \mathbf{F}_{25}, \mathbf{F}_{26})$ to 1. Since we assume the target positions are constant (i.e. stationary targets), we can set all the terms corresponding to them to zero. Thus we have

$$\mathbf{F} = \begin{bmatrix} \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & -1/T_a\mathbf{I} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix} \tag{82}$$

where $\mathbf{I}$ is a 3×3 identity matrix and $\mathbf{0}$ is a 3×3 matrix of zeros. The $-1/T_a$ is present in order to constrain the wander of the vehicle acceleration. In general, terrestrial

vehicles cannot have acceleration wander off to infinity. Thus this term continually draws the acceleration back towards zero. Combined with the appropriate noise term in $\mathbf{w_a}(t)$, this approach is called a first-order Gauss-Markov acceleration (FOGMA) model, and is commonly used to allow dynamic acceleration constrained to realistic motion dynamics [20, 45]. In particular, we want

$$\mathbf{w_a}(t) = \begin{bmatrix} 0 \\ 0 \\ \sigma_{FOGMA}(t) \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{83}$$

where σ_{FOGMA} defines the strength of the FOGM acceleration noise.

In order to estimate the states via an EKF, we need a discrete time dynamics model (see Section 2.4.9). Thus we want to find an equivalent discrete time dynamics model such that

$$\mathbf{a}_k = \mathbf{\Phi}\mathbf{a}_{k-1} \tag{84}$$

where $\mathbf{a}_k = \mathbf{a}(t_k)$ for $k \in 1, 2, \ldots$.The state transition matrix $\mathbf{\Phi}$ can be calculated from $\mathbf{F}$ [20]:

$$\mathbf{\Phi}(\Delta t) = e^{\mathbf{F}\Delta t} \tag{85}$$

where Δt is the propagation time, which is set equal to Δt_p in this dissertation.

Let

$$\mathbf{Q_a}(t) = E[\mathbf{w_a}(t)\mathbf{w_a}(t)^T] \tag{86}$$

then we can transform $\mathbf{w_a}$ into a discrete-time random process by calculating

$$\mathbf{Q_a}_d(k) = \int_0^{\Delta t} \boldsymbol{\Phi}(\zeta)\mathbf{Q_a}(k\Delta t_p)\boldsymbol{\Phi}^T(\zeta)d\zeta \tag{87}$$

Our discrete-time dynamics model is then

$$\mathbf{a}_k = \mathbf{a}_{(k-1)}\boldsymbol{\Phi}(\Delta t_p) + \mathbf{w_a}_d(k) \tag{88}$$

where $E[\mathbf{w_a}_d\mathbf{w_a}_d^T] = \mathbf{Q_a}_d$.

3.7   Error-State Navigation Filter Dynamics Model

In this section we pose the problem of navigation with a combined INS/radar system in such a way that we can apply the EKF equations of Section 2.4.9. We assume tracked radar reflectors with associated range/Doppler observations are available (as described in Section 3.4) to be used as measurements.

In contrast to Section 3.6, we propose here to model the AV position via error-states instead of direct-states (i.e. instead of having states to keep track of the AV's position, we have states to keep track of the error in the INS's estimate of the AV's position). This approach is useful as it greatly simplifies the dynamics modeling problem. In Section 3.6, we assumed the vehicle dynamics were a FOGM process. In most realistic scenarios a better model than this could be used. As an example, consider an AV in a coordinated turn with control inputs indicating the pilot intends such a maneuver. In this case, we could use both knowledge of the control inputs (via an extra term in (79) and the typical error in that particular maneuver to inform our dynamics model.

However, this approach is quite complicated; not only must we know about the varying characteristics of our platform (plane types, car types, etc.), but we must also actively attempt to understand the maneuver and operation mode the vehicle is in and provide integration with control systems being used. The error-state model requires

57

none of this: we simply ignore the vehicle position and focus on estimating the INS error, which is primarily a function of the particular INS chosen. Although the effects of the vehicle maneuvering can have an impact on the INS error (INSs exhibit different error for different operating conditions), these affects are typically second-order [1]. Thus the INS error state model is used in this section for INS integrated navigation.

3.7.1 State Model. This navigation filter uses an INS error state model. Thus the state vector **a** contains the INS position error $\delta\mathbf{p}$, velocity error $\delta\dot{\mathbf{p}}$, tilt error $\boldsymbol{\psi}$, accelerometer bias $\delta\mathbf{b}_a$ and gyroscope bias $\delta\mathbf{b}_b$, each of which are in three dimensions. The state vector also contains estimates of tracked and confirmed target positions $\mathbf{T}_1, \mathbf{T}_2, \cdots$ in three dimensions. The state vector is then

$$\mathbf{a}(t) = \begin{bmatrix} \delta\mathbf{p}(t) \\ \delta\dot{\mathbf{p}}(t) \\ \boldsymbol{\psi}(t) \\ \delta\mathbf{b}_a(t) \\ \delta\mathbf{b}_b(t) \\ \mathbf{T}_1(t) \\ \mathbf{T}_2(t) \\ \vdots \end{bmatrix} \tag{89}$$

The position $\mathbf{T}_k(l\Delta t_p)$ directly maps to the tracked targets in $\mathbf{t}_k$ used in the tracking algorithm, such that $\mathbf{T}_k(l\Delta t_p)$ is the position of the kth target in $\mathbf{t}_l$.

The tilt errors are assumed to be small angle errors in attitude, defined by [13]

$$\hat{\mathbf{C}}_b^n = [\mathbf{I} - (\psi\times)]\mathbf{C}_b^n \tag{90}$$

where $\hat{\mathbf{C}}_b^n$ is the INS-estimated direction cosine matrix from the body to the navigation frame, $\mathbf{C}_b^n$ is the true value of the same quantity, and $\psi\times$ is the skew-symmetric form of ψ [13].

3.7.2 Dynamics Model.   Our system dynamics are given by

$$\dot{\mathbf{a}}(t) = \mathbf{F}\mathbf{a}(t) + \mathbf{w_a}(t) \tag{91}$$

where $\mathbf{F}$ is the continuous-time dynamics and $\mathbf{w_a}$ is a column vector of noise inputs. $\mathbf{F}$ can be written as

$$\mathbf{F} = \begin{bmatrix} \mathbf{F}_{INS} & \mathbf{0} \\ \mathbf{0} & \mathbf{F}_{radar} \end{bmatrix} \tag{92}$$

where $\mathbf{F}_{INS}$ is a 15x15 matrix and $\mathbf{F}_{radar}$ is a $3N_{tt}$x$3N_{tt}$ matrix, and N_{tt} is the number of tracked targets at the current time. Since the reflectors are modeled as stationary, we have

$$\mathbf{F}_{radar} = \mathbf{0} \tag{93}$$

From [13] we can write

$$\mathbf{F}_{INS} = \begin{bmatrix} \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{C}_e^n \mathbf{G} \mathbf{C}_n^e & -2\mathbf{C}_e^n \mathbf{\Omega}_{ie}^n \mathbf{C}_n^e & (\mathbf{f}^n \times) & \mathbf{C}_b^n & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & -(\mathbf{C}_e^n \omega_{ie}^e) \times & \mathbf{0} & -\mathbf{C}_b^n \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & -\frac{1}{\tau_a}\mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & -\frac{1}{\tau_b}\mathbf{I} \end{bmatrix} \tag{94}$$

Each term is a 3x3 matrix: ω_{ie}^e is the angular rate vector between the i and e frames expressed in the e frame, $\mathbf{\Omega}_{ie}^e$ is the skew symmetric form of ω_{ie}^e, $\mathbf{C}_{(\cdot)}^{(\cdot)}$ is the direction cosine matrix from the bottom frame to the top frame, and $\mathbf{f}^n$ is the specific force represented in the n frame. $(\cdot)\times$ is the skew-symmetric form of the quantity.

We also need to describe the noise sources in (91). Let

$$\mathbf{w_a} = \begin{bmatrix} \mathbf{w}_{INS} \\ \mathbf{w}_{radar} \end{bmatrix} \tag{95}$$

Then $\mathbf{w}_{radar} = \mathbf{0}$ and $\mathbf{w}_{INS}$ is [13]

$$\mathbf{w}_{INS} = \begin{bmatrix} \mathbf{0} \\ \mathbf{C}_b^n \mathbf{w}_a \\ -\mathbf{C}_b^n \mathbf{w}_b \\ \mathbf{w}_{a_{bias}} \\ \mathbf{w}_{b_{bias}} \end{bmatrix} \tag{96}$$

where $\mathbf{w}_a$, $\mathbf{w}_b$, $\mathbf{w}_{a_{bias}}$, and $\mathbf{w}_{b_{bias}}$ are the gyro random walk, accelerometer random walk, gyro bias, and accelerometer bias random processes respectively. These processes are described in Section 3.5.

3.7.3 Discretization.   The rest of the process of converting this system into a discrete form usable with an EKF is described in Section 3.6.

3.8   *Measurement Model*

The measurements used as updates to the EKF are radar-only. The radar measurements are ranges between the position of the AV and the radar targets.

The measurement model is of the form

$$\mathbf{z}_k = \mathbf{h}(\mathbf{a}_k) + \mathbf{w_z} \tag{97}$$

where $\mathbf{h}(\cdot)$ is given by

$$\mathbf{h}(\mathbf{a}_k) = \begin{bmatrix} d(\mathbf{T}_1, k) \\ d(\mathbf{T}_2, k) \\ \vdots \end{bmatrix} \tag{98}$$

The notation $d(\mathbf{T}_\zeta, k)$ denotes the distance between target $\mathbf{T}_\zeta$ and the AV at time t_k. Finding this distance depends on the type of filter implementation. For the direct-state approach, the AV position and the target position are both states in the EKF state vector and are readily available. For the error-state approach, the AV position is not directly stored in the EKF state vector $\mathbf{a}$. Instead, we must use some information from the state vector combined with external information to fully describe the measurement model:

$$d(\mathbf{T}_\zeta, k) = \sqrt{(\mathbf{T}_\zeta(t_k) - [\hat{\mathbf{p}}(t_k) + \delta\mathbf{p}(t_k)])^T (\mathbf{T}_\zeta(t_k) - [\hat{\mathbf{p}}(t_k) + \delta\mathbf{p}(t_k)])} \qquad (99)$$

where $\hat{\mathbf{p}}(t)$ is the INS estimate of the AV at time t. The bracketed term is the true position:

$$\mathbf{p}(t) = \hat{\mathbf{p}}(t) + \delta\mathbf{p}(t) \qquad (100)$$

The $\delta\mathbf{p}$ and $\mathbf{T}$ terms are all contained in the state vector but the $\hat{\mathbf{p}}$ terms are not, as they must be pulled from an independently calculated INS solution. Thus as previously mentioned the $\mathbf{h}$ function must use information from an external source, and is not solely a function of the state vector as it is in the direct-state EKF implementation.

In order to use an EKF we must also calculate the Jacobian of $\mathbf{h}$ as it is a non-linear function. Thus we have

$$\mathbf{H}(\mathbf{a}_k) = \begin{bmatrix} \gamma_{k,1,1} & \gamma_{k,1,2} & \cdots & \gamma_{k,1,N_{tt}} \\ \gamma_{k,2,1} & \gamma_{k,2,2} & \cdots & \gamma_{k,2,N_{tt}} \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_{k,N_o,1} & \gamma_{k,N_o,2} & \cdots & \gamma_{k,N_o,N_{tt}} \end{bmatrix} \qquad (101)$$

where $\gamma_{k,l,m}$ is the derivative of $\mathbf{h}(\mathbf{a}_k)$'s lth row with respect to the mth state. That is,

$$\gamma_{k,l,m} = \frac{\partial[\mathbf{h}(\mathbf{a}_k)]_l}{\partial[\mathbf{a}_k]_m} \tag{102}$$

Note that because the $\mathbf{h}$ function is not solely a function of the state for the error-state model, the Jacobian of $\mathbf{h}$ is also dependent on the INS solution. Thus for code implementations of an EKF using this model the EKF will be coupled with the INS solution, as it must extract INS position estimates in order to compute the $\mathbf{H}$ matrix. For the results in this study, the Jacobian was evaluated numerically and linearized about the most current estimated INS position solution and state vector.

From experimental data collected in [42], the radar processing method used in this dissertation generated ranges with a standard deviation of 10 meters. Thus the noise strength of $\mathbf{w_z}$ is set as

$$\sigma_{\mathbf{w_z}} = 100 \tag{103}$$

3.9 *Filter Implementation*

The described dynamics/measurement model in this section is now fully described and ready to be estimated. The dynamics model is linear and the measurement model is non-linear, necessitating the use of a non-linear filter. In Chapter IV we will estimate this system with an EKF for various configurations.

There is one standing issue with the implementation of an EKF that is not yet addressed. Each state in the EKF must be initialized with approximately correct values in order for the linearization error due to the Jacobian to be minimized. The initial INS errors are set to be zero with complete certainty at the beginning of the simulation. However, as tracks are confirmed new states must be added and initialized into the EKF state vector, a process which continuously happens as the simulation continues.

Suppose a new track is confirmed at position $\mathbf{p}_k$, and the state vector $\mathbf{a}_k$ has N_{tt} targets being tracked in the EKF. The new track will have M past observations associated with it, which were used by the M/N detector to confirm it. For notational simplicity we will assume the M observations all were collected in the last M propagations. We thus need to perform the initialization at $k - M$, and then re-propagate up to the current time k.

Then the augmented state vector will be

$$\mathbf{a}'_{k-M} = \begin{bmatrix} \mathbf{a}_{k-M} & \mathbf{T}_{new} \end{bmatrix} \tag{104}$$

where $\mathbf{a}'_k$ is a 1x$N_{tt}+3$ vector. Since the new states will initially have no information, we use the inverse covariance form of the EKF during initialization. The augmented information matrix is then

$$[\mathbf{P}'_{\mathbf{a}}]^{-1}(k - M) = \begin{bmatrix} [\mathbf{P}_{\mathbf{a}}]^{-1}(k - M) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \tag{105}$$

where

$$\mathbf{P}'_{\mathbf{a}}(k) = E[\mathbf{a}_{k-M}\mathbf{a}^T_{k-M}] \tag{106}$$

The 3x3 $\mathbf{0}$ matrices in (105) indicate that we have no information initially of the new target. We now perform an update with the first of the M observations. Using the inverse covariance form, we have [45]

$$\tilde{\mathbf{a}}'_{k-M} = \mathbf{a}^{nom}_{k-M} - \left[\mathbf{J}^T\mathbf{R}^{-1}\mathbf{J} + [\mathbf{P}'_{\mathbf{a}}]^{-1}(k - M) \right] \mathbf{J}^T\mathbf{R}^{-1}[\mathbf{h}(\mathbf{a}^{nom}_{k-M}) - \mathbf{z}_{new}] \tag{107}$$

where $\mathbf{J}$ is the Jacobian of $\mathbf{h}$, $\mathbf{R} = E[\mathbf{w_z}\mathbf{w_z}^T]$, and $\mathbf{z}_{new}$ is the observation associated with the new track at $k - M$. $\mathbf{a}^{nom}_{(\cdot)}$ is the nominal state estimate, which is initially

63

set to the estimated target position using Doppler measurements. After calculating $\tilde{\mathbf{a}}'$ via (107), the result is used as the nominal and (107) is recomputed again. This process is repeated iteratively until the calculated nominal converges. This estimate is optimal in the minimum mean squared error (MMSE) sense, as long as the initial Doppler estimate of the target position is accurate enough to minimize linearization error. The covariance matrix associated with the updated state is [45]

$$\tilde{\mathbf{P}}'_{\mathbf{a}}(k - M) = [\mathbf{J}^T \mathbf{R}^{-1} \mathbf{J} + [\mathbf{P}'_{\mathbf{a}}]^{-1}(k - M)]^{-1} \tag{108}$$

The state $\tilde{\mathbf{a}}'_{k-M}$ and state uncertainty $\tilde{\mathbf{P}}'_{\mathbf{a}}(k - M)$ now replaces the original state calculated by the regular EKF at $k - M$. The EKF is then propagated forward as it was before, with the exception that the other $M - 1$ observations associated with the newly added target are now included as updates at the times they were observed. When the filter updates to time k, the old filter is replaced with the repropagated one and the navigation filter continues as normal.

3.10 Non-Bayesian Navigation Filter

Although state-based Bayesian methods are the typical approach to estimation of dynamic navigation problems with quasi-Gaussian measurement errors, an alternative approach is to use classic estimation methods such as LMA. The disadvantage of using the classical approach on a non-linear system is the lack of guarantees of optimality (see Chapter 2.4 for more explanation). However, classical methods tend to perform well in practice, and it is instructive to consider them as an alternative to typical navigation filters.

Our problem can be described by a set of non-linear equations. Namely, one equation for each range to each target at each time of data collection, i.e.

$$
\begin{bmatrix}
\sqrt{(\mathbf{p}(t_1) - \mathbf{T}_1(t_1))^T(\mathbf{p}(t_1) - \mathbf{T}_1(t_1))} = d(\mathbf{T}_1, t_1) \\
\sqrt{(\mathbf{p}(t_2) - \mathbf{T}_1(t_2))^T(\mathbf{p}(t_2) - \mathbf{T}_1(t_2))} = d(\mathbf{T}_1, t_2) \\
\sqrt{(\mathbf{p}(t_3) - \mathbf{T}_1(t_3))^T(\mathbf{p}(t_3) - \mathbf{T}_1(t_3))} = d(\mathbf{T}_1, t_3) \\
\vdots \\
\sqrt{(\mathbf{p}(t_1) - \mathbf{T}_2(t_1))^T(\mathbf{p}(t_1) - \mathbf{T}_2(t_1))} = d(\mathbf{T}_2, t_1) \\
\sqrt{(\mathbf{p}(t_2) - \mathbf{T}_2(t_2))^T(\mathbf{p}(t_2) - \mathbf{T}_2(t_2))} = d(\mathbf{T}_2, t_2) \\
\sqrt{(\mathbf{p}(t_3) - \mathbf{T}_2(t_3))^T(\mathbf{p}(t_3) - \mathbf{T}_2(t_1))} = d(\mathbf{T}_2, t_3) \\
\vdots
\end{bmatrix}
\tag{109}
$$

where again $d(\mathbf{T}_1, t_1)$ denotes the distance between $\mathbf{T}_1$ and the vehicle at time t_1. For a realistic problem, we will not have all of the ranges listed above (i.e. $d(\mathbf{T}_1, t_1), d(\mathbf{T}_1, t_2)$, $\ldots, d(\mathbf{T}_2, t_1), d(\mathbf{T}_2, t_2), \ldots$) available from our sensors; this would require us to observe all targets at all times. As targets go by, we will see some ranges to a target while the vehicle is near that target. Thus we will only have a subset of these equations available for navigation.

We then want find a solution for the unknowns in these equations. If we assume the targets are non-moving:

$$
\mathbf{T}_k(t_m) = \mathbf{T}_k(t_n), \forall k, m, n
\tag{110}
$$

then our unknowns are the set of target positions $\mathbf{T}_1(0), \mathbf{T}_2(0), \ldots$ and the vehicle position at each time epoch $\mathbf{p}(t_0), \mathbf{p}(t_1), \ldots$. The right-hand side of each equation is known, as it is the range measurement from the sensor. The problem of navigation is thus equivalent to solving this system of non-linear equations, which can be performed via LMA or another solver.

IV. Simulation Results

In this chapter, we provide simulated performance results for the methods discussed in the previous chapter in airborne scenarios. For indoor navigation results, see the experimental results in Section 5.1. We begin by analyzing the performance of the radar processing algorithms on simulated data. We then compute navigation solutions for various geometric configurations, estimation methods, and sensor availability.

4.1 Radar Signal Processing Results

The simulations in this section use the parameters defined in Table 3. Fig. 15 shows an example MF output for three reflectors positioned at (0,0,0), (100,100,-10), and (300,300,-10) with an SNR of 10dB, $\phi_c = 0$, and $\Delta t_p = 1/\dot{\mathbf{p}}(0)$. We see that the features are easily detected and tracked at this SNR. In addition, large parts of the space of MF data contains no information, which is the motivation for using stochastic extensions in the previous section. With three reflectors present, the stochastic tracking algorithm can evaluate a local search space around the paths in Fig. 15 and perform much better computationally than a full FFT-based MF implementation. Fig. 16 shows the same scenario with an SNR of -10dB and $\phi_c = 0.01$rads. The modulation and thermal noise makes it difficult to detect peaks during any single detection. However, the M/N detector with $M = 30$ and $N = 100$ can still operate under these conditions (MF SNR of 10dB).

Fig. 17 shows the EKF calculated positions of a set of 10 reflectors randomly distributed 1km east of the AV. The simulation used an SNR of -10dB, $\phi_c = 0.01$rads, $N_{skip} = 10$, and $\Delta t_p = 10$m. We see that 5 reflectors are tracked and have received good position estimates. Fig. 18 shows the same setup but with 50 reflectors. We see that the the error in the $\mathbf{x}$ direction grows as a function of the reflector distance from the AV initial flight path, which is the expected result from examination of the measurement model Jacobian. The large quantity of missed reflectors is due to $N_{skip} = 10$ and local search space tracking. However, the the number of tracked scatterers is still sufficient for navigation, as seen in the next section.

Table 3: Simulation Parameters	
$2D_{av}f_s/c$	$100,000$
UWB pulse width $(1/\Delta f)$	128ns
OFDM channel configuration (η_k)	$\sim^{iid} N(0,1)$
Mie scattering sphere radius	0.5m
INS Grade	Commercial
Fast-time sample rate $(1/\Delta t = f_s)$	1 GHz
SNR	0 dB
Number of collection points (N_{av})	25
Kalman propagation time (Δt)	1 s
Initial position of AV	(0,0,0) m
Initial velocity of AV	(200,0,0) m
Initial acceleration of AV	(0,0,0) m

4.2 Classical Estimation

In this section we use the classic estimation approach described in Section 3.10 to navigate via the LMA solver discussed in Section 2.4.6. For simplicity, point targets (i.e., perfect reflections) are assumed for the radar targets in lieu of the development in Section 3.4. The AV model and scenario described in Sections 3.2 and 3.3 respectively are used, with the exception that the AV and targets are constrained to a 2D plane and the AV moves along a fixed axis (the X axis). Fig. 19 shows the overall approach used. An initial LMA is given the truth to initialize the positions of the AV and initial targets (we may assume this is provided from another source at startup). As the AV moves forward, LMA is used to solve the system of non-linear equations as seen in Section 3.10 to find the new AV position. LMA is then used to update the estimates of the positions of the reflectors. This process repeats iteratively, going back and forth between estimating the AV location and the targets' locations, using the ranges from the radar processing algorithm (STM) when available.

Note that this approach does not use information about the vehicle dynamics in any form, and thus it is not claimed to be an optimal estimation method compared to
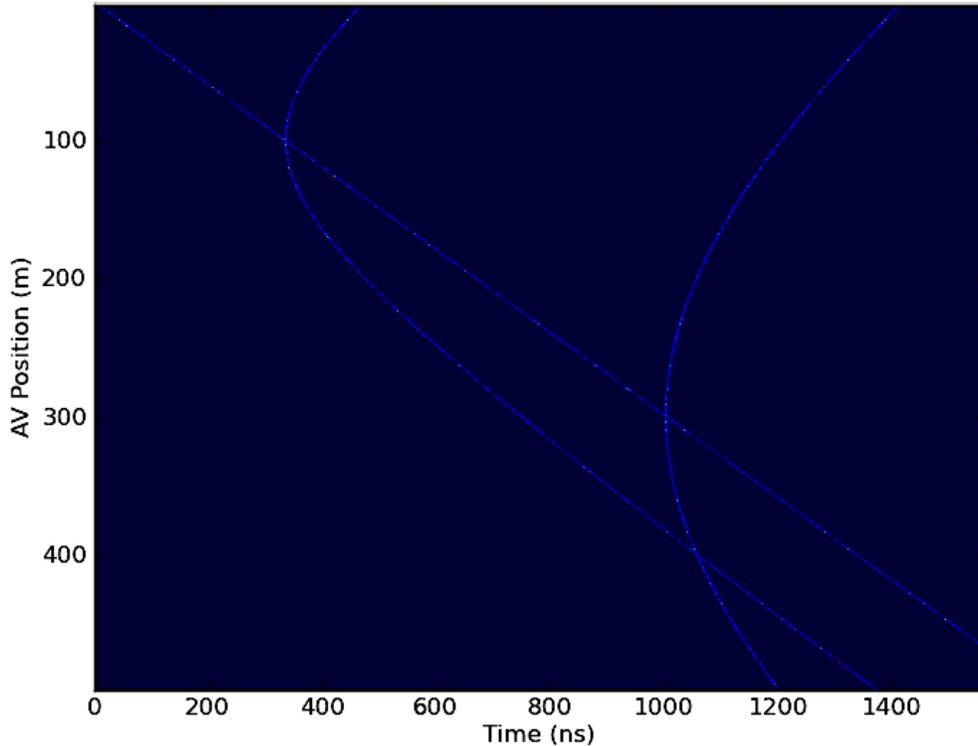
Figure 15: Example set of MF outputs for 3 reflectors located at (0,0,0), (100,100,-10), and (300,300,-10), with SNR=10 dB

a more complex EKF-based approach. However, the LMA-based approach is simple, able to adapt to a wider set of problems, and is instructional in understanding the error sources inherent in UWB-OFDM range-based navigation.

4.2.1 LMA Error Analysis.   LMA is used to calculate the positions of the reflectors and the AV platform. LMA errors are critical contributors of the solution errors. We shall analyze the LMA error for an example scenario: The AV flies to the points $(0\text{m}, 0\text{m}), (10\text{m}, 0\text{m}), \ldots, (0\text{m}, 100\text{m})$, and a single reflector is placed in view of the radar.

We first examine the LMA computed scatter position error for a single reflector in white noise with strength σ_w. Fig. 20 and 21 show a histogram of the calculated reflector location over 4,000 trials for $\sigma_w = 1\text{m}$ and $\sigma_w = 20\text{m}$ respectively when the
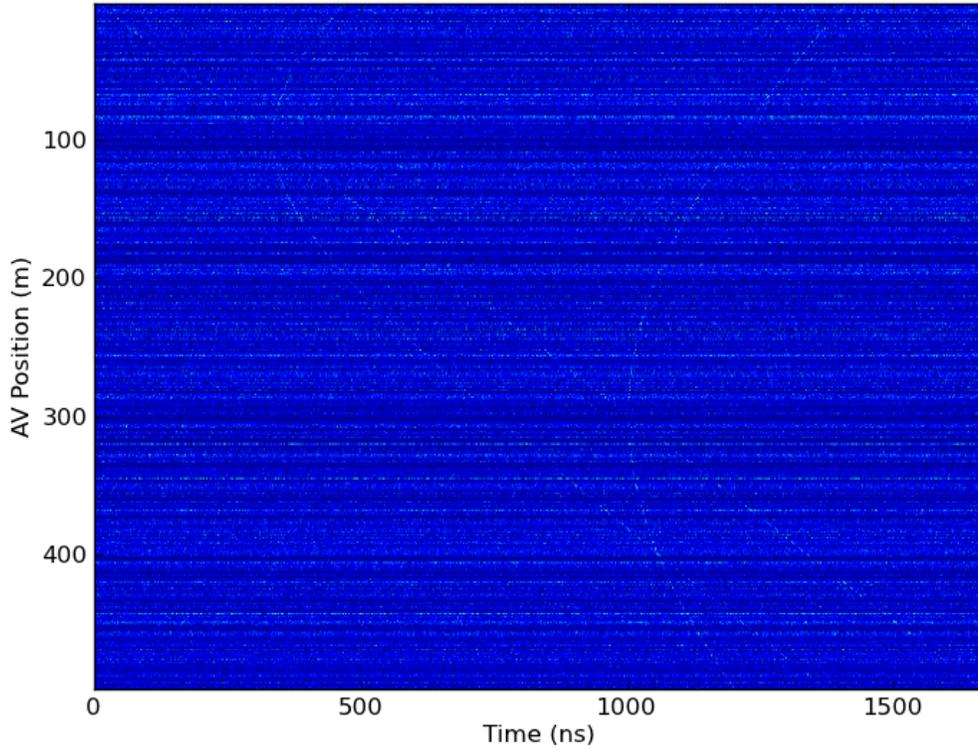
Figure 16: Example set of MF outputs for 3 reflectors located at (0,0,0), (100,100,-10), and (300,300,-10), with SNR=0 dB

reflector is located at the point $(100\text{m}, 200\text{m})$. We see that with zero-mean AWGN, an increase to the noise variance causes the mean position calculated to drop and mean error to become non-zero. This is due to the non-linearity of the system of equations, which causes LMA to be a a biased estimator. Since the mean error is not zero for high σ_w, we expect our positions solution error mean to be non-zero for high σ_w.

Fig. 22 shows the mean and standard deviation of a target as a function of σ_w which was allowed to vary from 0 to 50. We see the non-linear downward mean position solution trend and the more negative skew as σ_w increases. The standard deviation of the solution also increases as σ_w increases, as expected. For $\sigma_w = 50\text{m}$, the standard deviation is about 60m.
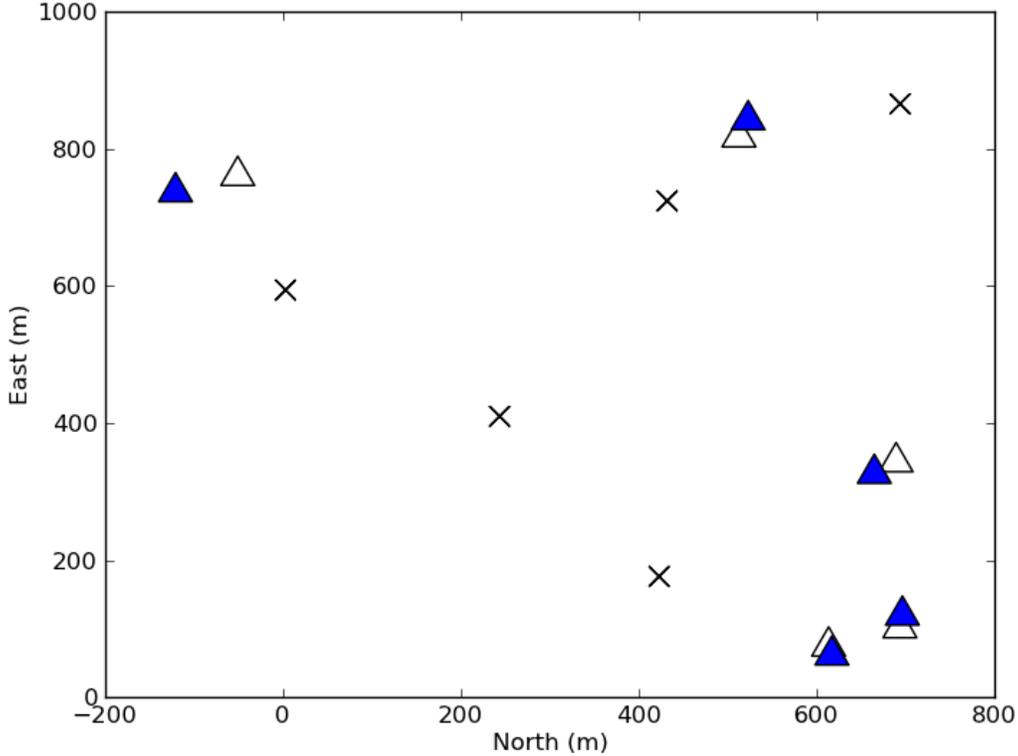
Figure 17: Reflector positions with 10 reflectors available. The X's are missed reflectors, the hollow triangles are the true locations of tracked reflectors, and the solid triangles are the EKF computed estimates of the tracked reflectors.

The LMA error is also dependent on the reflector position. Let $\sigma_w = 1$m. Fig. 23 and 24 show a sweep of the target's X coordinate and the calculated values of a target's X and Y coordinates respectively. We see that as the true target's X coordinate is more distant from the AVs X coordinate, the calculated error of the target's Y coordinate grows. By contrast, the change does not affect the error in the calculated target's X coordinate. Note that in Fig. 24 the actual position is subtracted from the calculated position. This is to remove the linear trend due to the swept variable. This is the expected result, and is the result of GDOP as explained in Section 2.3.

Next, we examine the impact of the reflector Y coordinate on its position estimation. Let the target's X coordinate be 50m. Fig. 25 and 26 show the calculated
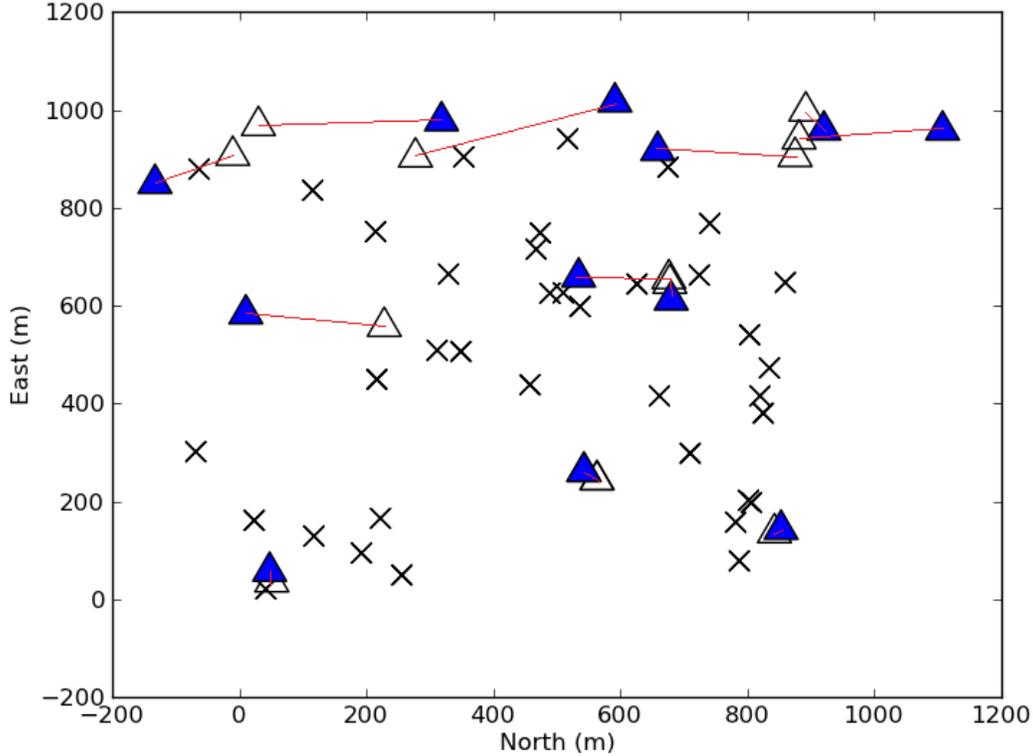
Figure 18: Reflector positions with 50 reflectors available. The X's are missed reflectors, the hollow triangles are the true locations of tracked reflectors, and the solid triangles are the EKF computed estimates of the tracked reflectors. The lines show the which true targets generated the EKF computed estimates.

reflector y and x coordinates as functions of the target's true Y coordinate respectively. Note in Fig. 25, the computed reflector X coordinate standard deviation increases almost linearly as the reflectors Y coordinate value increases, while Fig. 26 indicates that the reflectors Y coordinates error remains steady as the reflector y coordinate increases.

4.2.2 Navigation Solution Error Analysis. The full AV platform navigation solution was simulated using the iterative approach outlined in Fig. 19. The AV travels to each location in $(0m, 0m), (10m, 0m), \ldots, (0m, 1km)$, the targets are in a line 200m from the AV and spaced 10m apart from each other. Fig. 27 and 28 show the mean and standard deviations of the AV positions for this configuration. As

71
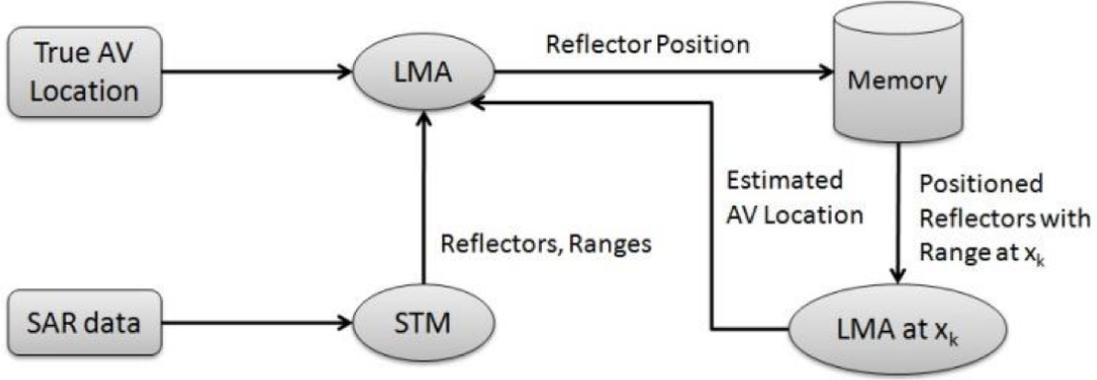
Figure 19:   Flowchart of iterative algorithm to track AV using classical estimation.

expected from our analysis of LMA, the higher σ_w causes the mean position to shift upwards linearly. The standard deviation is also a function of the noise, but even with $\sigma_w = 20$, the solution standard deviation was within 20m after 1km of flight.

Fig. 29 and 30 are generated using the same configuration as that of Fig. 27 and 28 except the targets are separated by 20m. The reflectors are sparser in this scenario which causes a higher standard deviation at each noise level. As the noise increases the standard deviation is affected more, with more than 30m of error when $\sigma_w = 20$. In addition, the mean error is magnified, with the $\sigma_w = 20$ case at 1km having a mean position of 1086m, or an 86m bias. Clearly, the mean error is a function of both σ_w and the number of reflectors available.

Finally, we consider the case where the reflectors are located farther away. Let the targets be in a line 700m away from the AV and still spaced 20m apart. Fig. 31 and 32 show the mean and standard deviation for this configuration. For high noise, the standard deviation grows to 70m when $\sigma_w = 10$. This is due to the distance of the reflectors limiting the number of available ranges per reflector, as the reflection strength follows the inverse square law. Increasing the number of reflectors decreases this error, as it did with the reflectors at 200m.
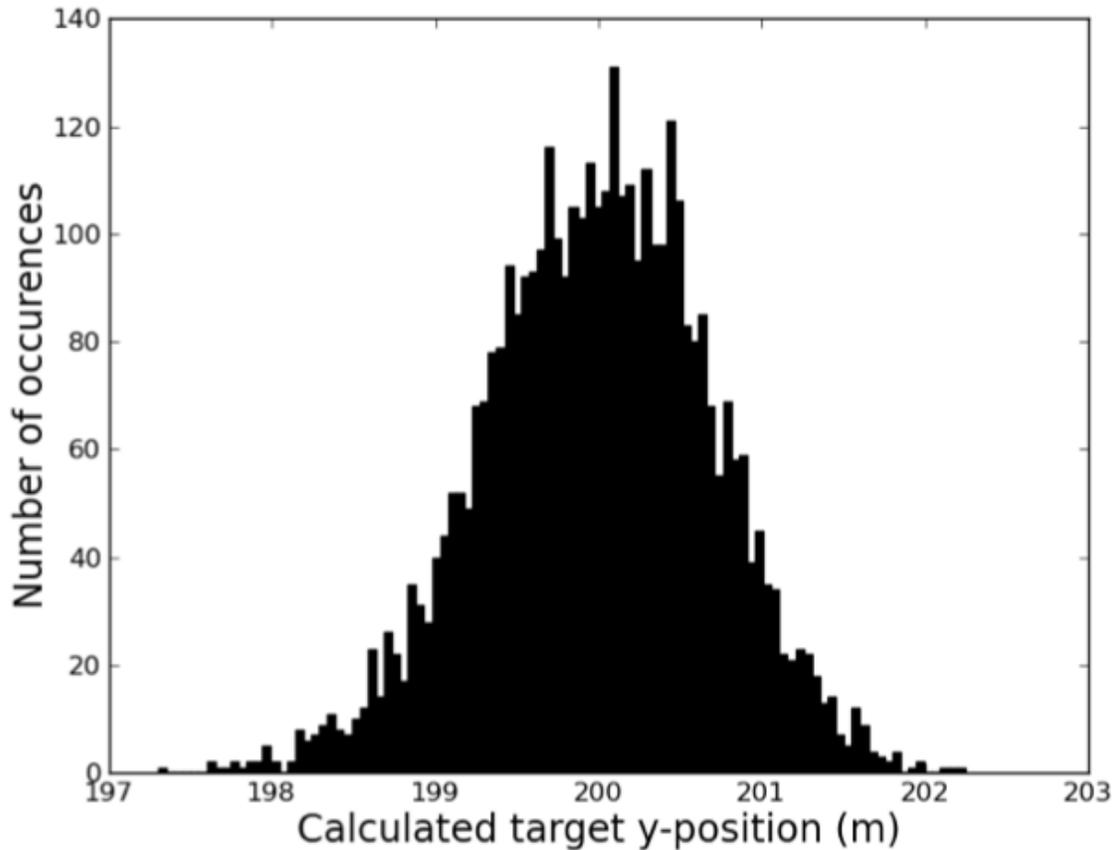
Figure 20:   Histogram of calculated position of reflector at 200m with $\sigma_w = 1$.

4.3   *EKF with Only Radar Available*

This section contains simulated navigation results using the direct-state models developed in Section 3.6, the radar model developed in Section 3.4, and the AV model and scenario described in Sections 3.2 and 3.3 respectively.

The initial position, velocity, and acceleration of the AV is assumed to be known with complete certainty, with the matrix values for these 9 quantities set to zero initially. The reflector locations are initialized using Doppler and range to get a crude estimate of each locations. The method described in Section 3.9 is used to initialize newly discovered stationary persistent radar reflectors. The **P** matrix values for the target locations are set to 2500 m^2, so that the EKF does not trust the crude measurement and converge to the wrong location during initial linearization. However

73
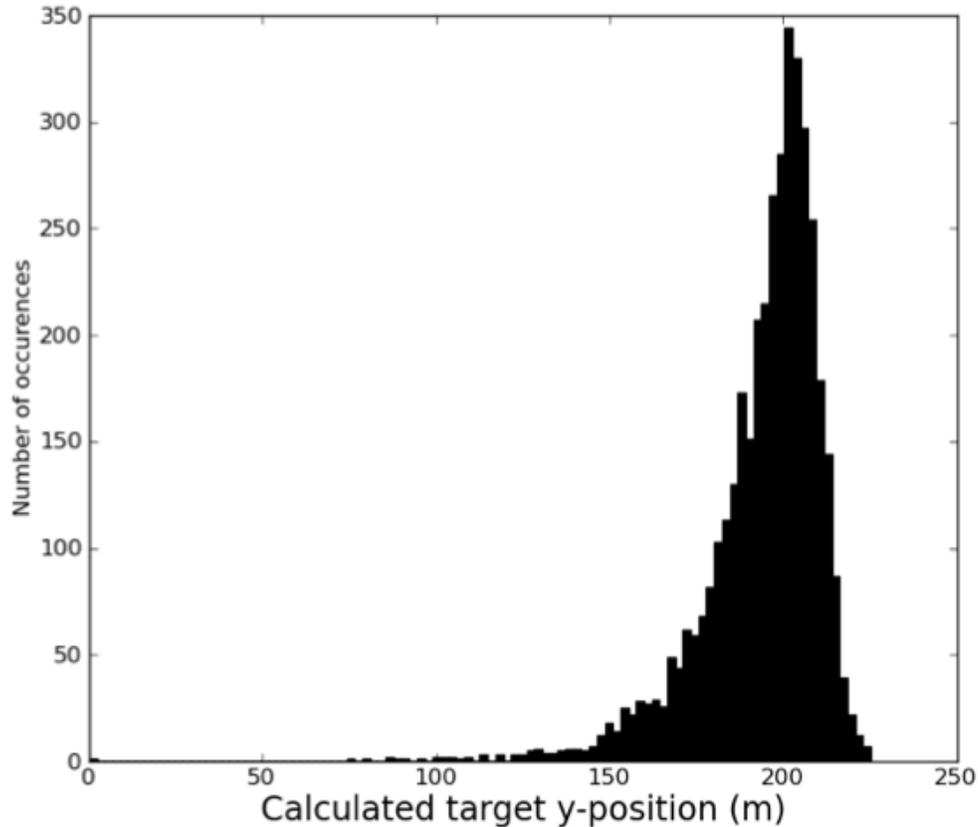
Figure 21:   Histogram of calculated position of reflector at 200m with $\sigma_w = 20$. Distribution displays a negative skew.

the initial crude estimate is still critical, as it affects the linearization point in the Jacobian computation and can cause divergence in the EKF solution if it is a bad estimate.

The parameters listed in Table 3 were used for all simulations in this section. These were chosen to mimic those of the experimental radar prototype system. The geometry considered is a set of reflectors arrayed in a line parallel to the initial direction of travel of the AV. Table 4 details the different configurations of the lines and the associated name of each configuration.

4.3.1   Single Trial Simulations.   In order to illustrate the performance of the filter, a single trial was performed in the LVNDG configuration.
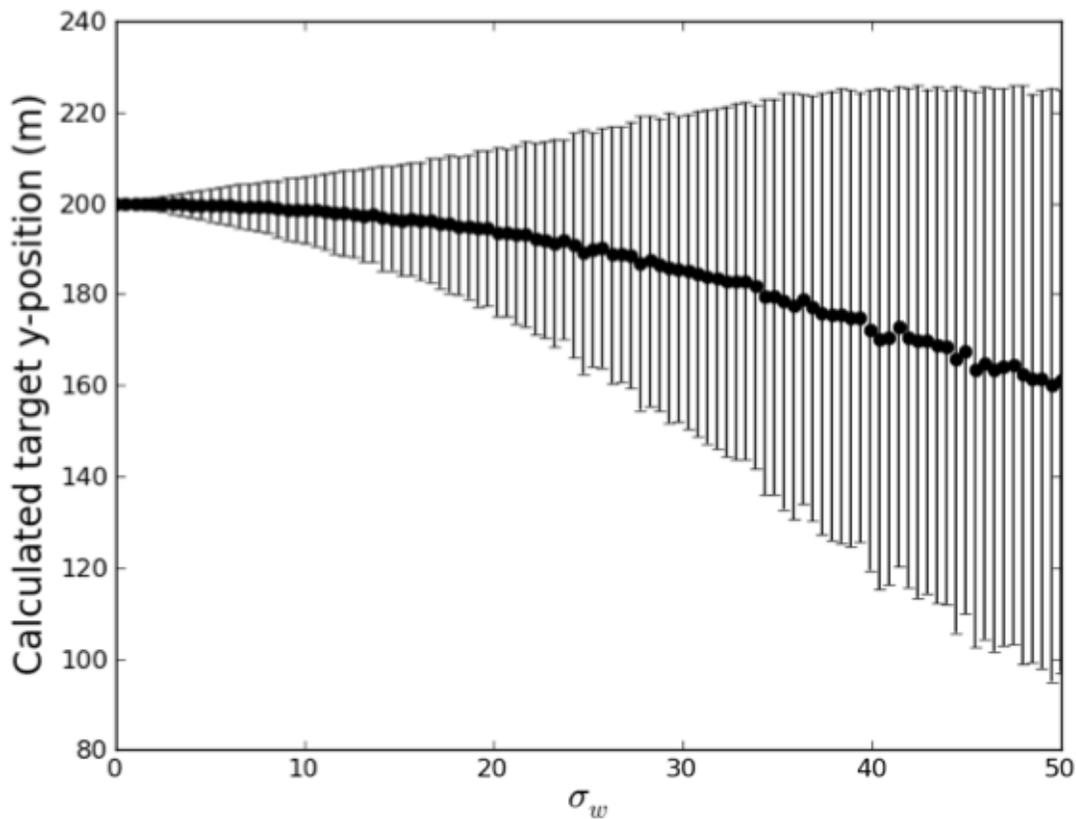
Figure 22: Calculated position of reflector at 200m with varying σ_w, averaged over 500,000 trials. We see a negative bias in the calculated solution as the noise level increases.

Table 4: Line configurations considered. Δx is the distance between each scatterer. Δy is the distance of the line of scatterers from the AV.

Configuration Name	Δx	Δy	N_r
LF (line far)	Δx=800 m	Δy=1500 m	50
LFD (line far dense)	Δx=300 m	Δy=1500 m	50
LND (line near dense)	Δx=300 m	Δy=800 m	50
LVND (line very near dense)	Δx=300 m	Δy=300 m	50
LVNDG (LVND w/ gap)	Δx=300 m	Δy=300 m	50

Fig. 33 shows the X and Y axis position estimate and true value. Fig. 34 shows the AV position error between the estimate and truth. We observe less than 5m of position error in both directions although the AV has traveled 3.5km in the X
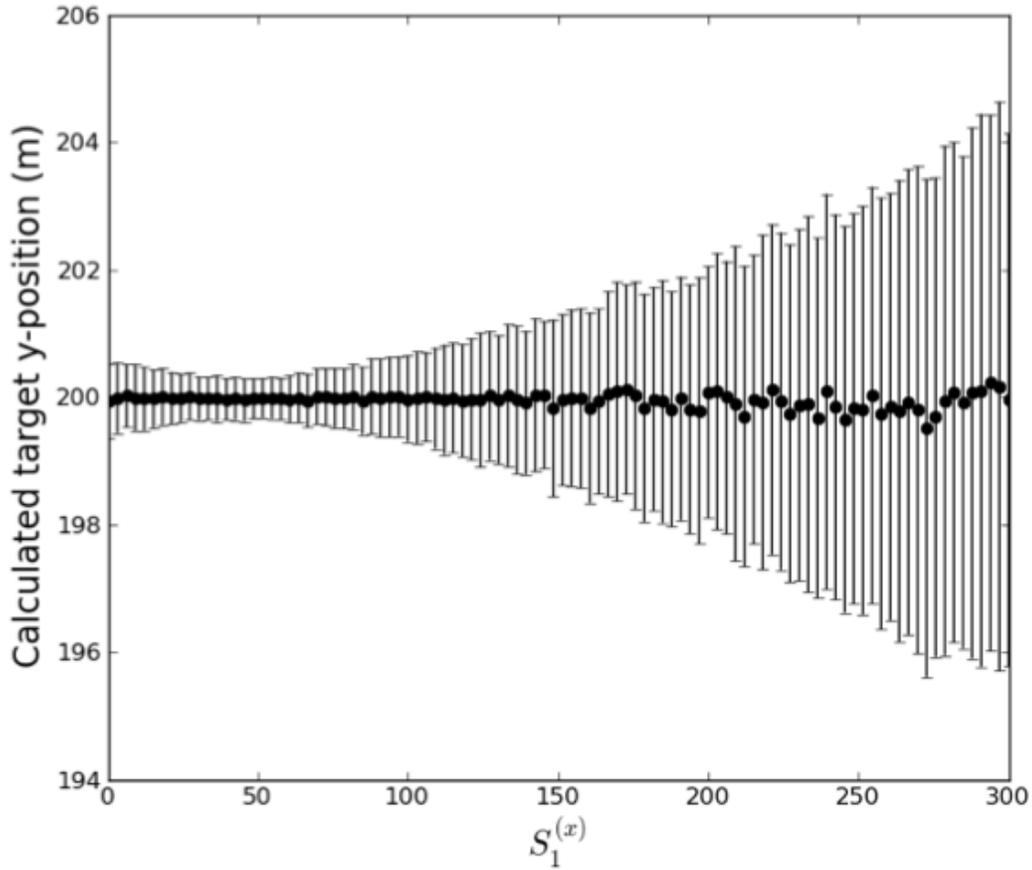
Figure 23:   Calculated reflector y-coordinate as a function of its x-coordinate. We see the mean of the estimate is not a function of the reflector's x-coordinate; the variance is however, showing very fast growth.

direction and 1.2km in the Y direction due to the FOGM acceleration in this trial. The initial velocity in the Y direction is 0m, which means that the 800m travelled in the positive Y direction is completely random but is still tracked by the Kalman filter with high accuracy.

Fig. 35 and 36 show the velocity and acceleration states for X and Y respectively. The error between the truth and the Kalman output is shown in each. We can observe that the error generally falls within the bounds of the P-value. This validates the P value, as it is the standard deviation of the observed states. In this configuration, the velocity error is within 10 m/s, indicating a reasonable error which is expected for our measurement noise model and process noise. We see that the filter is able to

76
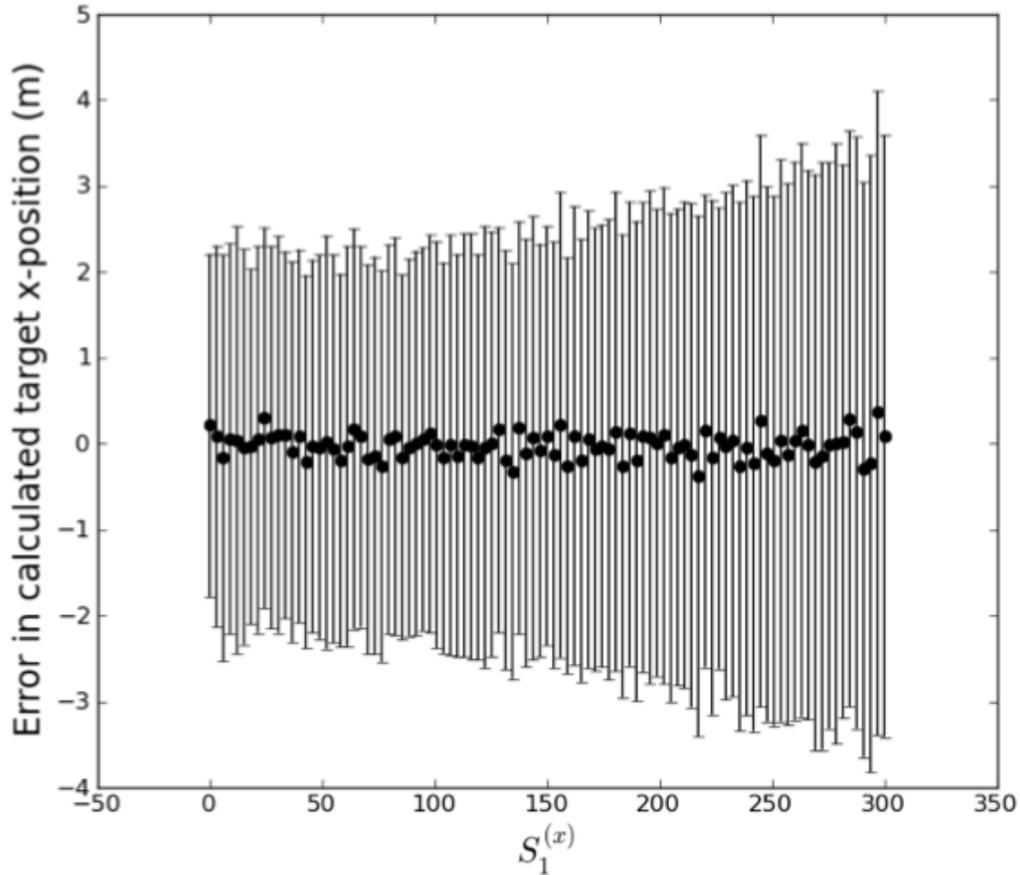
Figure 24:   Calculated reflector X coordinate error as a function of its true X co-ordinate. The Y axis is the difference between the calculated position and the true position.

track the FOGM random acceleration of the plane well with this scatterer availability. The Kalman output tracks the true acceleration with a slight delay, most likely due to the double integration between position (which the ranges contain information about) and acceleration (which is being estimated). We see that in LF configuration the P value has a standard deviation of 10 m/s^2 in both directions. This indicates that the short term error in the state model is low for both axes.

4.3.2 Ensemble Statistics.   Each of the configurations listed in Table 4 was simulated with 1,000,000 trials. Fig. 37 shows the standard deviations calculated for each of these configurations for both the X and Y axes. We immediately notice that

77
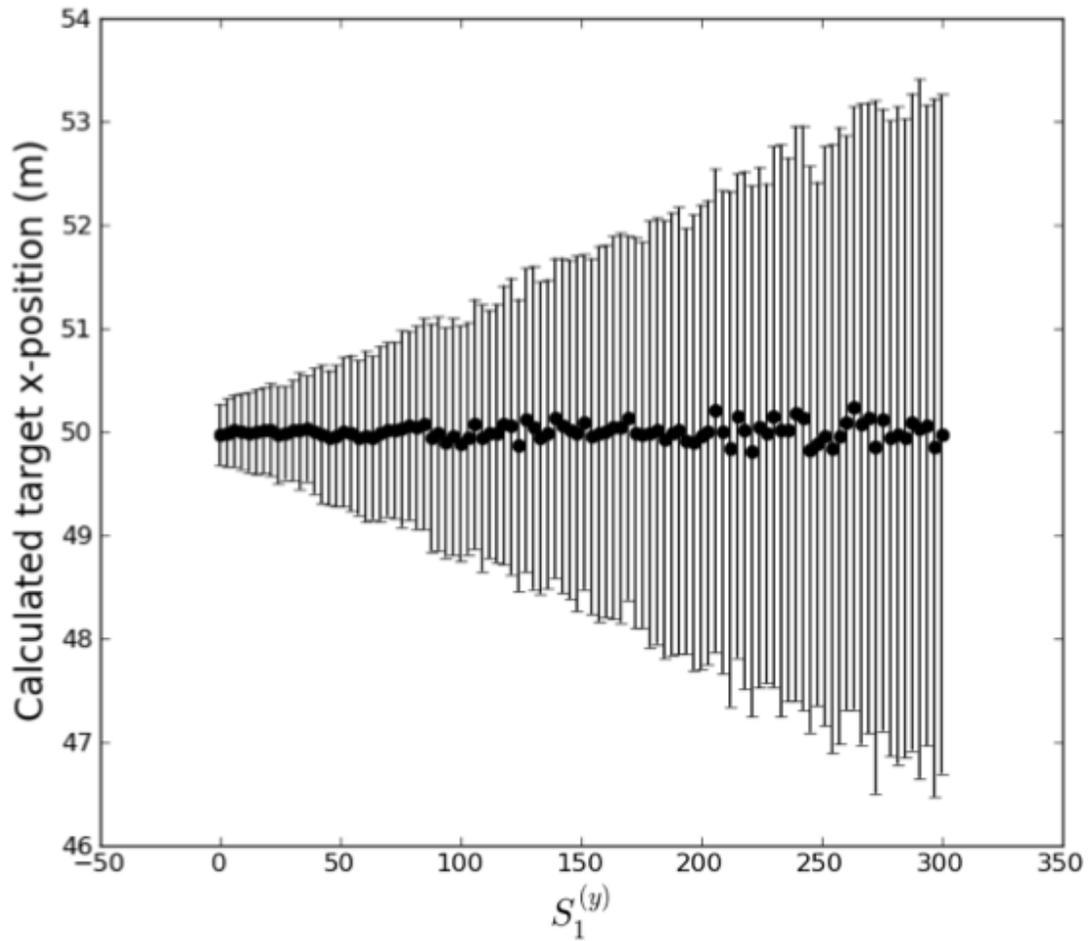
Figure 25:   Calculated reflector X coordinate as a function of its Y coordinate.

the variation between the configurations in the X axis is very small, with less than 0.5m separation. Since each configuration is a set of reflectors in a long line parallel to the X axis, there is always good observability in the X direction. Thus for the rest of this section we will focus our analysis on the Y axis. Note that by inverting the line to be parallel to the Y axis, the reverse results will be computed, with good observability along the Y axis and little observability along the X axis.

The LF configuration maintains a standard deviation of less than 10 meters for the entire 21 second duration. This configuration represents a good geometry, with visible reflectors far away in both the cross and down ranges. The LFD configuration is

78
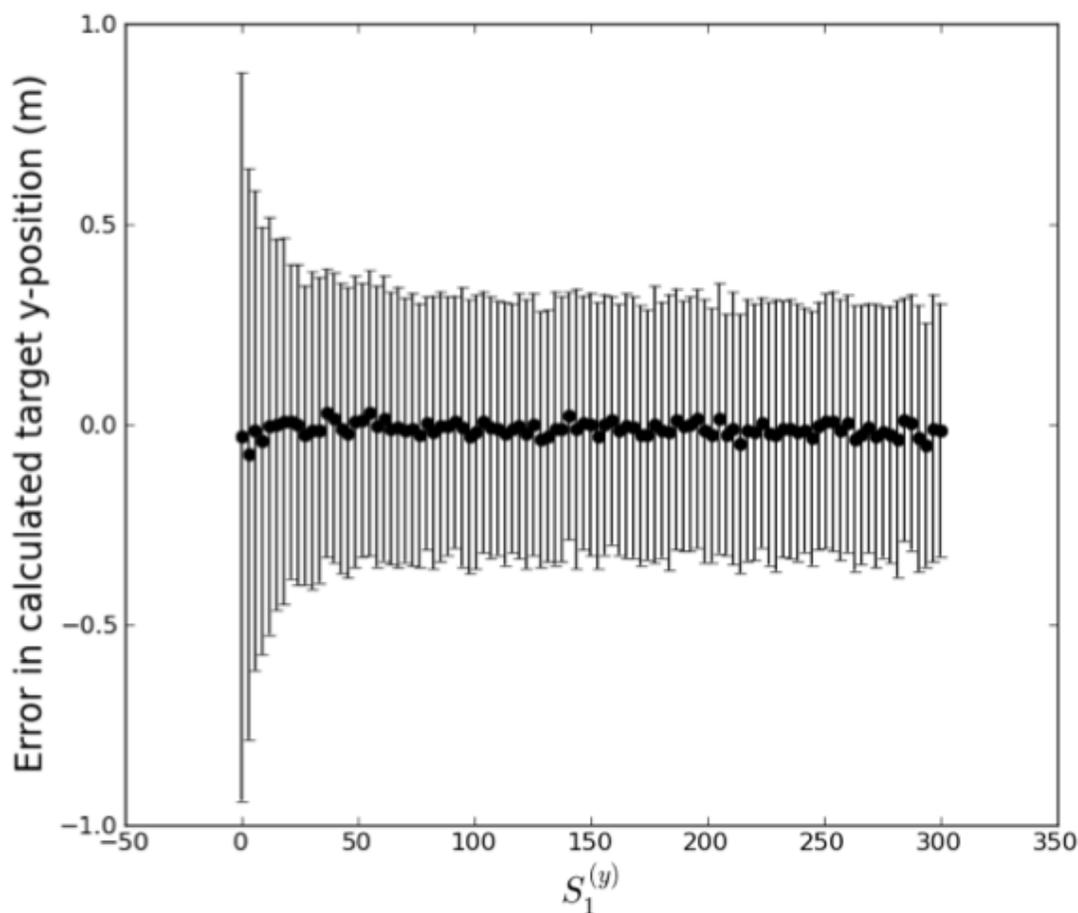
Figure 26: Calculated reflector Y coordinate error as a function of its true Y co-ordinate. The Y axis is the difference between the calculated position and the true position.

the same as LF except with reflector density increased. We see the expected decrease in Y axis error, from 7m to 4m.

The last three configurations all diverge eventually in their estimate of the po-sition. All three of these configurations are either Near or Very Near, denoting that the line of reflectors is close to the beginning track of the plane (the X axis). This illustrates that if the only observable reflectors are too near to the AV, there is not enough information to estimate the AVs Y coordinate well, and bad estimates are made. These bad estimates are then used to calculate the Jacobian which in turn causes divergence in the position solution. Since this linearization error is invisible to

79
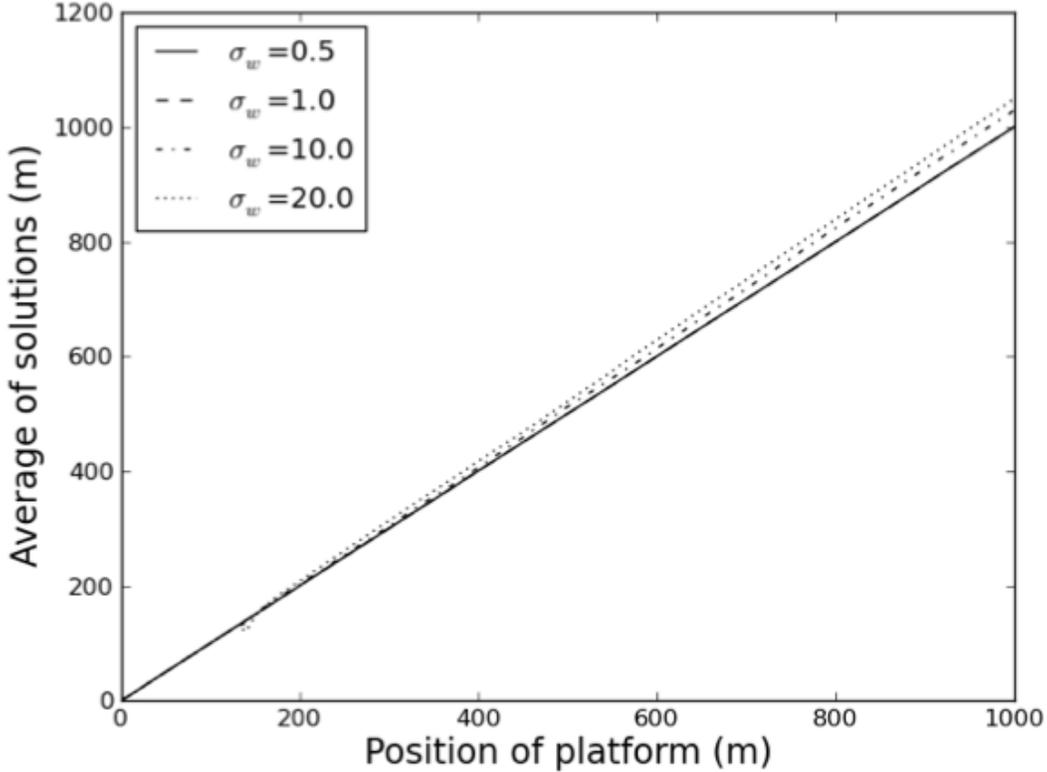
Figure 27:   Tracking AV position as it travels 1km with reflectors every 10m, at a distance of 200m from the AV. Mean is shown over 1 million trials.

the EKF, the P value remains small. This results in the filter trusting bad (high error) information more than it should, and will cause a diverging position estimate that may be worse than simply not including the radar measurements themselves. Usage of the processing method described in this dissertation is only feasible if such bad geometries are detected and the filter recalibrated to take into account these anomalies. We note that the two Very Near configurations diverge quicker than the Near configuration, implying this relationship is a function of the distance of the reflectors to the AV track. The LVNDG configuration contains a gap in the line of reflectors which is encountered by the SAR processor at around 5 seconds. We note the Kalman filter takes into account the loss of information, as the P value for the Y axis for LVNDG rises to 35m. After the gap has passed, the EKF detects the good geometry returning, and the P value lowers again. However, the linearization error during the gap period

80
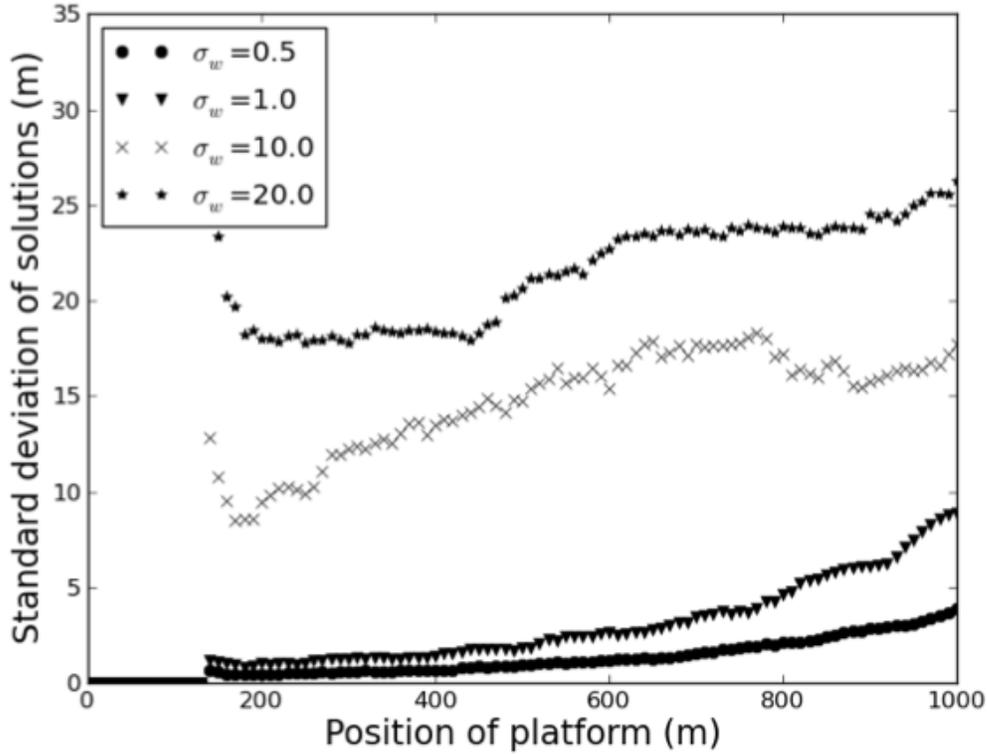
Figure 28:   Tracking AV position as it travels 1000m with reflectors every 10m, at a distance of 200m from the AV. Standard deviation is shown over 1 million trials.

causes divergence in the position solution, and the error diverges quicker than the LVND configuration.

Fig. 38 shows the mean values for each configuration. We see that they are all zero mean error as expected, except for the two Very Near configurations. The means of LVNDG and LVND are negative, which is an artifact of linearizing about an erroneous position.

4.4   EKF with Radar and INS Available

This section contains simulated navigation results using the error-state models developed in Section 3.7, the INS and radar models developed in Sections 3.4 and 3.5 respectively, and the AV model and scenario described in Sections 3.2 and 3.3 respectively. The simulations in this section use the INS parameters from Table 2, the radar parameters from Table 3, and the navigation parameters from Table 5. The

81
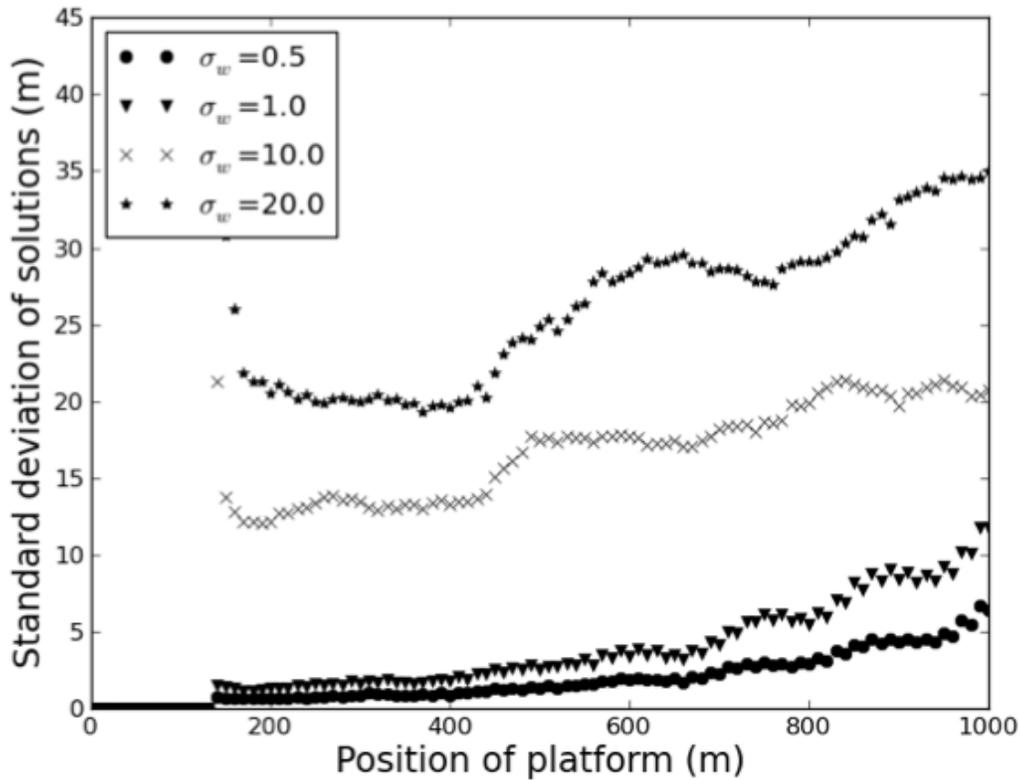
Figure 29: Tracking AV position as it travels 1km with reflectors every 20m, at a distance of 200m from the AV. Standard deviation shown over 1 million trials.

Mie reflectors are arranged in two staggered lines parallel to the direction of flight. The lines are located 300m and 600m from the initial flight path and have 300m spacing inbetween each reflector.

4.4.1 Single Trial and Multiple Trace Results. In order to illustrate the performance of the filter, a single trial was performed using a tactical grade INS. Fig. 39 shows the error in the error-state estimates with and without measurements turned on for a tactical-grade INS. The measurements are turned off by forcing the Kalman gain to zero, showing the actual error states without UWB-OFDM SAR aiding. We see that for a tactical-grade INS, the UWB-OFDM aiding greatly improves the position estimate. The errors for the case with SAR aiding are less than 3m with a tactical grade inertial, whereas they diverge to greater than 1000m for the no-
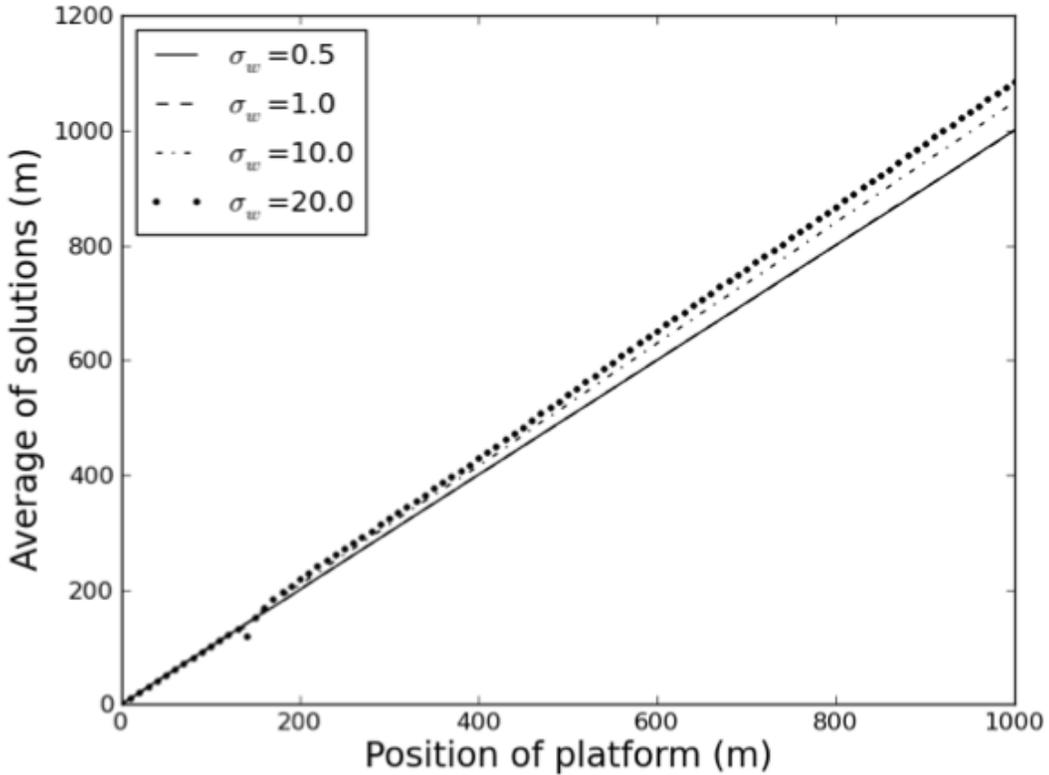
Figure 30:   Tracking AV position as it travels 1km with reflectors every 20m, at a distance of 200m from the AV. Mean is shown over 1 million trials.

measurements trial. Fig. 40 repeats the simulation in Fig. 39 30 times and show the results for each realization. We see that the single-run results shown in Figure 39 are not outliers, and that the trials are on the same order as the filter-computed computed from the P matrix.

4.4.2 Ensemble Statistics.   The ensemble statistics in this section are calculated by running 1 million trials. Fig. 42 shows the ensemble statistics for a tactical grade INS with and without radar measurements. The results here mirror what we saw in the single trial simulations, with the radar measurements significantly decreasing the error. For the x-axis, we see 1200m of error for the no-measurements case and around 1.2m of error with measurements after 600 seconds. For the y-axis and
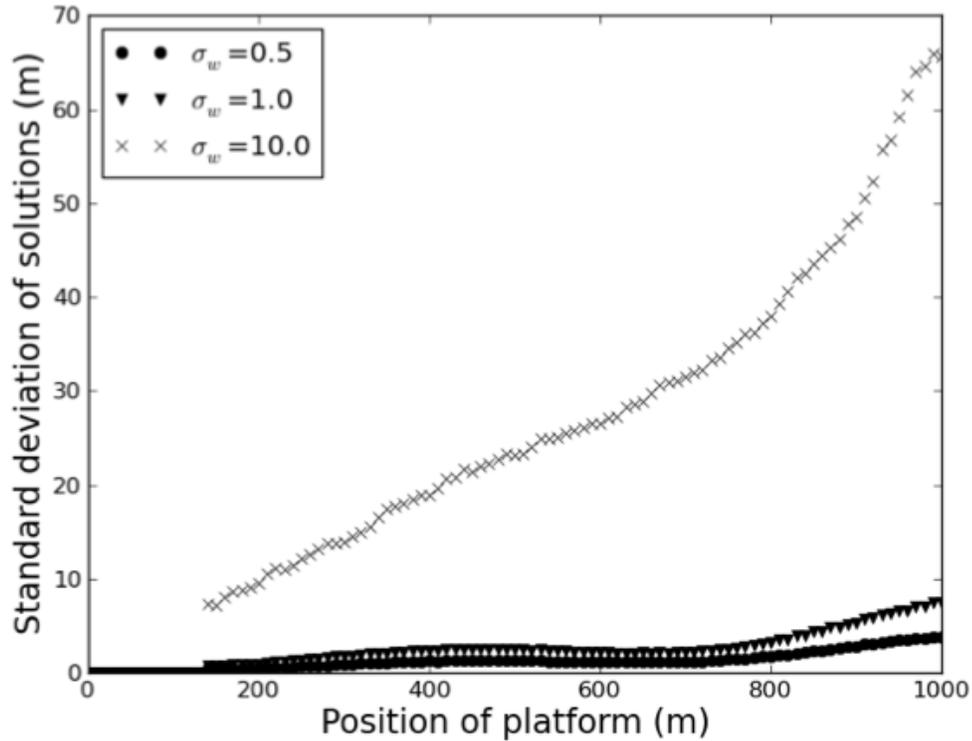
Figure 31:   Tracking AV position as it travels 1km with reflectors every 20m, at a distance of 700m. Standard deviation is shown over 1 million trials.

z-axis, we see that the decrease in error is not as extreme, with 8m of error and 100m of error for the y and z axes respectively when using SAR aiding.

In order to gauge the effects of INS grade on the results, commercial and navigation grade INS were simulated. Fig. 41 shows the ensemble statistics for a commercial grade INS with and without radar measurements. We see decreased performance of the stand-alone INS when using a cheaper commercial grade package compared to tactical grade. However, with radar measurements the performance is only slightly worse than the tactical grade with measurements. The x-axis error after 600 seconds is 1.2m and 5m and the y-axis error is 8m and 20m for aided tactical and commercial grades respectively. Thus the radar aiding is extremely beneficial for lower quality commercial and tactical INS devices.
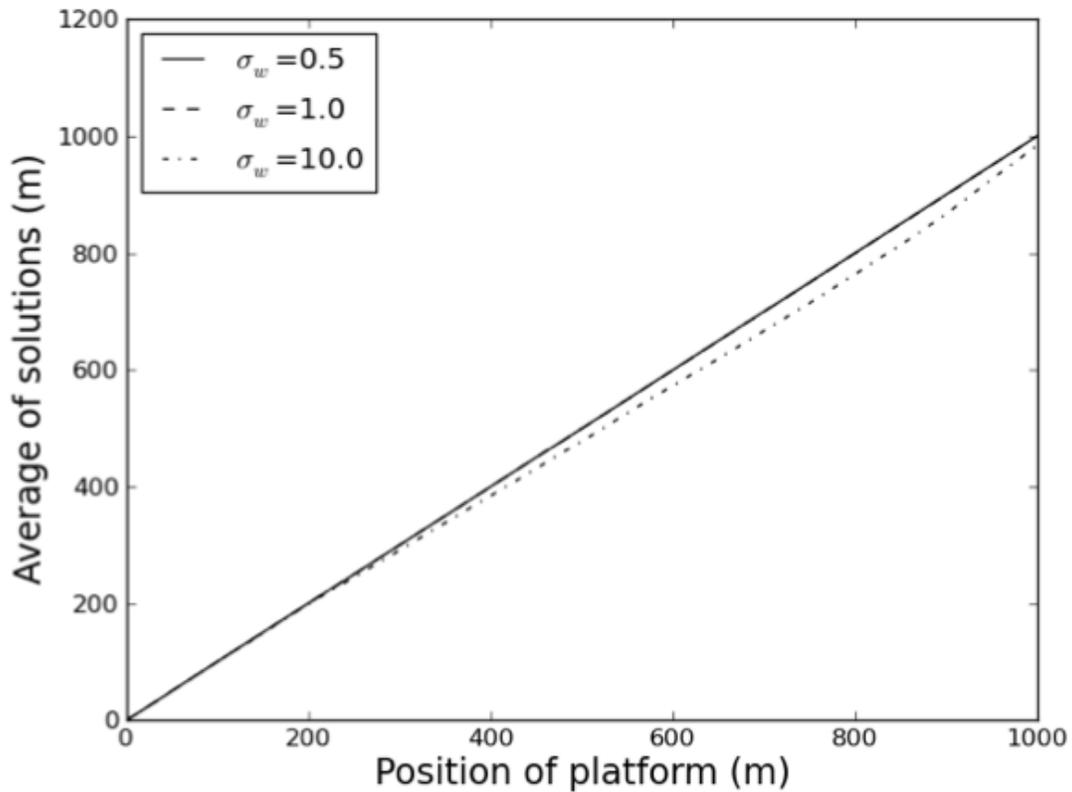
Figure 32:   Tracking AV position as it travels 1km with reflectors every 20m, at a distance of 700m. Mean is shown over 1 million trials.

Fig. 43 shows the ensemble statistics for a navigation grade INS with and without radar measurements. We see the baseline error of a stand-alone navigation grade INS after 600 seconds is around 1-20m depending on the axis. Since the baseline error is so small, we expect a navigation grade INS to not benefit greatly from SAR aiding, which is seen here.

Tables 6 and 7 summarize the results from Figures 41-43. Note that we omit the z-axis, as the results in the z-axis may be inaccurate due to our use of an idealized gravity model. We see that the reduction in error when using SAR aiding is more significant for low-grade INS devices. However, even with a navigation-grade INS, the error reduction is two orders of magnitude.
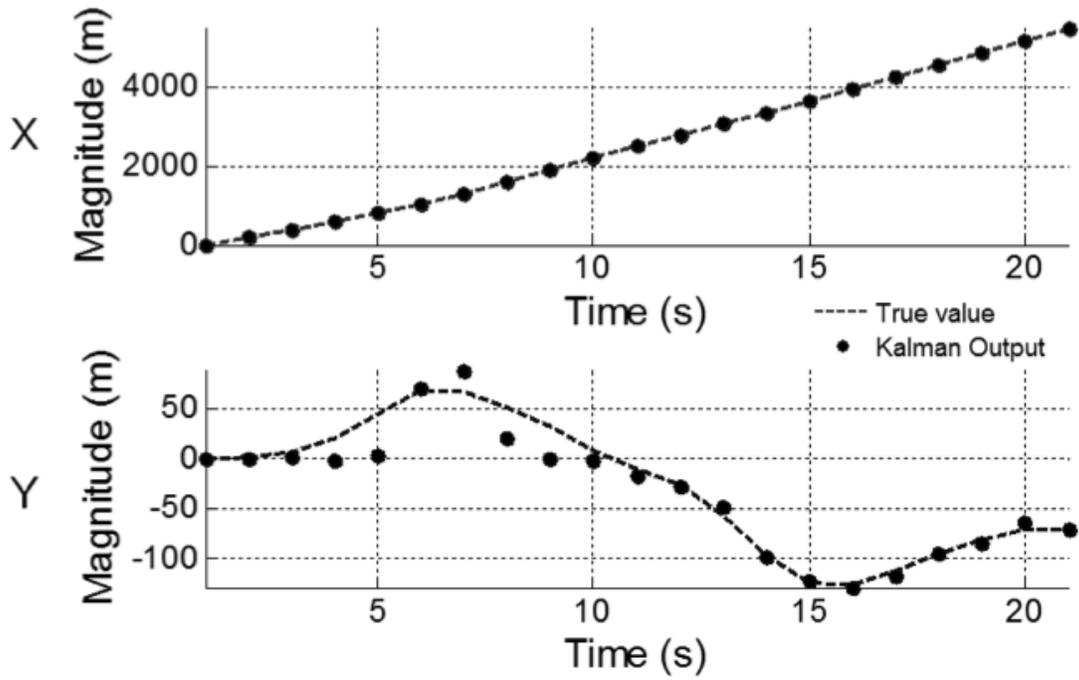
Figure 33:   Absolute position estimate and truth. (top) X-axis (bottom) Y-axis.
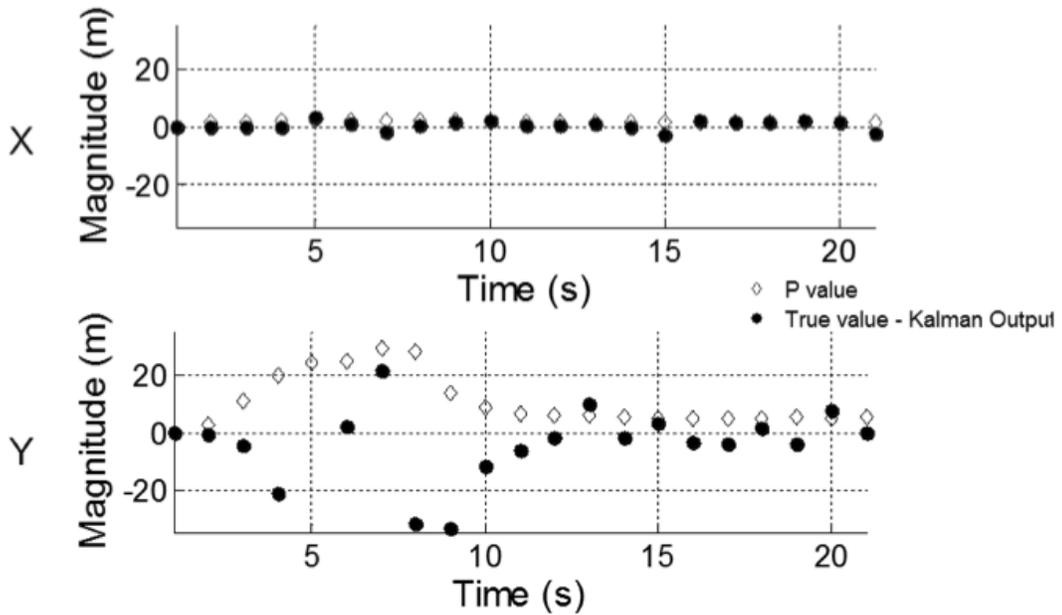
Figure 34:   Position state error and expected P value. (top) X-axis (bottom) Y-axis.

86
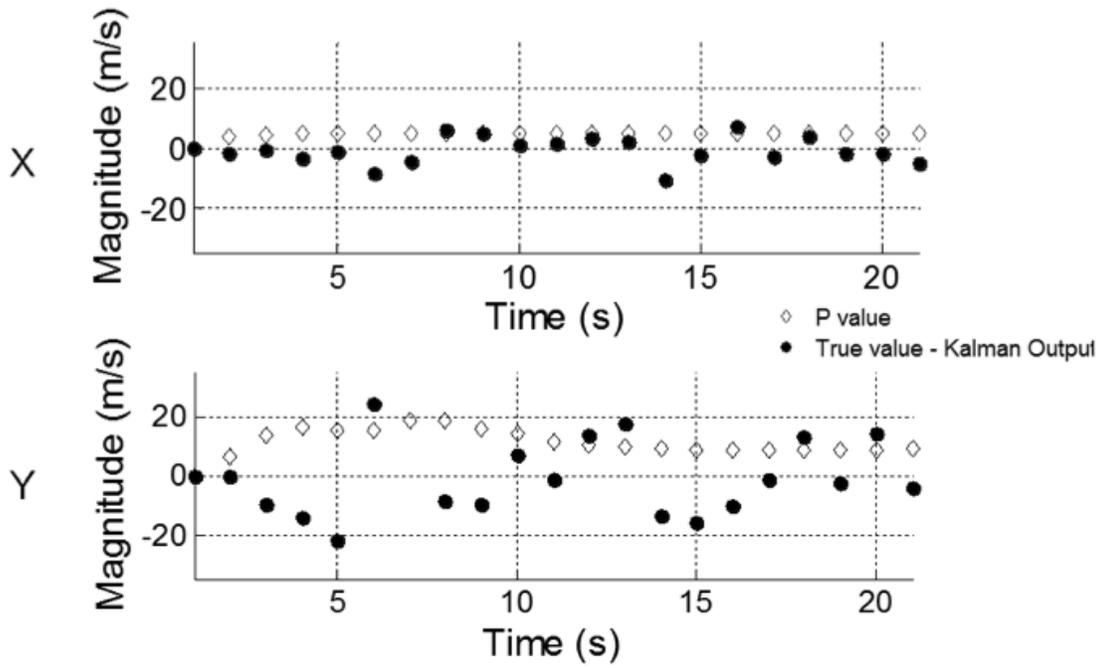
Figure 35:   Velocity state error and expected P value. (top) X-axis (bottom) Y-axis.
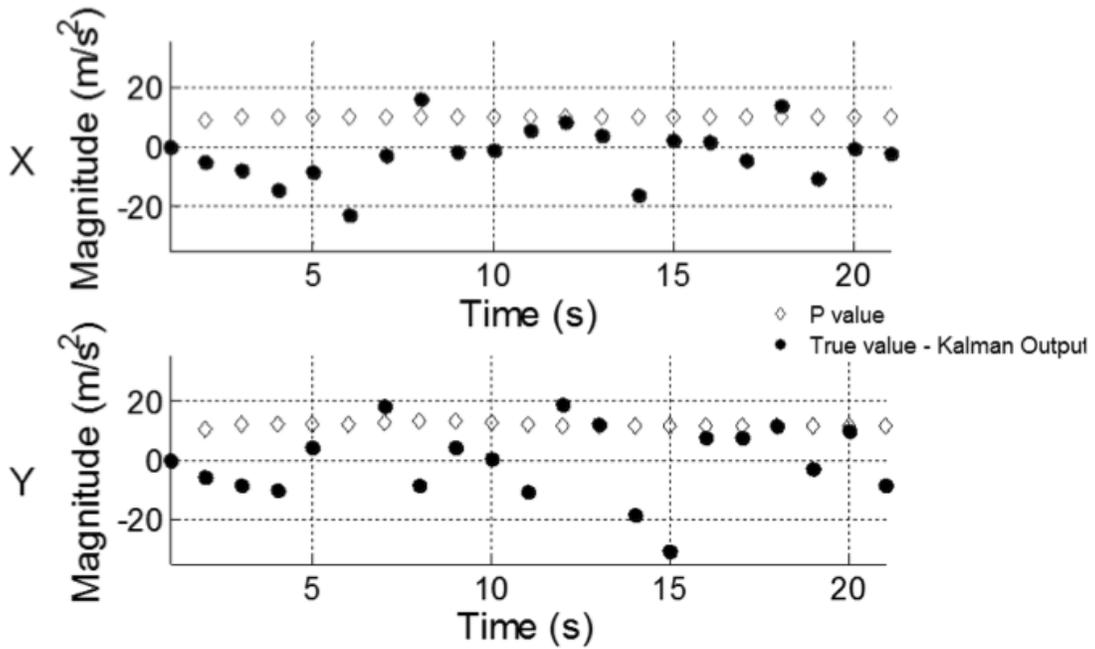
Figure 36:   Acceleration state error over time and expected P value. (top) X-axis (bottom) Y-axis.
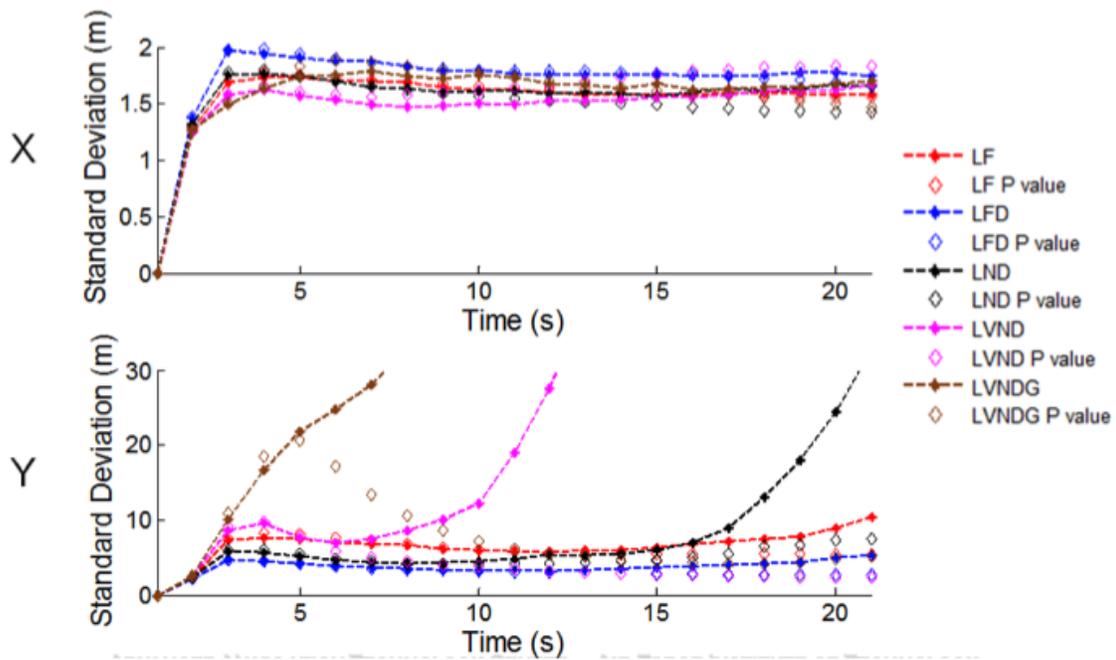
Figure 37: Standard deviation over 1 million trials for each configuration.
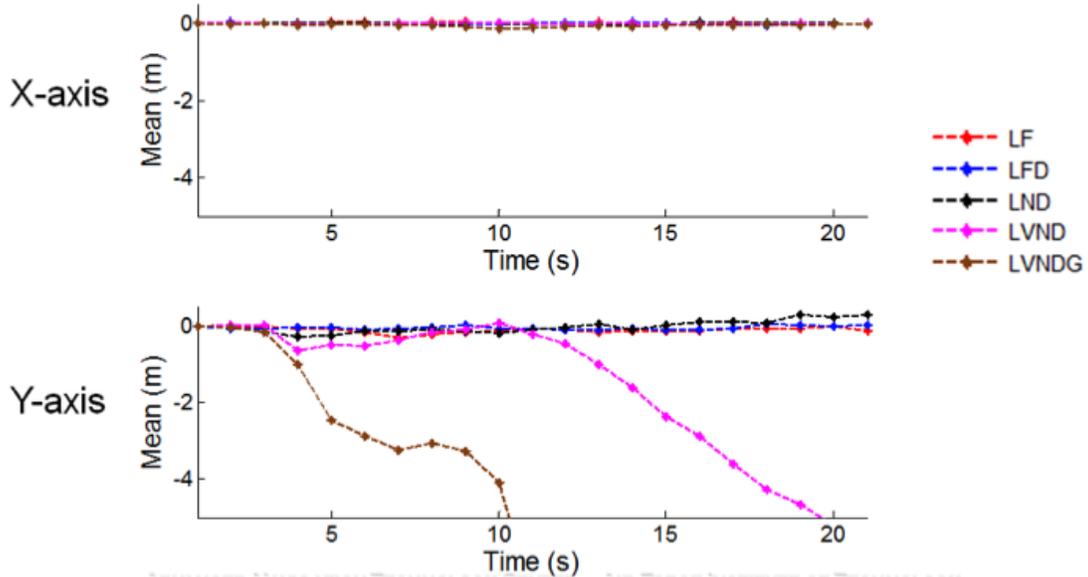
Figure 38: Mean over 1 million trials for each configuration.

Table 5: Simulation Parameters for Navigation Filter

EKF Propagation Time	1s
Initial Position	(0,0,0) m
Initial Velocity	(200,0,0) m/s
Initial Acceleration	(0,0,0) m/s^2

Table 6: Summary of X-axis error for each configuration after 600 seconds of flight.

	Measurements	No Measurements
Commercial Grade	4m	600km
Tactical Grade	1.2m	1700km
Navigation Grade	0.5m	15m

Table 7: Summary of Y-axis error for each configuration after 600 seconds of flight.

	Measurements	No Measurements
Commercial Grade	15m	600km
Tactical Grade	10m	1700km
Navigation Grade	4m	15m
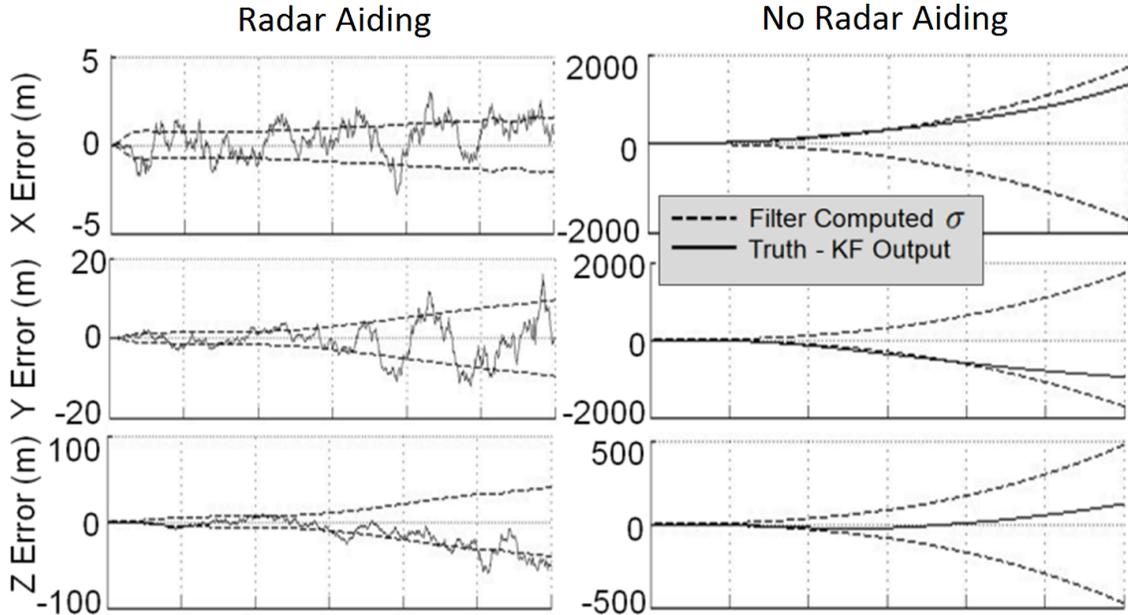
Figure 39: Difference between true inertial error and Kalman estimate of inertial error for tactical-grade INS with and without radar measurements. A single trial is plotted for each axis.
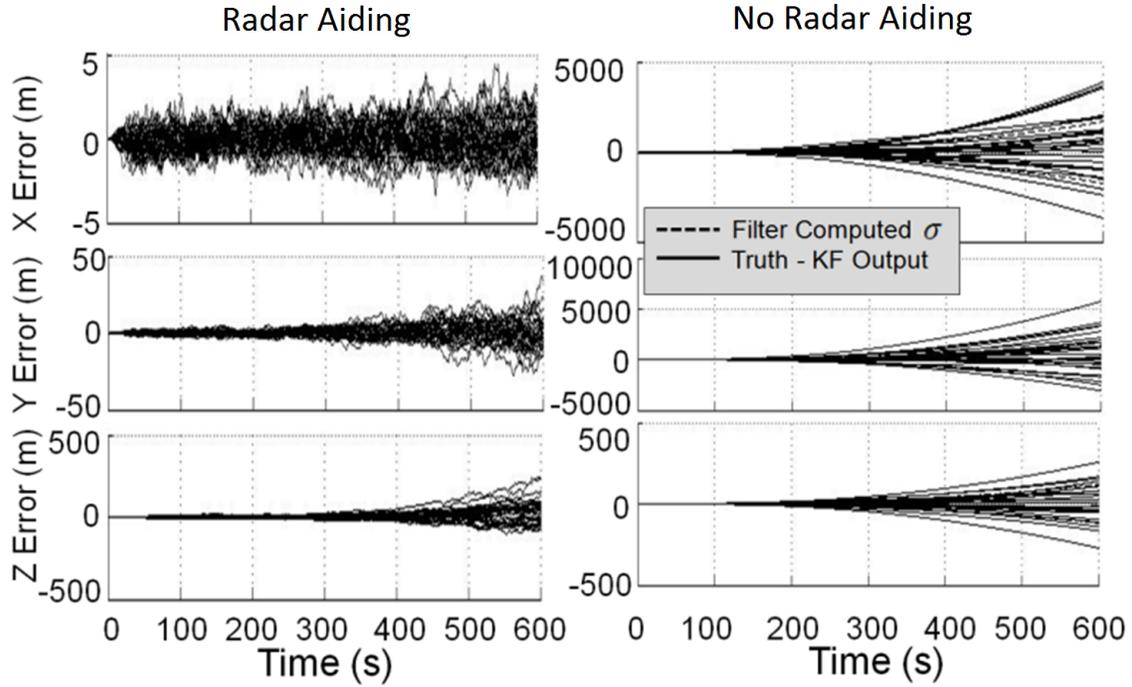
Figure 40: Difference between true inertial error and Kalman estimate of inertial error for tactical-grade INS with and without radar measurements. 30 trials are plotted for each axis.
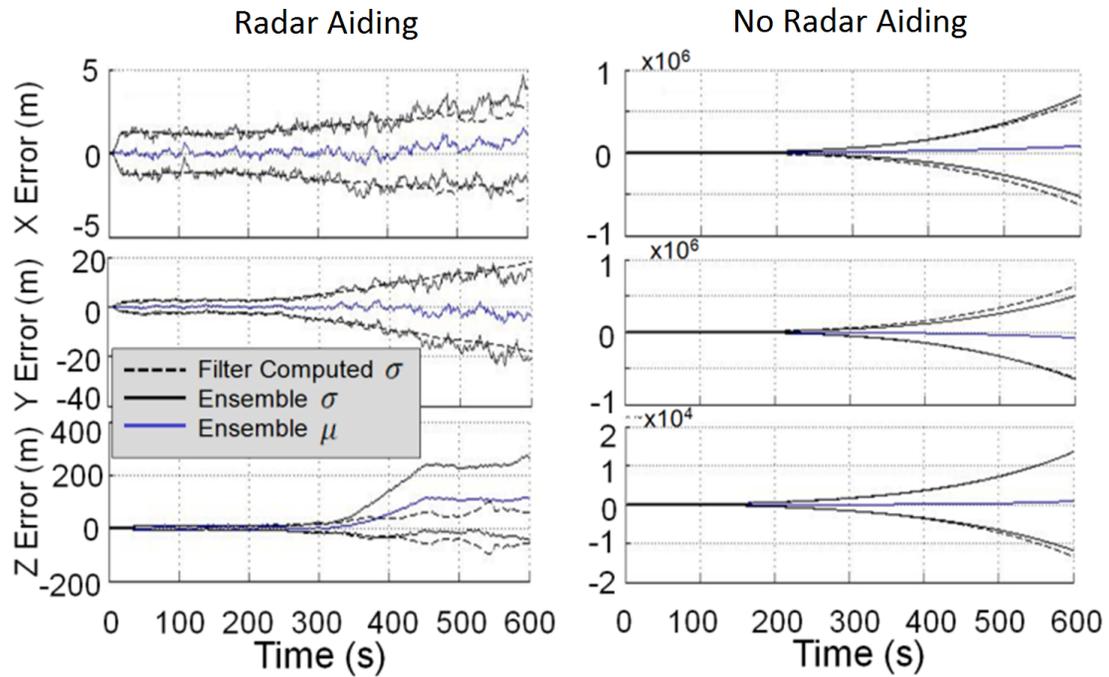
Figure 41: Ensemble mean and standard deviation versus filter computed standard deviation for commercial grade INS with and without radar measurements.
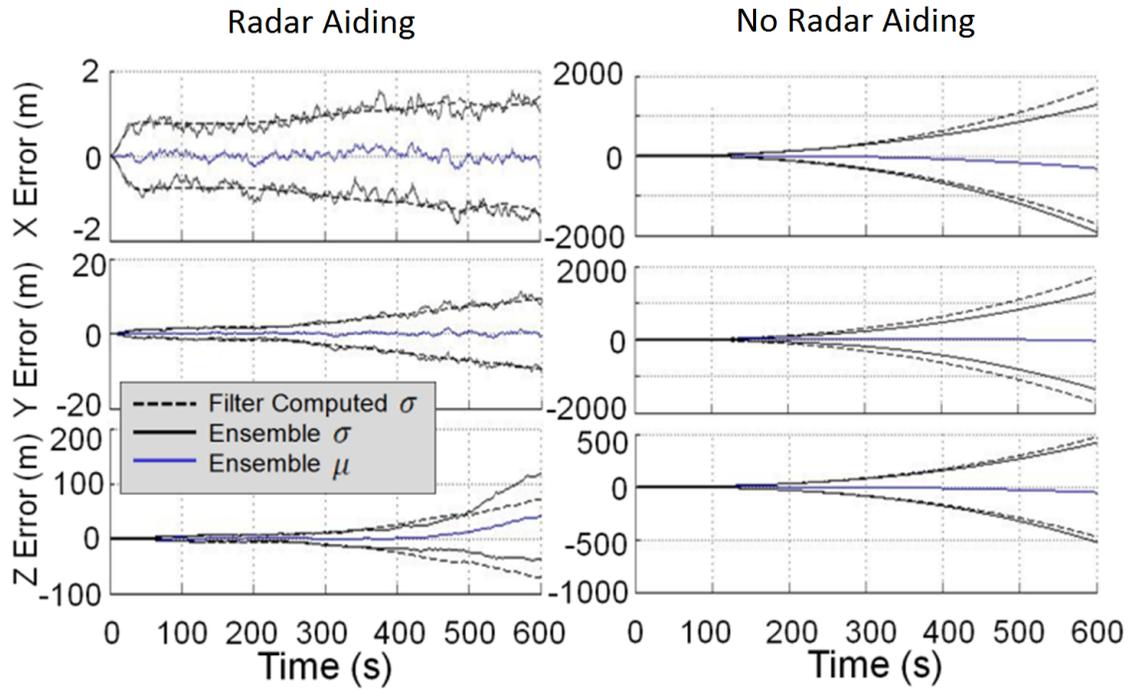
Figure 42: Ensemble mean and standard deviation versus filter computed standard deviation for tactical grade INS with and without radar measurements.
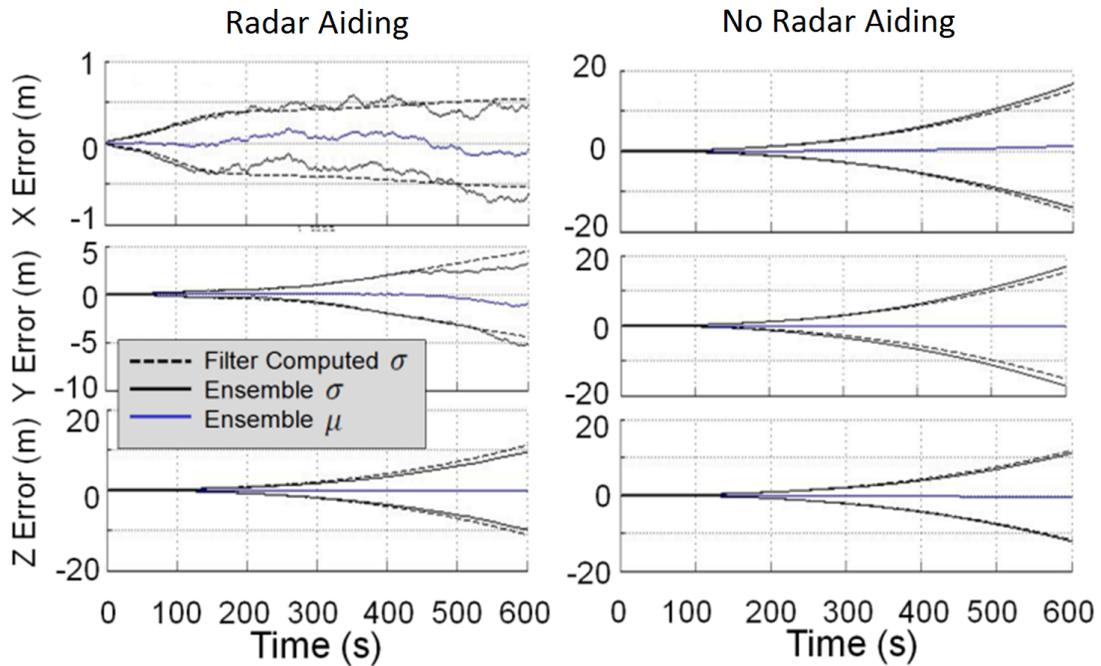
Figure 43: Ensemble mean and standard deviation versus filter computed standard deviation for navigation grade INS with and without radar measurements.

V. Experimental Results

In this chapter, we provide experimental performance results for the methods discussed in Chapter III. The experimental results were collected using an experimental UGV UWB-OFDM system prototype. We consider the navigation performance of the mobile system for different geometries and reflector availability.

5.1   Experimental System

This section details the software and hardware we developed in order to perform experimental validation of the navigation algorithms developed previously.

5.1.1   Radar System.   A hardware prototype was constructed in conjunction with other researchers at Miami University [42]. The system prototype is designed as a multi-purpose device capable of communications and radar imaging, as well as other utility roles such as acting as the role of an integrated navigation sensor considered in this study. The original system design is shown in Fig. 44 and the device itself is pictured in Fig. 45. Fig. 46 shows an example of a backprojected SAR image. Fig. 47 shows the bit error rate (BER) at a 5m distance when the device is used as communications device.

In order to facilitate usage of the system in a navigation context several improvements were made to the system. The original receiver board was replaced by a Tektronix TDS3201D board as it:

- Decreased the timing jitter in the trigger assembly to less than 200ps. This fine Tx/Rx synchronization allows us to perform true absolute ranging to targets, critical for real-time autonomous navigation. The previous board exhibited 50ns of timing jitter, which allows only for relative ranging between targets in the same coherent processing interval (CPI).

- Allowed for a 5 GS/s sampling rate. This enhancement allows for precise ranging of isolated targets.
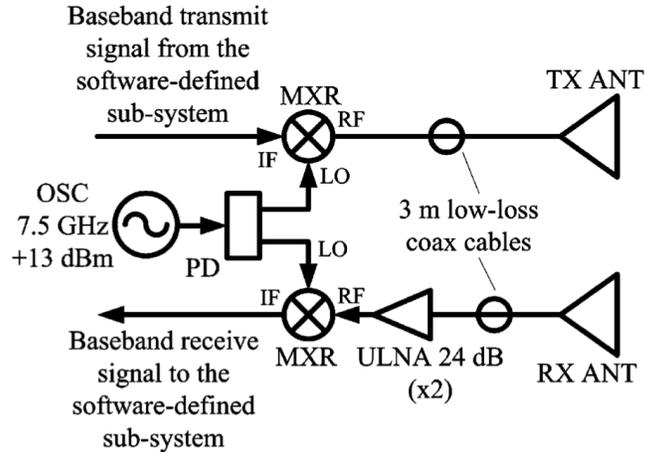
Figure 44:   Block diagram of experimental UWB-OFDM radar system.

- Enabled GPS-synchronized timing on data collections. This change facilitates integration with the INS, which time stamps its collected data via GPS time periods.

In addition, a 2-channel I/Q demodulator was used in order to allow video signal capture at baseband. This allows for continuous tracking of targets at any range, as opposed to signal-channel systems which will have dead-ranges where no return energy is detected (see Coherent Demodulation section in the background section for more details).

5.1.2   Navigation Platform.   The radar system from the previous section was designed to be used in stationary indoor environments. It was therefore modified to fit on a mobile cart along with an ATX computer, Honeywell HG1700 INS, SpanSE front-end system, and supporting electronics. The combined navigation system is pictured in Fig. 48.

5.1.3   Real-Time Processing Software.   A software suite was developed in order to implement the algorithms from Chapter III in real-time. All results in this chapter used this software. The real-time radar data processing (including detection and ranging) was performed using a massively parallel implementation on a GPU.
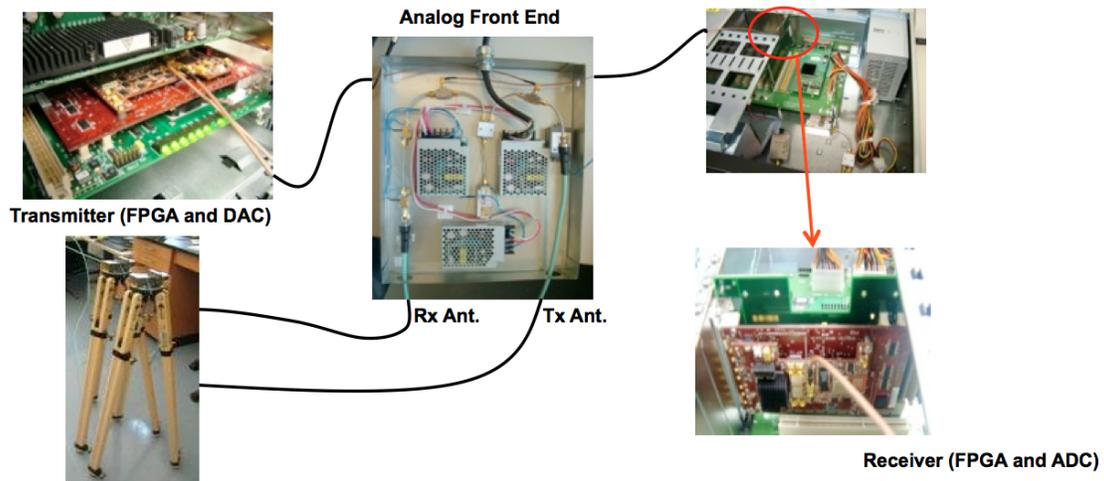
Figure 45:   Picture of components in experimental system.
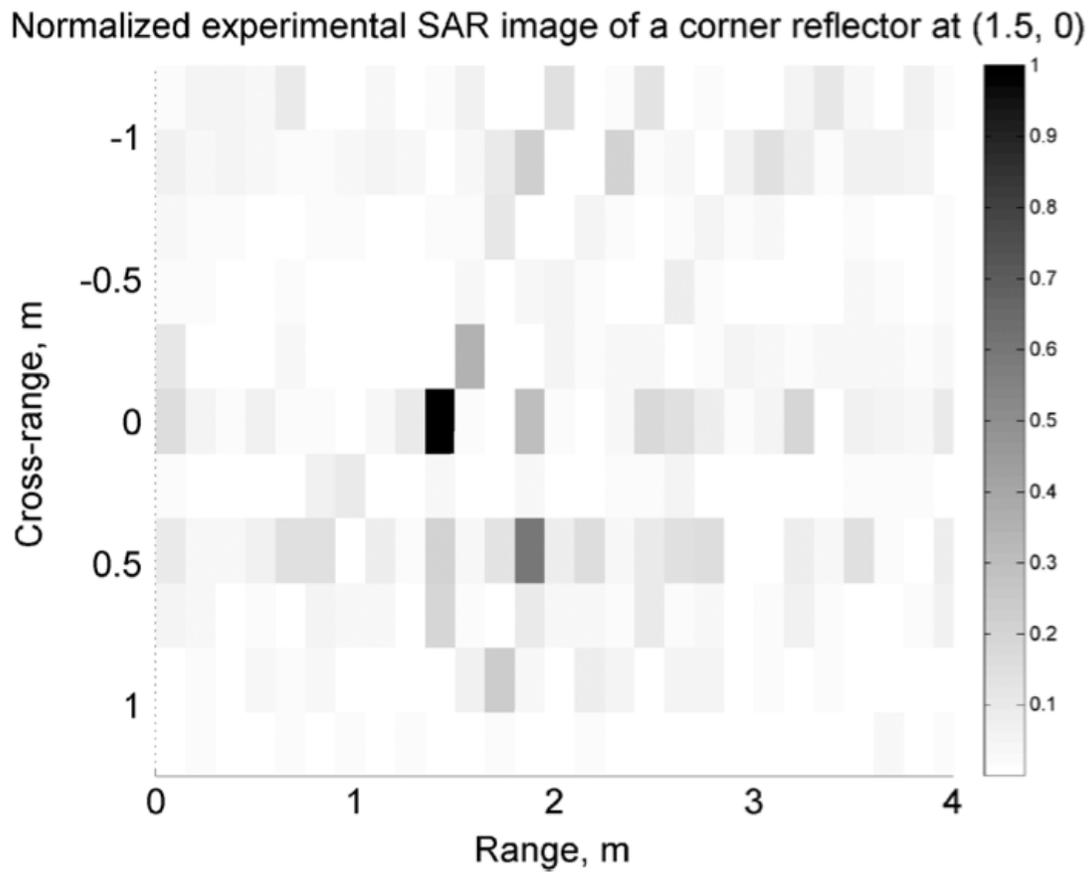
Figure 46:   SAR image captured with experimental system via backprojection.
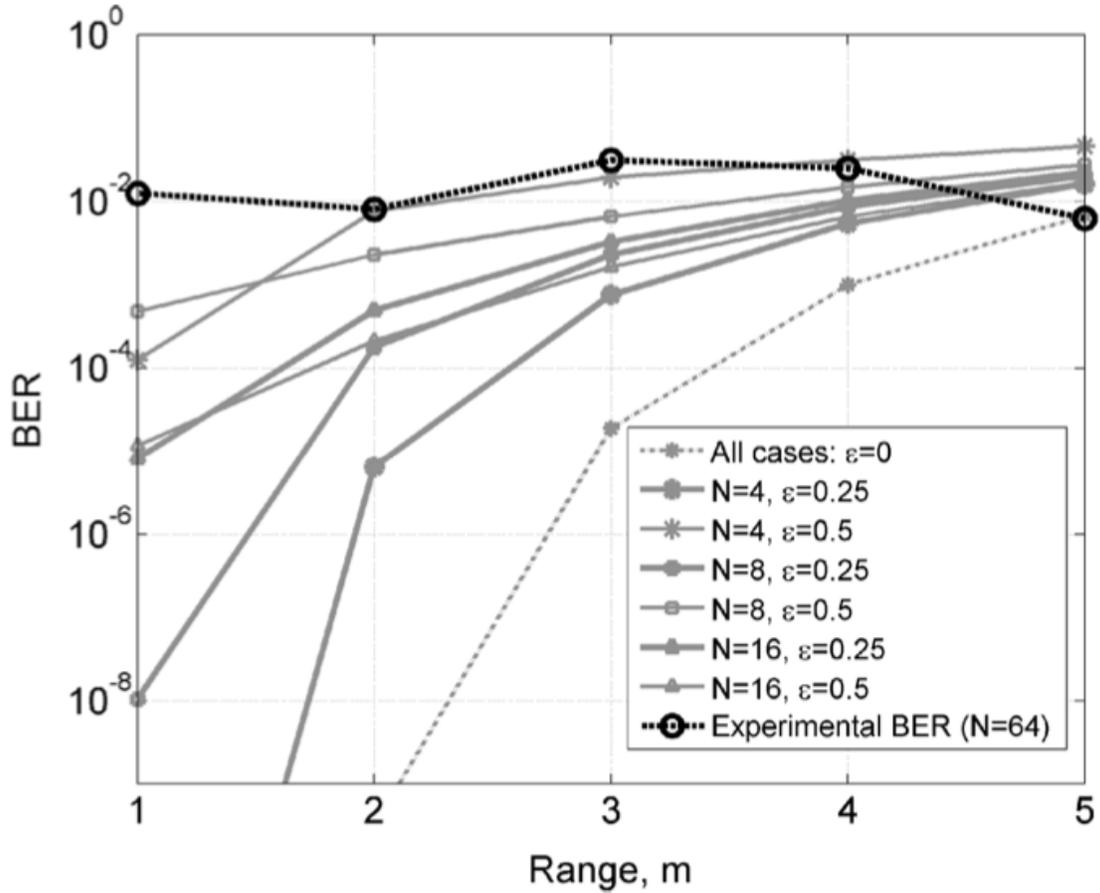
Figure 47: BER of experimental system transmitting at a theoretical data rate of 57Mb/s.

Data collection from each device was synchronized via a multi-threaded monitor program written in Qt. Table 8 illustrates the capabilities of the software suite.

Table 8: Real-time processing software characteristics.

Radar maximum CPI length	$80\mu s$
Radar data detect/track throughput	20×10^6 samples/sec
INS data rate	60 Hz
Synchronous sensor data collection jitter	< 10 ms
Real-time display and sensor failure detection	Radar, INS, camera sensors

Figure 48: Picture of mobile platform including an ATX computer, INS, UWB-OFDM radar system, and analog front-ends for both sensors.

5.2 Hallway Results

The experimental system was tested in an indoor hallway with metallic reflectors placed within the environment. Examples of the reflectors used are given in Fig. 49 and 50. Fig. 51 shows an overview of the data processing methods used to generate the results in this section. The INS data is processed in the fixed local-level frame to generate an unaided position estimate of the UGV. The radar data was processed

using the methods developed in Section 3.4 and summarized here in Fig. 52. The resulting INS solution and radar tracks/ranges are processed in a navigation filter. The navigation filter used is the error-state model EKF developed in Section 3.7.

5.2.1 Radar Data Processing Results. Fig. 53 shows a set of radar data collections for a single stationary corner reflector target at a distance of 1m. We see that we have near perfect time synchronization and very little noise present, as the returns are nearly identical over time. Fig. 54 shows a single slow-time bin of Fig. 53 after pulse compression. We see a greater than 16 dB SNR gain at the range bin of the reflector. Using sub-sample interpolation and alignment, the calculated range of the target from Fig. 54 is 0.995m, showing less than 1cm of error. Performing this experiment multiple times yielded a consistent ranging error of less than 6 cm. Thus with the new receiver board our ranging precision is on the order of 6 cm. Fig. 55 shows the entire data set from Fig. 53 after pulse compression.

Fig. 56 shows the I-channel only pulse compressed data for a moving corner reflector target. The target is initially stationary. It moves 0.5m towards the radar, then it is stopped again. We observe the clear decreasing range from 10-20ns as the target is moving. The target fades in and out of the data set during the time it is moving. This is due to us looking at a single channel. The temporary absence of returns is handled by our M/N detector, which allows reflectors to not appear in every data collection and still be tracked. Thus radar target tracking is still feasible with a single-channel system UWB-OFDM system.

Fig. 57 shows the calculated ranges from the target tracking algorithm for the data set in Fig. 56. We see that sufficient ranges are still available to detect and estimate motion. Fig. 58 shows the phase history of a single corner reflector in a hallway with the radar moving along the hallway wall. We can clearly see the wall feature at 54ns, and the reflector forming a parabolic path at 50ns. From this data set we see that wall guiding would be possible for an indoor environment. However, in this dissertation only point target tracking is used.
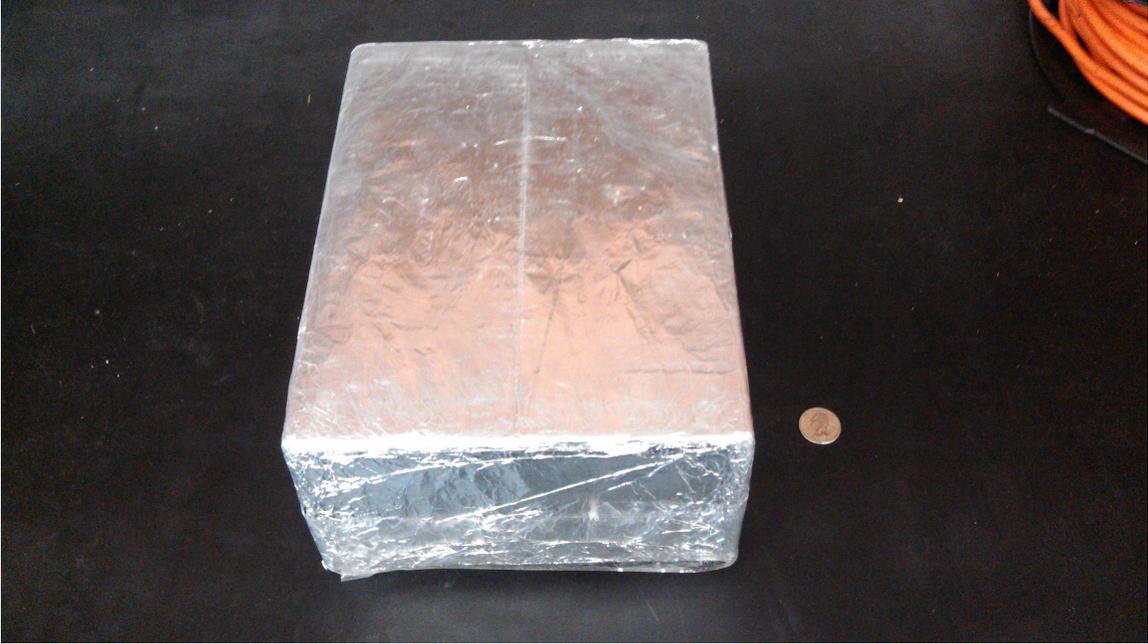
Figure 49:   Rectangular reflectors used in experimental results.

5.2.2   SAR Navigation Results.   The left side of Fig. 60 shows the geometry used for the navigation experiment. The radar is moved through a hallway with 6 reflectors scattered in front of it. (The reflector positions are not initially known, but are calculated as part of the SLAM algorithm). At each turn, the radar is rotated clockwise (from above), so that the antennas are in view of the reflectors. Fig. 59 shows an extract from the data collected in this configuration. The wall feature is visible at 59 ns (Fast Time), and two reflectors are present at 15s and 50s (Slow Time). The reflector at 15s is closer to the wall than the one at 50s, and thus is easily to see.

Fig. 60 shows the computed position of the radar platform for the experimental setup on the left of Fig. 60. We see that with radar aiding we have a significant increase in the accuracy of the position solution. The INS-only solution begins drifting before the platform begins moving. The INS-radar aided solution does not show drift, as the radar has multiple reflectors visible while in initial stationary position which tells the EKF that it is not moving. The stand-alone INS solution has an overall southward bias. The INS also drifts eastward over time, which is corrected by the
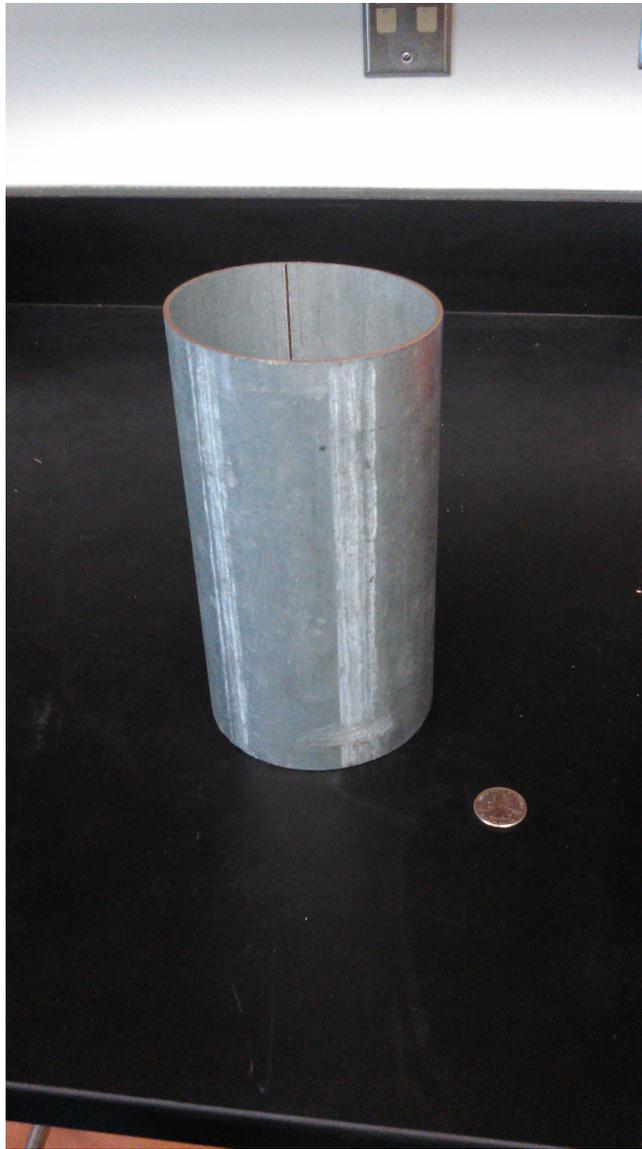
Figure 50:   Cylindrical reflectors used in experimental results.

radar aiding. The final error of the radar-aided solution is 0.3m north and .2m east, compared to the stand-alone INS solution with errors of 3m south and 4m east. Also, due to the nature of INS which involves double integrating acceleration measurements to obtain position, the stand-alone INS solution will typically have increasing error over the short term. In contrast, the radar-aided solution tends to grow as a linear function of distance traveled.
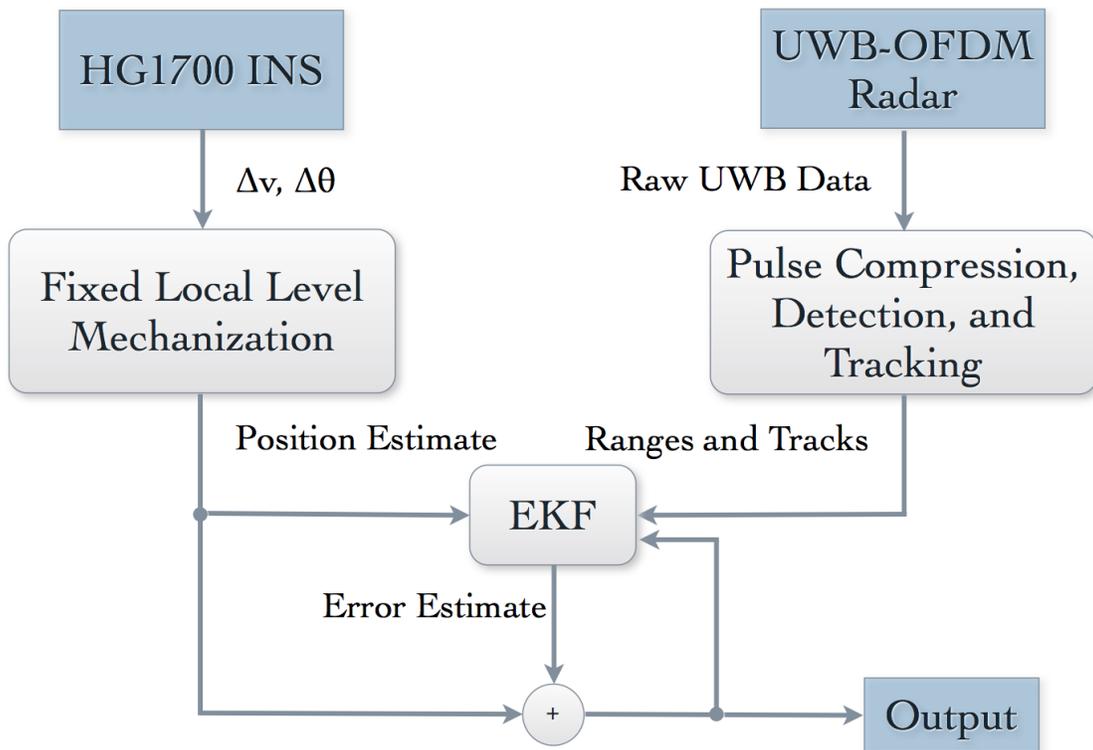
Figure 51:   Overview of the navigation system implemented.

Raw SAR Data → **Feature Extraction** — Pulse Compression, SNR Thresholding

Feature Tracking — Global Nearest Neighbor (GNN)

Feature Range Estimation — Oversampled Matched Filter (Fast-Time Domain)

Track Confirmation — M/N Detector

← Ranges & Associations
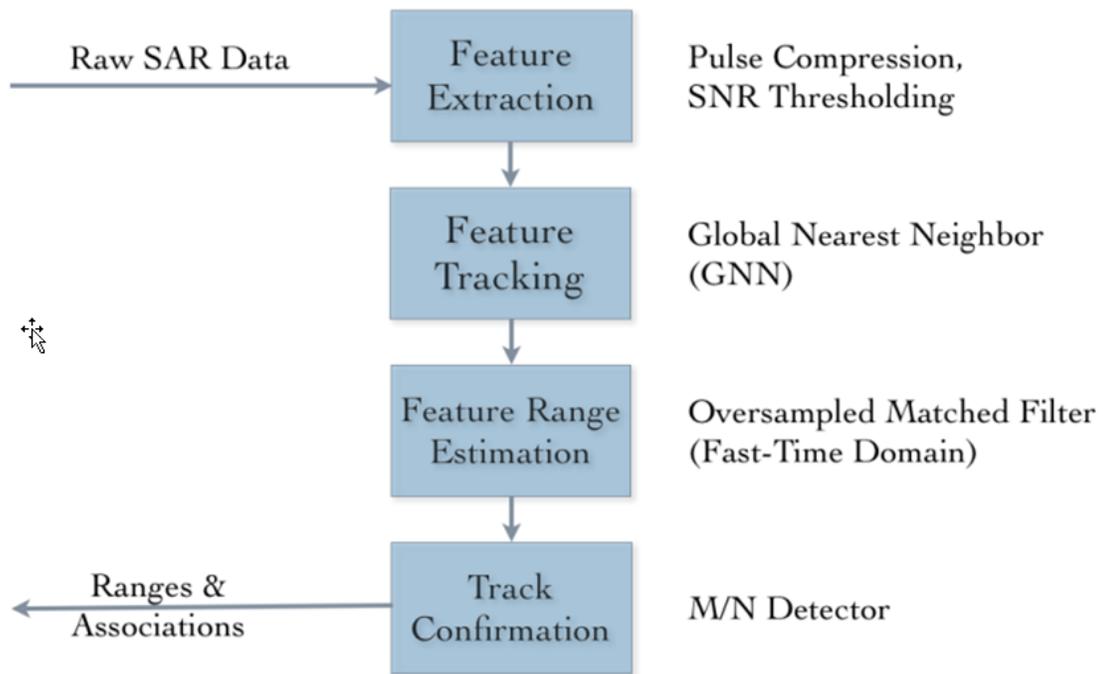
Figure 52:   Overview of radar signal processing method.
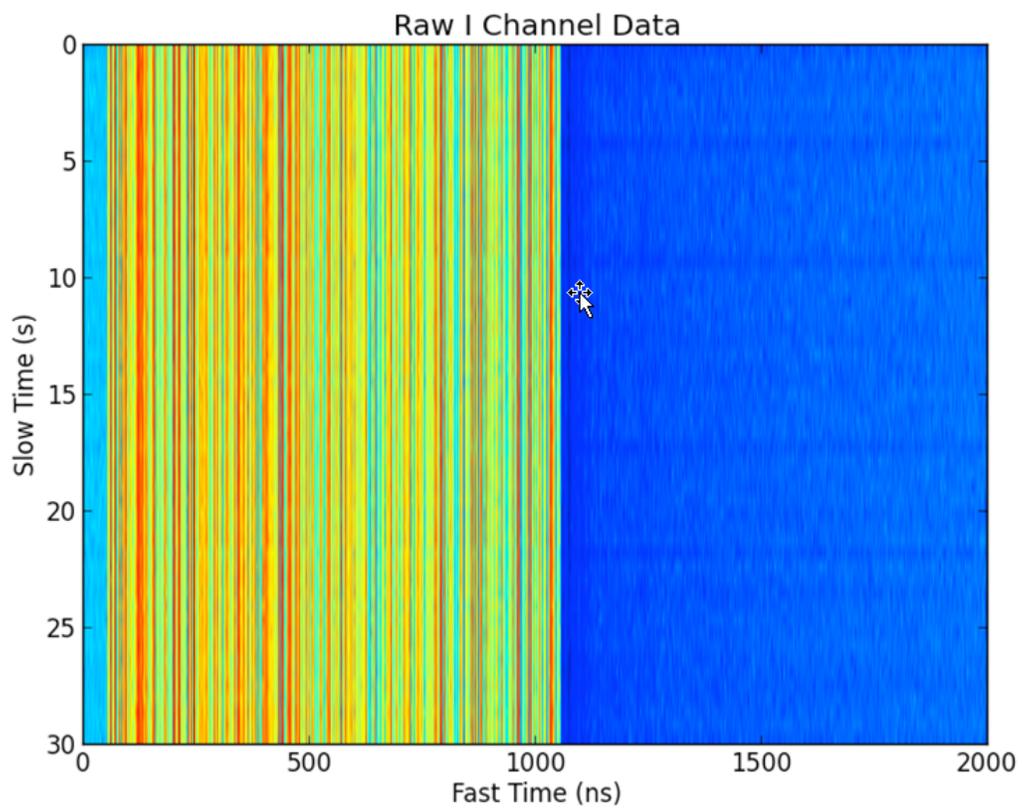
Figure 53: SAR phase history magnitude (observing a single stationary corner reflector).
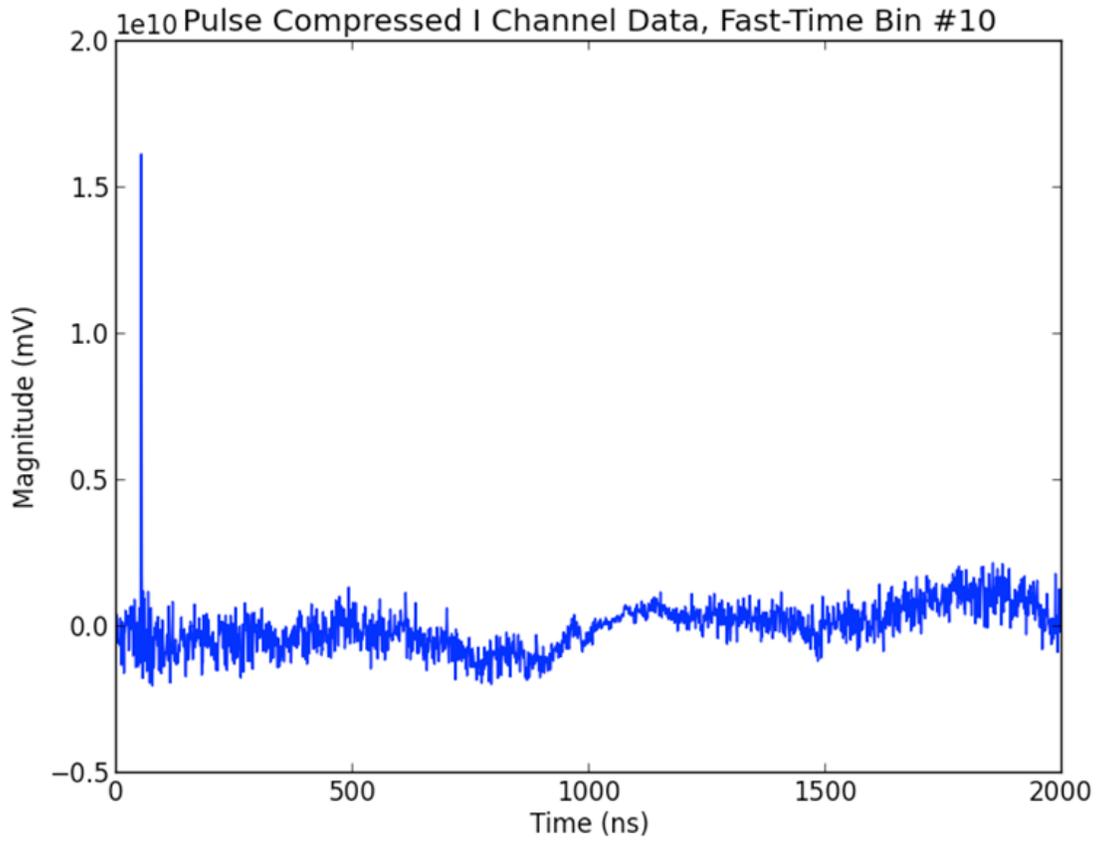
Figure 54:   Fast-time collection after pulse compression (observing a single stationary corner reflector).
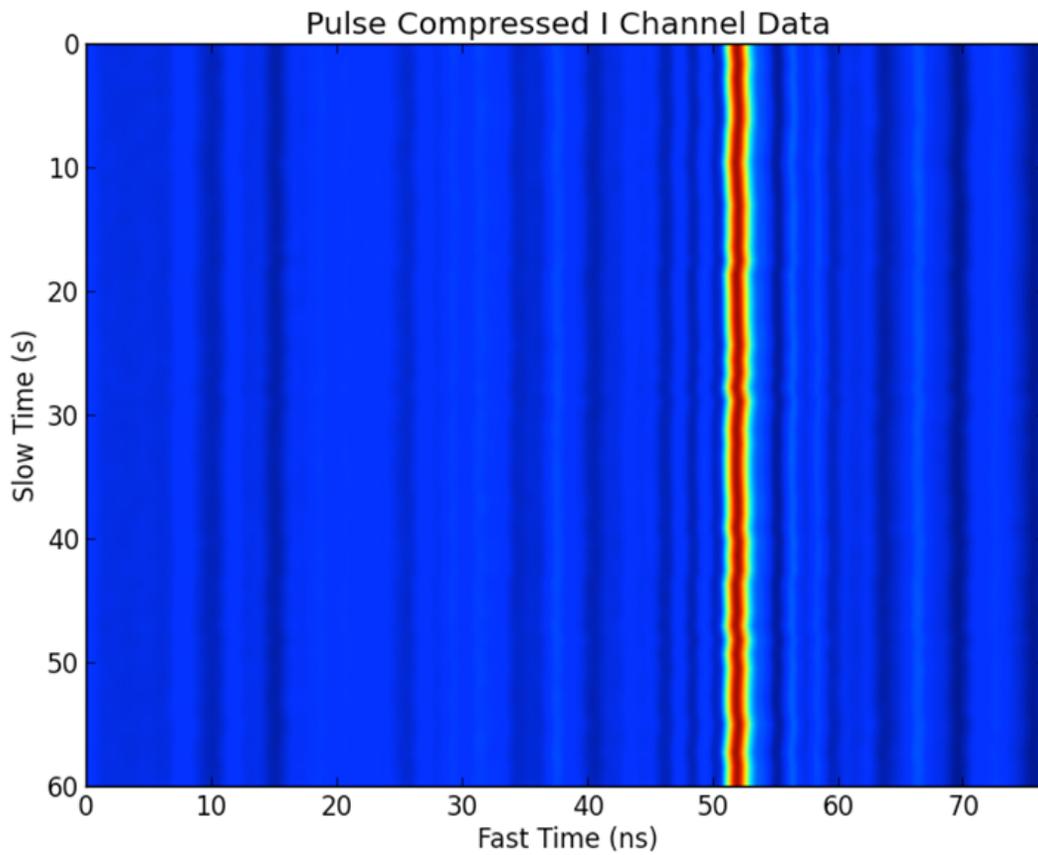
Figure 55: Phase history after pulse compression (observing a single stationary corner reflector).
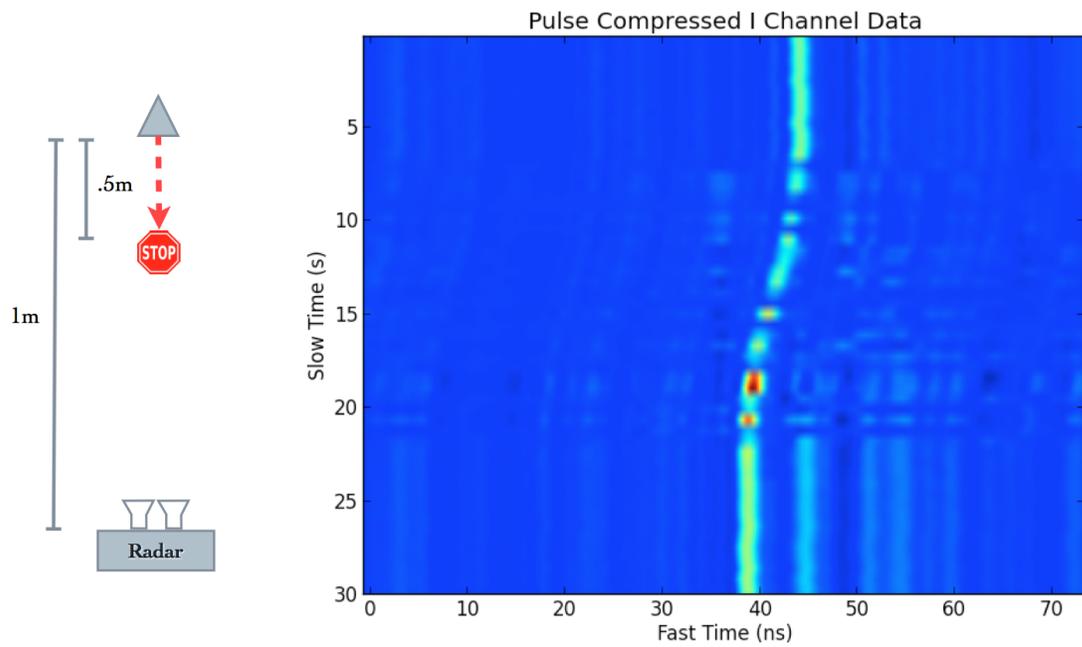
Figure 56: Phase history after pulse compression (observing a single corner reflector moving towards the radar).
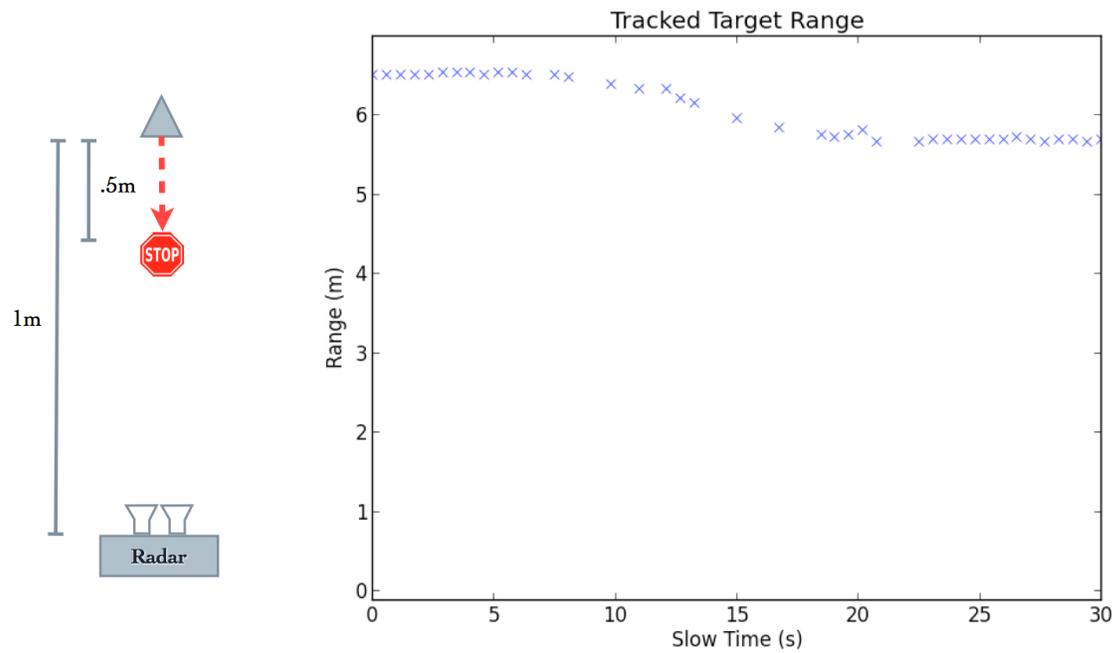
Figure 57: Single track extracted range history for data set in Fig. 6.
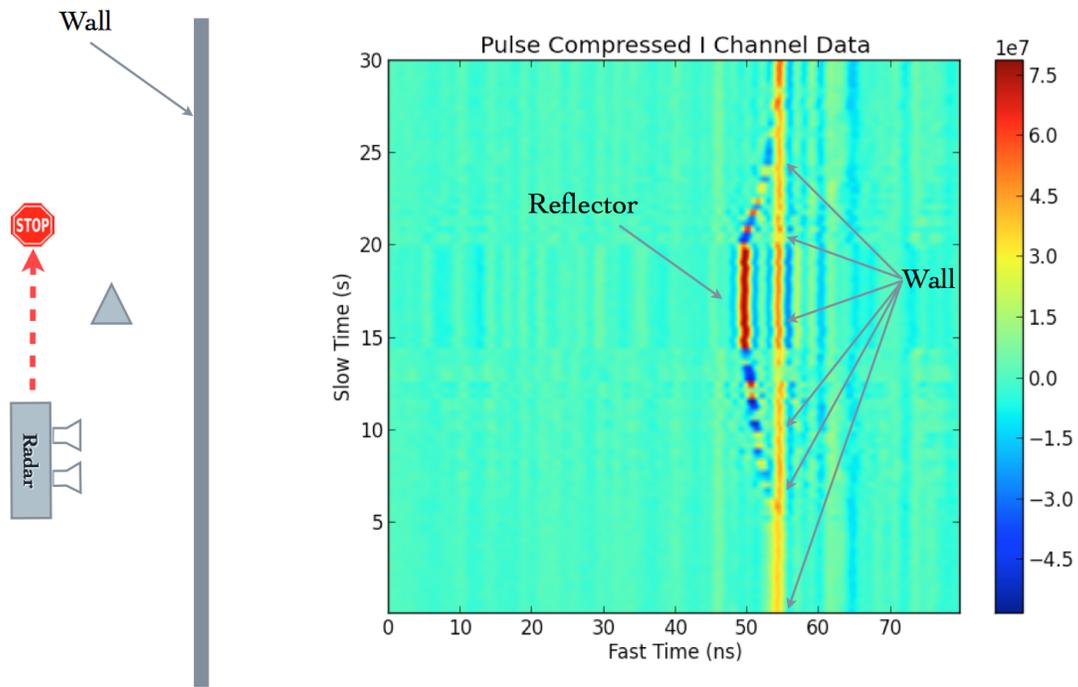
Figure 58:   Phase history after pulse compression for moving radar in hallway with single stationary corner reflector.
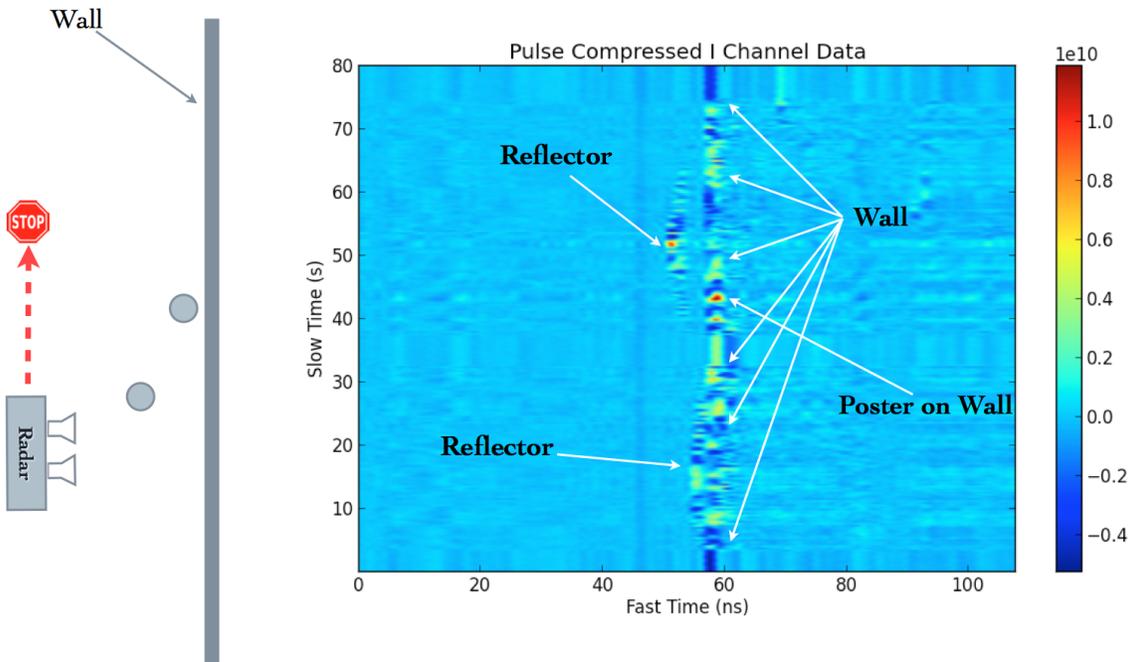
Figure 59:   Phase history after pulse compression. Short sample taken from SAR navigation data set.
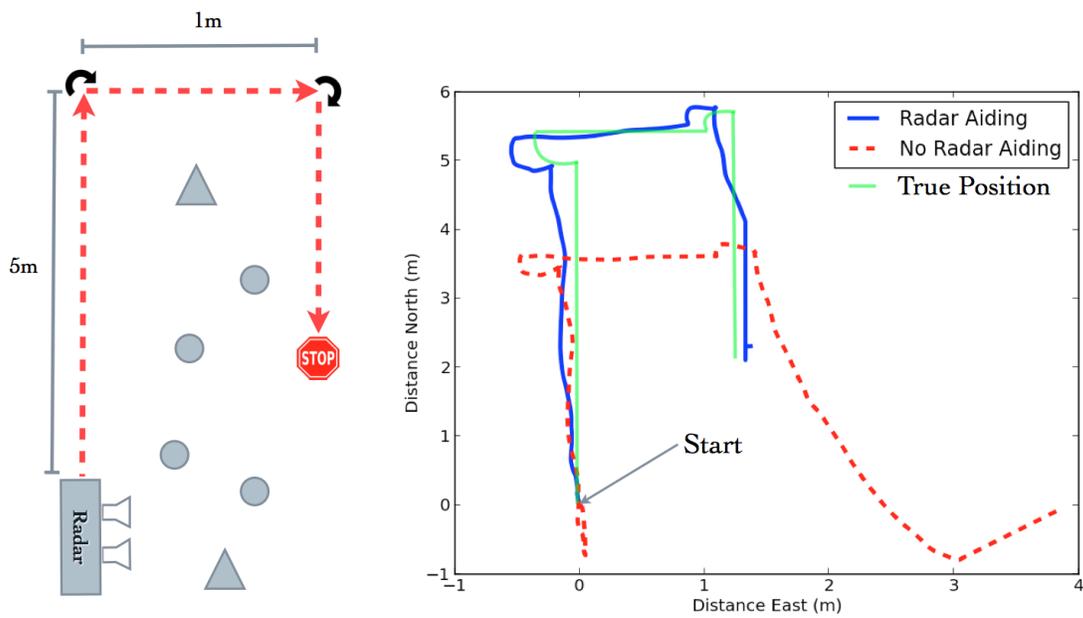
Figure 60: SAR data set computed navigation solutions, shown with and without radar aiding along with true trajectory.

VI. Conclusion

This dissertation presents a novel sensor to allow for always-present, jamming-resistant, active navigation in unknown environments. This chapter presents the conclusions of the study and future potential research in the area of UWB-OFDM and general radar-based navigation. The conclusions of this study support the claimed research contributions stated in Section 1.3.

6.0.3 Conclusions. The goal of this research is to improve the performance of an INS-based navigation platform when other absolute sources of information such as GPS are not available. This is done by augmenting an INS with a UWB-OFDM radar, and using the combined system as an integrated navigation platform. The primary conclusion of this work is the experimental demonstration of the viability of a developed set of algorithms to process the data into a navigation solution. This was done in three processing stages with tight integration: process the INS measurements, process the radar measurements, and then combine the two data sources in an EKF.

In order to implement the SAR data processing stage, a set of algorithms were developed. These algorithms process the raw SAR data by detecting and tracking reflectors embedded within the environment. Range and Doppler measurements associated with the tracked reflectors are recorded for use in the EKF. The INS measurements were processed via a simple fixed local-level mechanization implementation. The mechanization integrates the raw INS measurements (accounting for effects such as the earth rate) into a position/velocity/orientation solution. Lastly, several EKFs were implemented which use the INS solution and/or the measurements to the tracked SAR reflectors to compute an optimal estimate of the AV, depending on sensor availability.

The combined navigation algorithms for the three processing stages were then tested with a variety of airborne simulation and indoor experimental setups. In simulation, we saw an order of magnitude improvement of the navigation solution error when comparing the standalone INS solution and the combined radar/INS solution.

For example, with a tactical grade INS the standalone solution showed an error standard deviation of 2km in the direction of flight after 10 minutes, whereas the integrated solution showed a standard deviation of 2m. For a navigation grade INS, the error went from 20m to sub-meter error. These results show theoretical viability of the system as an airborne navigation sensor.

In addition to the simulation results, we desired to experimentally validate the combined INS/UWB-OFDM navigation filter. In order to achieve this goal, an UWB-OFDM system prototype was first constructed and tuned for use as a navigation sensor. An existing INS platform was then used to generate a navigation solution. In order to process the large SAR data set from the radar, a software suite was developed which controls the hardware systems and implements the previously developed navigation algorithms. The developed software suite runs in real-time on the experimental navigation platform. This demonstrates not only that high data throughput UWB radar navigation is possible, but that it is feasible for real-time navigation with existing hardware. The resulting INS/UWB-OFDM navigation solution using experimental data shows a large improvement over the INS-only solution. The large drift seen in INS-only navigation platforms is greatly reduced by the radar when it has constant access to strong persistent stationary reflectors that do not leave its field of view. For example, when the vehicle was driven down a hallway making a U-shaped turn the INS solution drifted off to greater than 20m of error within a minute, whereas the integrated solution showed less than 0.5m of error. The difference in error grows as a function of time, as the UWB-OFDM showed no drift when the same reflectors were in sight. These results show the large improvement in position solution when using the integrated filter.

As is expected, this solution only performs well when radar has constantly tracked targets. When limited scatterer availability occurs or when the radar must continually track new scatterers, the radar solution will drift like any dead-reckoning system. However, this drift will be in a different direction than that of the INS drift, resulting in a lowered error in the overall navigation solution. Thus the combined

platform will still perform better than a stand-alone INS system, as the drift will be minimized by the addition of the radar.

In summary, this research develops a navigation method using an combined INS/radar system and experimentally validates its use as a navigation sensor. The method shows significant improvement over a stand-alone INS and demonstrates navigation-grade solution errors while using relatively low-cost sensors.

6.0.4 Future Work. There are many areas of further study, including:

- Sensor fusion with additional sensors and more comprehensive sensor platforms, such as image, lidar, or sonar sensors. Since radar is inherently a range-based sensor there is a large opportunity to combine it with an angle-based sensor such as a camera, exploiting the diversity of the different sensors to minimize the navigation error.

- The simulation studies in Chapter IV demonstrate the potential for use in airborne scenarios. The application of these techniques on an airborne SAR system would potentially allow for aerial vehicles to have improved navigation without reliance on GPS or terrain maps.

- The radar signal processing algorithms in this dissertation were computed in a GPU. There is a huge potential for massive parallelization techniques to be applied to this problem, greatly enhancing the tracked radar targets.

- If a fast enough parallelization were developed and implemented, full SAR image construction (instead of partial, as used here) in real-time would be feasible. The techniques here could be extended and used in conjunction with image-aided navigation research to use raw UWB SAR data as an image-based sensor. This would allow for UWB SAR to be used as both an angle and range-based sensor, opening a huge field of potential research.

- OFDM is capable of being used in adaptive anti-jamming scenarios. One problem of interest is the evaluation of UWB-OFDM navigation in hostile envi-

ronments where an intelligent agent actively seeks to interrupt or corrupt the navigation solution.

- Since OFDM is employed in many commonly available communications protocols, there is the potential for using an active OFDM communications system already present on a vehicle as a dual-use communications/navigation sensor. Such a system would transmit information in the OFDM symbols while simultaneously using the multi-path fading as a radar echo response, mapping the environment and performing navigation.

Bibliography

1. D. Titterton and J. Weston, *Strapdown Inertial Navigation Technology.* London, United Kingdom: Peter Peregrinus Ltd., 1997.

2. R. Goddard, "The use of the gyroscope in the balancing and steering of airplanes," *Scientific American*, Jun. 1907.

3. NASA. (2012, Oct.) Nasa - nssdc - spacecraft - details. [Online]. Available: http://nssdc.gsfc.nasa.gov/nmc/spacecraftDisplay.do?id=1978-020A

4. K. Simonsen, M. Suycott, R. Crumplar, and J. Wohlfiel, "Loco gpsi: preserve the gps advantage for defense and security," *Aerospace and Electronic Systems Magazine, IEEE*, vol. 19, no. 12, pp. 3 –7, Dec. 2004.

5. M. Veth, R. Martin, and M. Pachter, "Anti-temporal-aliasing constraints for image-based feature tracking applications with and without inertial aiding," *Vehicular Technology, IEEE Transactions on*, vol. 59, no. 8, pp. 3744 –3756, Oct. 2010.

6. M. Veth and J. Raquet, "Fusion of low-cost imaging and inertial sensors for navigation," in *Proc. ION GNSS*, Fort Worth, TX, Sep. 2006, pp. 1093–1103.

7. T. Hall, "Radiolocation using am broadcast signals: The role of signal propagation irregularities," in *Position Location and Navigation Symposium, 2004. PLANS 2004*, Apr. 2004, pp. 752 – 761.

8. G. Shippey, M. Jonsson, and J. Pihl, "Position correction using echoes from a navigation fix for synthetic aperture sonar imaging," *Oceanic Engineering, IEEE Journal of*, vol. 34, no. 3, pp. 294 –306, Jul. 2009.

9. M. Rabinowitz and J. Spilker, J.J., "A new positioning system using television synchronization signals," *Broadcasting, IEEE Transactions on*, vol. 51, no. 1, pp. 51 – 61, Mar. 2005.

10. A. Varshavsky, M. Y. Chen, E. de Lara, J. Froehlich, D. Haehnel, J. Hightower, A. LaMarca, F. Potter, T. Sohn, K. Tang, and I. Smith, "Are gsm phones the solution for localization?" in *Mobile Computing Systems and Applications, 2006. WMCSA '06. Proceedings. 7th IEEE Workshop on*, Apr. 2006, pp. 20 – 28.

11. C. Toth, D. Grejner-Brzezinska, and Y.-J. Lee, "Terrain-based navigation: Trajectory recovery from lidar data," in *Position, Location and Navigation Symposium, 2008 IEEE/ION*, May 2008, pp. 760 –765.

12. M. Skolnik, *Radar Handbook*, 3rd ed. New York: McGraw-Hill, 2008.

13. M. Veth, "Fusion of imaging and inertial sensors for navigation," Ph.D. dissertation, AFIT, Dayton, OH, 2006.

14. C. Jakowatz, D. Wahl, P. Eichel, D. Ghiglia, and P. Thompson, *Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach.* NY: Springer Science, 1996.

15. M. Skolnik, *Introduction to Radar Systems*, 3rd ed. New York: McGraw-Hill, 2001.

16. J. Schuerger and D. Garmatyuk, "Deception jamming modeling in radar sensor networks," in *Proc. Military Communications Conference (MILCOM)*, San Diego, CA, Nov. 2008.

17. M. Frigo and S. Johnson, "The design and implementation of fftw3," *Proceedings of the IEEE*, vol. 93, no. 2, pp. 216 –231, feb. 2005.

18. D. E. Knuth, "Big omicron and big omega and big theta," *SIGACT News*, vol. 8, pp. 18–24, April 1976. [Online]. Available: http://doi.acm.org/10.1145/1008328. 1008329

19. Z. C. K. H. TODO, "Fusion of imaging and inertial sensors for navigation," Ph.D. dissertation, AFIT, Dayton, OH, 2006.

20. R. Brown and P. Hwang, *Introduction to Random Signals and Applied Kalman Filtering*, 3rd ed. NJ: Wiley, 1996.

21. M. Soumekh, *Synthetic Aperture Radar Signal Processing.* NJ: Wiley, 1999.

22. D. Gabor, "A new microscope principle," *Nature*, vol. 161, no. 777, 1948.

23. L. Cutrona, E. Leith, L. Porcello, and W. Vivian, "On the application of coherent optical processing techniques to synthetic-aperture radar," *Proc. IEEE*, vol. 54, pp. 1026–1032, Aug. 1966.

24. D. Ausherman, A. Kozma, J. Walker, H. Jones, and E. Poggio, "Developments in radar imaging," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 20, 1984.

25. J. A. Parker, *Image Reconstruction in Radiology.* CRC Press, 1990.

26. R. Mersereau, "Recovering multidimensional signals from their projections," *Computer Graphics and Image Processing*, vol. 1, pp. 179–195, Oct. 1973.

27. J. Dezert, "Improvement of strapdown inertial navigation using pdaf," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 35, no. 3, pp. 835 –856, Jul. 1999.

28. P.-J. Nordlund and F. Gustafsson, "Marginalized particle filter for accurate and reliable terrain-aided navigation," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 45, no. 4, pp. 1385 –1399, Oct. 2009.

29. A. Henley, "Terrain aided navigation: current status, techniques for flat terrain and reference data requirements," in *Position Location and Navigation Symposium, 1990. Record. The 1990's - A Decade of Excellence in the Navigation Sciences. IEEE PLANS '90., IEEE*, Mar. 1990, pp. 608 –615.

30. L. Hostetler and R. Andreas, "Nonlinear kalman filtering techniques for terrain-aided navigation," *Automatic Control, IEEE Transactions on*, vol. 28, no. 3, pp. 315 – 323, Mar. 1983.

31. J. Kim and S. Sukkarieh, "Autonomous airborne navigation in unknown terrain environments," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 40, no. 3, pp. 1031 – 1045, Jul. 2004.

32. R. H. Mayer, "Doppler navigation for commercial aircraft in the domestic environment," *Aerospace and Navigational Electronics, IEEE Transactions on*, vol. ANE-11, no. 1, pp. 8 –15, march 1964.

33. P.-M. Lee, B.-H. Jun, K. Kim, J. Lee, T. Aoki, and T. Hyakudome, "Simulation of an inertial acoustic navigation system with range aiding for an autonomous underwater vehicle," *Oceanic Engineering, IEEE Journal of*, vol. 32, no. 2, pp. 327 –345, april 2007.

34. N. Braverman, "Self-contained navigation aids and the common system of air traffic control," *Aeronautical and Navigational Electronics, IRE Transactions on*, vol. ANE-4, no. 2, pp. 52 –56, june 1957.

35. D. Sheen and L. Johnston, "Statistical and spatial properties of forest clutter measured with polarimetric synthetic aperture radar (sar)," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 30, no. 3, pp. 578 –588, may 1992.

36. R. Kapoor, A. Banerjee, G. Tsihrintzis, and N. Nandhakumar, "Uwb radar detection of targets in foliage using alpha-stable clutter models," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 35, no. 3, pp. 819 –834, jul 1999.

37. B. Werle, "Sea backscatter, spikes and wave group observations at low grazing angles," in *Radar Conference, 1995., Record of the IEEE 1995 International*, may 1995, pp. 187 –195.

38. H. Melief, H. Greidanus, P. van Genderen, and P. Hoogeboom, "Analysis of sea spikes in radar sea clutter data," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 44, no. 4, pp. 985 – 993, april 2006.

39. J. Marcum, "A statistical theory of target detection by pulsed radar," *Information Theory, IRE Transactions on*, vol. 6, no. 2, pp. 59 –267, april 1960.

40. D. Shnidman, "Generalized radar clutter model," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 35, no. 3, pp. 857 –865, jul 1999.

41. A. Farina, A. Russo, and F. Studer, "Coherent radar detection in log-normal clutter," *Communications, Radar and Signal Processing, IEE Proceedings F*, vol. 133, no. 1, pp. 39 –53, february 1986.

42. D. Garmatyuk, J. Schuerger, and K. Kauffman, "Multifunctional software-defined radar sensor and data communication system," *Sensors Journal, IEEE*, vol. 11, no. 1, pp. 99 –106, Jan. 2011.

43. D. Garmatyuk, "Simulated imaging performance of uwb sar based on ofdm," in *Proc. IEEE ICUWB*, Waltham, MA, 2006.

44. D. Garmatyuk and K. Kauffman, "Radar and data communication fusion with uwb-ofdm software-defined system," in *Ultra-Wideband, 2009. ICUWB 2009. IEEE International Conference on*, Sep. 2009, pp. 454 –458.

45. P. Maybeck, *Stochastic Models, Estimation and Control.* Arlington, Virginia: Navtech Press, 1994, vol. 1.

46. P. Swerling, "Probability of detection for fluctuating targets," *Information Theory, IRE Transactions on*, vol. 6, no. 2, pp. 269 –308, Apr. 1960.

47. S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems.* Norwood, MA: Artech House, 1999.

48. D. Shnidman, "Binary integration for swerling target fluctuations," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 34, no. 3, pp. 1043 –1053, Jul. 1998.

49. D. Chung, J. G. Lee, C. G. Park, and H. W. Park, "Strapdown ins error model for multiposition alignment," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 32, no. 4, pp. 1362 –1366, oct 1996.

50. B. Scherzinger and D. Reid, "Modified strapdown inertial navigator error models," in *Position Location and Navigation Symposium, 1994., IEEE*, apr 1994, pp. 426 –430.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704–0188

1. REPORT DATE (DD–MM–YYYY)	2. REPORT TYPE	3. DATES COVERED (From — To)
31-12-2012	Doctoral Dissertation	July 2009 – December 2012

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Radar Based Navigation in Unknown Terrain	
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)	5d. PROJECT NUMBER
Kauffman Kyle J	JON10-516,10-11,11-516,12-518
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/ENY) 2950 Hobson Way WPAFB OH 45433-7765	AFIT/ENG/DS/12-03

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
DAGSI Dr. Elizabeth Downie edownie@dagsi.org	
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

13. SUPPLEMENTARY NOTES
This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

14. ABSTRACT

There is a great need to develop non-GPS based methods for positioning and navigation in situations where GPS is not available. This research focuses on the development of an Ultra-Wideband Orthogonal Frequency Division Multiplexed (UWB-OFDM) radar as a navigation sensor in GPS-denied environments. A side-looking vehicle-fixed UWB-OFDM radar is mounted to a ground or aerial vehicle continuously collecting data. A set of signal processing algorithms and methods are developed which use the raw radar data to aide in calculating the vehicle position and velocity via a simultaneous localization and mapping (SLAM) approach. The radar processing algorithms detect strong, persistent, and stationary reflectors embedded in the environment and extract range/Doppler measurements to them. If the radar is the only sensor available, the measurements are used to directly compute the vehicle position. If an existing navigation platform is available, the measurements are combined with the other sensors in an EKF. The developed algorithms are tested via both a series of airborne simulations and a ground-based experiment. The computed navigation solution performance is analyzed with the following sensor availability: radar-only, INS-only, and combined radar/INS. In both simulation and experimental scenarios, the integrated INS/UWB-OFDM system shows significant improvements over an INS-only navigation solution.

15. SUBJECT TERMS
Radar, navigation, SLAM, SAR, EKF

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE	UU	131	Dr. John Raquet AFIT/ENG
U	U	U			19b. TELEPHONE NUMBER (Include Area Code) (937)255-3636, ext 4580

Standard Form 298 (Rev. 8–98)
Prescribed by ANSI Std. Z39.18