Magnetometer Aided Navigation Filters for Improved Observability and Estimation on Ground Vehicles

by

Gabriel Morales

A thesis submitted to the Graduate Faculty of
Auburn University
in partial fulfillment of the
requirements for the Degree of
Master of Science

Auburn, Alabama
August 6, 2016

Keywords: Magnetometer, Extended Kalman Filter, Closely-Coupled

Approved by

David Bevly, Chair, Professor of Mechanical Engineering
George Flowers, Professor of Mechanical Engineering and Dean of the Graduate School
John Hung, Professor of Electrical and Computer Engineering

Abstract

Two modified extended Kalman filters, called the MGVA and MVVA filters, have been developed which combine the attitude determination capabilities of magnetometers with the standard extended Kalman filter. The standard extended Kalman filter is known to experience observability problems during driving which does not provide enough excitation to its sensors. The MGVA and MVVA filters attempt to remedy these problems by providing an attitude solution to the filter, via a process that utilizes magnetometer measurements. In order to show the effectiveness of the MGVA and MVVA filters, the filters were tested both in simulation, and experimentally for low excitation trajectories. It is shown in simulation that the two filters are able to improve attitude and accelerometer bias estimation in situations where the standard filter experiences estimation error. In addition to this, it is shown that both modified filters increase the observability rank of the system such that it is full rank. Experimental testing of the two modified filters reveals that both the MGVA and MVVA filters provide improved attitude and accelerometer bias estimation during low-excitation trajectories.

The MGVA and MVVA filters are also compared to a heading constrained filter that has been studied in the past. It is shown that the heading constrained filter produces more accurate estimates of heading than both of the modified filters. As a result, a further modification is implemented on both the MGVA and MVVA filters, and improved performance is shown for the MGVA filter. Finally, the MGVA and MVVA filters were tested on both simulated and experimental dynamic trajectories. This was done in order to test whether or not the filters are useful under typical trajectories that a ground vehicle might drive, not just low excitation trajectories. It is shown that under these circumstances the MGVA filter is suffers from the increased vehicle dynamics. In contrast, it is also shown that the MVVA

filter is, once again, able to estimate attitude and accelerometer bias well. It is concluded from these tests that the MVVA filter is, in general preferable to the MGVA filter for more dynamic trajectories.

Acknowledgments

I would like to begin by acknowledging my graduate advisor, Dr. Bevly, who has never ceased to challenge me both as a student and a researcher. Without him, I would never have aimed or dreamed of being where I am today. It is because of his faithful belief in his students, and his obvious desire to see them succeed, that they are capable of going above and beyond their greatest ambitions. It is for this that I am highly thankful.

I would also like to acknowledge my parents, who have never wavered in their faith and belief in me and my future. Their continual love, support, and sacrifice has been a foundation that has sustained me through all the seasons of life that I have ventured through till now. Additionally, I would like to acknowledge my brother and sister, who have believed in me and continually encouraged me to strive for better and brighter achievements in my own life.

It is not possible to overstate how grateful I am to the other GAVLAB graduate students who I have ventured through this season of life with. I would like to acknowledge Dan, Grant, and Scott, who have offered their own professional and technical insight on countless occasions throughout the work that has been done. I would also like to thank Steven, Tyler, Josh, and Tripp for the effort and assistance they provided in preparing this work for its final state. Additionally, I'd like to thank Josh and his wife Katie for being friends who continually demonstrate their friendship in word and deed. I'd like to thank Eick and Brently for keeping me and my coworkers laughing during strenuous days, and I'd like to thank Russell for making sure that everyone knows ping-pong is life to a GAVLAB student.

I'd like to also thank the brothers and sisters in Christ that I've had the privilege to share life with at the Auburn Wesley Foundation and at Auburn Community Church. Their friendship has been a true witness of the love of Christ to me during my time as a graduate

student. The time I've spent surrounded by them has challenged, molded, and shaped who I am in ways that can't be described.

Above all, I'd like to also give thanks to the Lord Jesus Christ, who gives meaning and purpose to every action and work of my life. Without His work of grace in me, nothing that has happened would be possible.

Table of Contents

List of Figures

Chapter 1

Introduction and Background

1.1 Motivation, Objectives, and Contribution

In recent years, there has been an extensive amount of research performed in the areas of GPS/INS based navigation systems. This is primarily due to an increased interest and need for high integrity navigation systems in both civilian and military applications on ground, aerial, and even aquatic vehicles [1–4]. However, as much as GPS/INS systems have been used to provide accurate navigation solutions for many applications, it is apparent that both GPS systems and INS systems suffer from several flaws that diminish their reliability in various contexts. GPS receivers require line-of-sight visibility to GPS satellites in order to form an accurate position, velocity, and attitude (PVA) solution and it is well known that GPS receivers will suffer from multipath, unavoidable signal outages in urban canyons, dense foilage, or other undesirable environments [5]. Additionally, most consumer grade INS systems will suffer from common sensor errors such as stochastic wide-band noise, sensor bias, and sensor drift that will cause the INS PVA solution to drift over time [1]. As a result, a large amount of literature has covered the use of alternative navigation methods for use in aiding or completely replacing GPS or INS systems during PVA solution degradation.

Several different measurement based systems have been used in the past to mitigate situations where GPS and INS systems become unreliable. In [6], the authors explored the effectiveness of terrestrial passive ranging systems, direct measuring equipment (DME), and multilateration based navigation systems on aerial vehicles as alternatives to GPS navigation. In [8], it was demonstrated that WiFi positioning could be used as a GPS alternative to correct for drift in INS systems during indoor pedestrian navigation. In addition to these works, it has more recently been proposed that novel landmarks in operating environments

may be used to perform real-time navigation during GPS outages [7]. Another substitute for INS dependent navigation is vehicle model based navigation. In [9], the authors implement and demonstrate the use of a four wheel vehicle model on a dual extended Kalman filter for ground vehicle navigation. Visual odometry is another method that has been explored in detail for aiding GPS/INS systems. In [10], the authors show the capability of a visual odometry system on a ground vehicle to produce accurate motion information for the vehicle without any assistance from GPS or INS systems. These methods are just a few of the many techniques which have been examined for aiding navigation systems during situations where GPS or INS systems become unreliable.

In addition to this list of GPS/INS aiding sensors and substitutes, magnetic sensors have often been used in literature to aid GPS/INS systems in heading stabilization and attitude determination. In [11], Caruso briefly describes the fundamentals of magnetic technology, and discusses their usefulness in heading calculation. In [12], a magnetic heading system was designed and tested that reduced heading estimate errors from several degrees of error to sub-degree errors. Additionally, in [13], an attitude determination system for low earth orbit satellites was developed which relied entirely on three-axis magnetometer measurements.

While this is far from an exhaustive list of methods which are available, it can be clearly seen from this list that GPS/INS systems are often in need of novel solutions to the inherent navigation problems that they face. Of all the problems that GPS/INS systems encounter, one of the most difficult problems is that of error source estimation and correction. In particular, many researchers have attempted to enhance estimation of stochastic sensor errors in inertial measurement units (IMUs). In [16], random constant, random walk, Gauss-Markov, and autoregressive stochastic models were all described, and their implementation on a loosely-coupled navigation filter for error estimation was demonstrated. In [17], the authors designed and tested a new multi-position sensor calibration method and then compared the results from their method with those from other traditional calibration methods.

A corresponding problem which follows from an investigation into GPS/INS system performance is the question of navigation filter observability. In particular, the topic of Kalman filter observability is a subject that has been extensively covered in the past. In [18], Goshen-Meskin and Bar-Itzhack developed a theoretical approach to the observability analysis of linear time-varying systems. In [20], Hong showed that attitude errors, gyro bias errors, and GPS lever arm errors cause position, velocity, and accelerometer bias estimate errors to become unobservable. In [21], Rhee investigated the state observability of the standard Kalman filter for various aerial maneuvers. It was shown that periods of straight motion at constant speed cause the attitude of the system to become unobservable to the filter, while excitation of various accelerometer modes, either by steady turning or by changes in acceleration, cause various attitude and bias states to become observable. Finally in [23], Ryan investigated the effect of the trajectories explored in [21] on a modified Kalman filter that only utilized sensors commonly found in electronic stability control (ESC) systems on ground vehicles. In doing so, Ryan emphasized several points which Rhee previously noted. First, one of the strongest inherent couplings in the standard Kalman filtering navigation equations exists between the attitude states of the filter and the accelerometer bias estimates. Following from this he also noted that this implies accurate estimation of accelerometer bias states if the attitude states of the filter can be estimated.

In light of all of the problems faced by GPS/INS navigation systems and the previous work which has been done to mitigate these errors, it is the goal of this thesis to provide a solution to the specific problems of sensor error estimation and filter observability. In this thesis, a magnetometer will be fused with a standard GPS/INS navigation filter to create two modified filters which shall be called the Magnetic and Gravity Vector Aided (MGVA) filter and the Magnetic and Velocity Vector Aided (MVVA) filter. In implementing these filters, this thesis seeks to provide better attitude estimation during trajectories where the attitude is otherwise unobservable. In doing so, it is theorized that this will lead to an increase in the ability of the filter to estimate accelerometer biases which corrupt the inputs to the system.

Finally, as a result of the enhanced estimation abilities of these modified filters, this thesis seeks to also improve the overall observability of the standard GPS/INS navigation filter with the proposed modified filters.

As can be seen, there is a wealth of prior art which precedes this thesis. The areas of GPS/INS navigation, error estimation, and sensor fusion have been studied in great detail. Additionally, the subject of attitude determination via vector observations has been researched thoroughly, and several of the methods used in this thesis have been used and suggested by others in the past. However, this work sets itself apart from the prior art with the following contributions:

- Observability analysis of the MGVA and MVVA filters for low-excitation trajectories on ground vehicles
- Analysis of the attitude and accelerometer bias estimation capabilities of the two filters on low-excitation trajectories
- Comparison of the two filters to a heading constrained filter studied by another author [23]
- Performance analysis of the two filters for dynamic ground vehicle trajectories

In order to give adequate background and context to the work which is to follow, it is essential to have a fundamental understanding of the magnetic sensing technology that this thesis is dependent on. As a result, the following sections will give a brief description of the technology, discuss its uses as a navigational sensor and also explain inherent flaws and error sources that these sensors suffer from.

1.2 Magnetic Sensing Technology

Magnetic sensors (or magnetometers) are sensors that, simply put, measure magnetic fields. There are, in general, two types of magnetometers: scalar magnetometers and vector

magnetometers. Scalar magnetometers are capable of measuring the magnitude of magnetic fields they are immersed in, while vector magnetometers are capable of measuring the magnitude and direction of a magnetic field. Magnetometers have been used extensively in the past for a wide range of applications such as space exploration, Earth magnetic field surveying, submarine detection, and mineral exploration. One of their most well known uses in the medical field is in magnetic resonance imaging (MRI) machines [24].

In addition to all of these applications, magnetometers are also very capable sensors in the navigation industry. It is well known that magnetometers are highly capable of producing accurate heading measurements [11]. In addition to their use as heading sensors, magnetometers have also been used frequently for map matching in various types of vehicle navigation and pedestrian navigation systems [25–27]. Finally, magnetometers have been frequently used for six degree of freedom (DoF) attitude determination in cooperation with other vector sensors [13], [14], [28]. For this thesis, magnetometers were used both for heading stabilization and for six DoF attitude determination. As a result, a brief explanation on the use of magnetometers for these two processes is given in the following sections.

1.2.1 Heading Calculation

The majority of the Earth's magnetic field is created internal to the center of the earth, and expresses itself similar to a large magnet with two poles as shown in Figure 1.1. This is, of course, a simplified way of discussing the Earth's magnetic field. The actual magnetic field is a variable field which changes over time and with location. It is important to note that the direction of the Earth's magnetic field does not always coincide with geographic north, and it is usually not horizontal to the surface of the earth. Because of this, the Earth's field is described by a declination angle (angular difference between magnetic north and geographic north), and an inclination angle (angular deviation from the horizontal).

Because magnetometers measure any magnetic field which they are immersed in, they are very useful sensors for measuring the Earth's magnetic field anywhere on the globe.

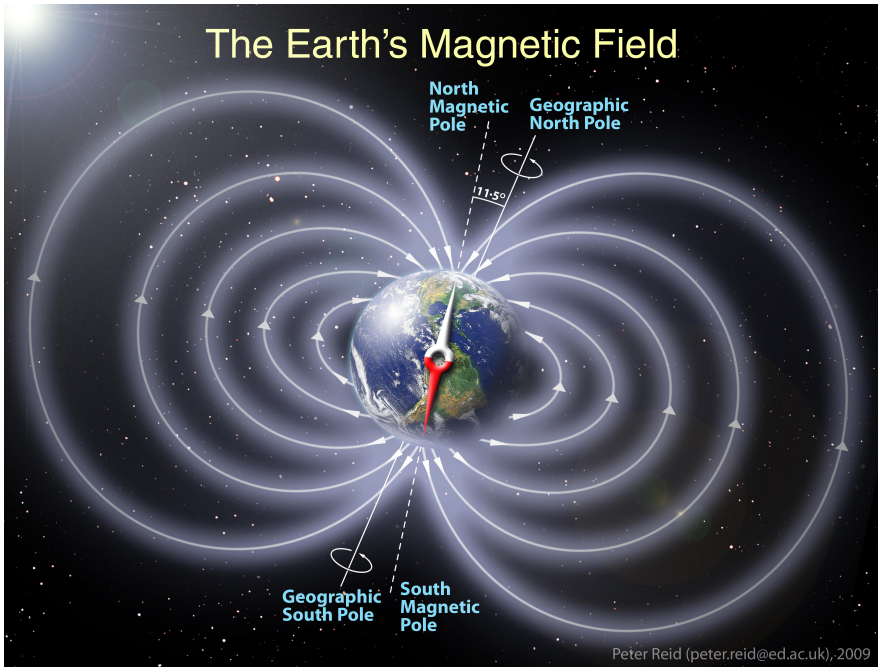

Figure 1.1: Illustration of the Earth's magnetic field [31]

In addition to this, vector magnetometers are particularly useful for navigation purposes because they can "see" the direction of the Earth's magnetic field in their own coordinate frame. This capability in turn allows them to estimate the angular difference between the Earth's field direction and their own orientation in the horizontal plane. Mathematically, this expresses itself as in Equation (1.1)

$$\psi = atan2(\frac{b_y}{b_x}) \tag{1.1}$$

where ψ is heading, b_y is the magnetic field measurement in the y-direction, and b_x is the magnetic field measurement in the x-direction.

While this approach is fairly straightforward, it is important to recognize that measurements made by the magnetometer in its own respective frame will often not reside perfectly in the horizontal plane. As a result, before heading estimates can be formed using Equation (1.1), all magnetometer measurements must be rotated into the horizontal plane via the roll

and pitch tilt angles of the system, as in Equations (1.2) and (1.3)

$$b_{Hx} = b_x cos(\theta) + b_y sin(\phi) sin(\theta) - b_z cos(\phi) sin(\theta) \tag{1.2}$$

$$b_{Hy} = b_y cos(\phi) + b_z sin(\phi) \tag{1.3}$$

where b_z is the z-direction magnetic field measurement, ϕ defines the roll of the body, and θ defines the pitch, and the b_H terms represent the magnetic measurements leveled in the horizontal plane.

1.2.2 Attitude Determination via Vector Observations

The problem of determining attitude from vector observations is well known and has been solved many times. In fact, there are many algorithms available for the process, and faster algorithms for real-time calculations continue to be sought [32]- [36]. The problem was first considered by Grace Wahba in 1965 and is commonly referred to as Wahba's problem.

The driving concept behind the problem and its solution is as follows. Given two coordinate frames in space, which are both able to observe two separate and distinct points, it is possible, via vector observations of these points, to form estimates of the angular rotation between the two coordinate frames. Figure 1.2 helps to shed some light on this concept. Figure 1.2a depicts two different coordinate frames. The first is oriented along the x, y, and z axes of the image, while the second is rotated by 45 degrees about each axis. Figure 1.2b depicts two separate and distinct points, A and B, in the first coordinate system. Point A is located along the system x-axis, and point B is located along the system y-axis. As a result, the two vectors which point to A and B are also aligned with the x and y axes. Figure 1.2c depicts points A and B as they are seen in the rotated coordinate frame, along with two different vectors which point to them. It can be shown (and will be shown in Chapter 3) that the same rotation between the first and second coordinate systems will also rotate the first

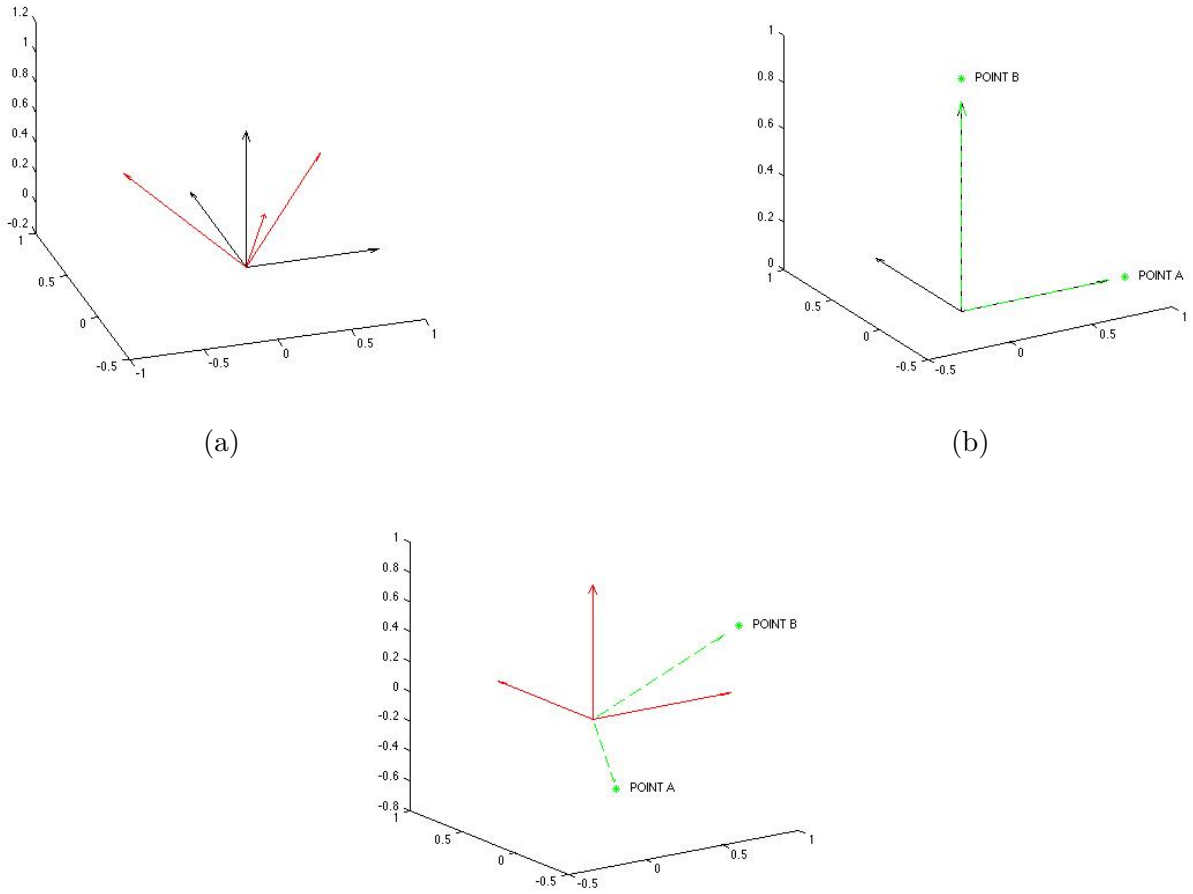

(a)

(b)

(c)

Figure 1.2: Illustration depicting vector measurements of two points in two different coordinate systems

set of vectors into the second set of vectors. Calculation of this rotation, therefore, becomes a problem of estimation given four different vector measurements.

In the navigation industry, the solution of this problem is particularly useful for predicting the rotation (or attitude) between a reference coordinate system and a user coordinate system. Vector measurement devices such as magnetometers, sun sensors, star sensors, and horizon sensors are all commonly used for spacecraft attitude determination [38]. Other inertial sensors have also been used for attitude determination under more static conditions [14].

1.2.3 Magnetometer Error Sources

However useful magnetic sensors are for attitude determination, like all sensors, they still suffer from many inherent flaws that must be accounted for if they are to be used for navigation. In this section, a brief discussion will be given on the several different kinds of errors that must be accounted for.

Many of the error sources which are encountered by magnetometers are very similar to those seen using inertial sensors. Like inertial sensors, magnetometers suffer from the effects of stochastic bias and wide-band noise on their measurements. Magnetometers also suffer from the effects of misalignment error. It is crucial that the axes of a given magnetometer align with the axes of the coordinate system it resides in; otherwise, the vector measurements it collects will not truly reside in the user coordinate system. Additionally, magnetometers may suffer from measurement drift as a result of temperature effects if the sensors are used outside of their operating range.

While all of these errors do serve to corrupt magnetic measurements, the errors that are most prevalent in magnetic sensors are those due to magnetic distortion. Magnetometers will always detect the Earth's magnetic field, but in addition to the Earth's magnetic field, they will also measure all other magnetic activity in their vicinity. Errors due to magnetic distortions are often categorized into hard-iron and soft-iron biases. Hard-iron biases are

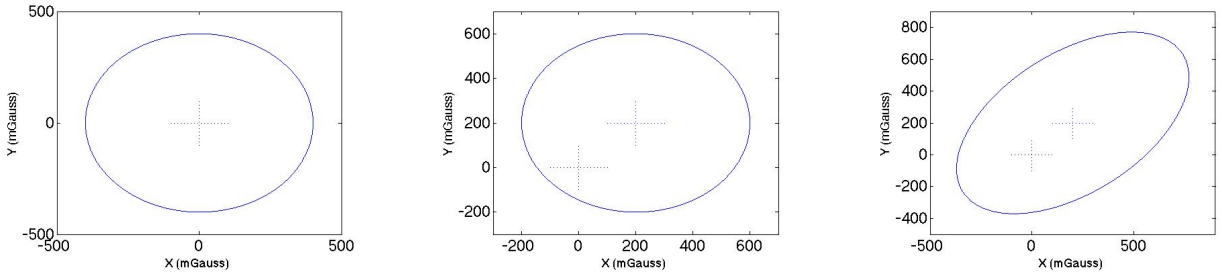

(a) Locus of perfect data set

(b) Locus of data corrupted by hard-iron bias

(c) Locus of data corrupted by hard and soft-iron bias

Figure 1.3: Effects of hard and soft-iron bias on magnetometer data

caused by any object which generates its own magnetic field in the vicinity of the magnetometer, while soft-iron bias is caused by any object which distorts a magnetic field which it is immersed in.

The nature of these error sources is made clearer when they are viewed through loci of magnetic field data. If a magnetometer is rotated 360 degrees only in the Earth's magnetic field, it will produce a locus of data like that shown in Figure 1.3a. However, hard-iron biases from objects which generate their own magnetic field will cause a constant shift in the data such as in Figure 1.3b. It should be noted that this will be the case for objects which are located permanently near the magnetometer. Finally, soft-iron bias will manifest itself as scale factor errors in the data such that the locus of data is stretched and rotated as in Figure 1.3c.

Despite all of their flaws, magnetometers still carry large potential in the navigation industry. Because of this, much work has been dedicated to correcting for the various error sources they encounter. The subject of the following chapters will briefly discuss methods of magnetometer error correction, the use of magnetometers in providing attitude information, and finally their implementation into the modified filters this thesis proposes.

1.3 Conclusion

In this chapter, some of the motivation and background information that supports this thesis was discussed. Several current issues with GPS/INS navigation systems were reviewed, and various works were presented which have sought to provide solutions to these issues. A brief review of literature was presented regarding the use of magnetometers for attitude determination, error source estimation and correction in GPS/INS systems, and observability analysis of estimation systems. Following this review, the goals of this thesis, along with its contributions, were presented.

Additionally, some of the fundamental information that supports magnetic sensing technology was discussed. The operation and functionality of magnetometers was defined at a basic level. The process of determining heading as well as 6 DOF attitude was also presented at an introductory level. Finally, a brief discussion of the error sources that corrupt magnetometers was given.

Chapter 2

Magnetometer Calibration Algorithms

In lieu of all of the errors which magnetometers are susceptible to, many have set forth to develop calibration algorithms that will correct for these errors in the data. In [39, 40], Gambhir presented the RESIDG algorithm which calibrates data using a centering approach that minimizes the weighted sum of squares of the residuals between measured and modeled magnetic fields. In [41], a fixed-point algorithm called RESIDT was successfully implemented on the AMPTE spacecraft. RESIDT simultaneously estimated the spin-axis attitude of the spacecraft as well as the inherent magnetometer bias on its sensors. In [34], Davenport developed an algorithm which first computes approximate values for magnetometer bias, and then fine tunes these values via an iterative solution. This approach has been used in the past on the Hubble Telescope. In [43] and [44], Alonso and Shuster presented their TWO-STEP algorithm which improved upon some of the limitations seen by RESIDG, RESIDT, and Davenport's algorithm. In [42], the authors give a survey of these various methods as well as a comparison of them to their own algorithm TWO-STEP.

Finally, in [29] and [30], Gebre-Egziabher developed two different algorithms; both of which work utilizing features of the modeled magnetic field locus to calibrate magnetic field data. In this thesis, several of the methods mentioned here were tested and validated both by simulated and experimental processes. However, the algorithm that was selected for use in the majority of this thesis was the algorithm presented in [29]. In the following section, a brief explanation of this algorithm is given, and then simulated and experimental results are shown which validate it's use for the rest of this thesis.

2.1 Gebre-Egziabher Magnetometer Calibration Algorithm

The algorithm described in [29] is a two-step estimator that begins by forming initial estimates of magnetometer error parameters, and then fine tunes these values using an iterative least-squares process. In particular, the algorithm seeks to estimate bias terms (b_x, b_y, b_z) and scale factor errors (γ_x, γ_y, γ_z) for a three-axis magnetometer. It should be noted that the bias terms estimated account for the effects of stochastic bias and hard-iron bias, while the scale factor terms account for the effects of scale factor errors and soft-iron bias. The algorithm does not, however, account for the effects of misalignment error in any way. Because of this, the use of this algorithm requires accurate alignment of the magnetometer axes with the body axes of a chosen system.

2.1.1 Calculating Initial Estimates

The process of calculating initial estimates of the error parameters begins by establishing an accurate error model for the system as in Equation (2.1)

$$h^2 = \left(\frac{\tilde{h_x^b} - b_x}{\gamma_x}\right)^2 + \left(\frac{\tilde{h_y^b} - b_y}{\gamma_y}\right)^2 + \left(\frac{\tilde{h_z^b} - b_z}{\gamma_z}\right)^2 \tag{2.1}$$

where h is the magnitude of a given earth magnetic field model, and $\tilde{h^b}$ represents magnetic measurements in each respective axis. Expanding Equation (2.1) we arrive at Equation (2.2)

$$h^2 = \frac{(\tilde{h_x^b})^2 - 2(\tilde{h_x^b})(b_x) + (b_x)^2}{\gamma_x^2} + \frac{(\tilde{h_y^b})^2 - 2(\tilde{h_y^b})(b_y) + (b_y)^2}{\gamma_y^2} + \frac{(\tilde{h_z^b})^2 - 2(\tilde{h_z^b})(b_z) + (b_z)^2}{\gamma_z^2} \tag{2.2}$$

Then, recognizing that many of the terms in Equation (2.2) are in fact vectors of data, Equation (2.2) can be re-written as Equation (2.3)

$$-\begin{bmatrix} (\tilde{h_x^b}(t_1))^2 \\ (\tilde{h_x^b}(t_2))^2 \\ (\tilde{h_x^b}(t_3))^2 \\ \vdots \\ (\tilde{h_x^b}(t_{k-1}))^2 \\ (\tilde{h_x^b}(t_k))^2 \end{bmatrix} = \begin{bmatrix} H_{11} H_{12} \end{bmatrix} \begin{bmatrix} b_x \\ \mu_2(b_y) \\ \mu_3(b_z) \\ \mu_2 \\ \mu_3 \\ \mu_4 \end{bmatrix} + \vec{\nu} \tag{2.3}$$

where H_{11} and H_{12} are defined as

$$H_{11} = \begin{bmatrix} -2\tilde{h_x^b}(t_1) & -2\tilde{h_y^b}(t_1) & -2\tilde{h_z^b}(t_1) \\ -2\tilde{h_x^b}(t_2) & -2\tilde{h_y^b}(t_2) & -2\tilde{h_z^b}(t_2) \\ -2\tilde{h_x^b}(t_3) & -2\tilde{h_y^b}(t_3) & -2\tilde{h_z^b}(t_3) \\ \vdots & \vdots & \vdots \\ -2\tilde{h_x^b}(t_{k-1}) & -2\tilde{h_y^b}(t_{k-1}) & -2\tilde{h_z^b}(t_{k-1}) \\ -2\tilde{h_x^b}(t_k) & -2\tilde{h_y^b}(t_k) & -2\tilde{h_z^b}(t_k) \end{bmatrix} \tag{2.4}$$

$$H_{12} = \begin{bmatrix} (\tilde{h_y^b}(t_1))^2 & (\tilde{h_z^b}(t_1))^2 & 1 \\ (\tilde{h_y^b}(t_2))^2 & (\tilde{h_z^b}(t_2))^2 & 1 \\ (\tilde{h_y^b}(t_3))^2 & (\tilde{h_z^b}(t_3))^2 & 1 \\ \vdots & \vdots & \vdots \\ (\tilde{h_y^b}(t_{k-1}))^2 & (\tilde{h_z^b}(t_{k-1}))^2 & 1 \\ (\tilde{h_y^b}(t_k))^2 & (\tilde{h_z^b}(t_k))^2 & 1 \end{bmatrix} \tag{2.5}$$

and the μ terms included in the state vector of the equation are defined as

$$\mu_1 = h^2 \gamma_x^2 \tag{2.6}$$

$$\mu_2 = \frac{\gamma_x^2}{\gamma_y^2} \tag{2.7}$$

$$\mu_3 = \frac{\gamma_x^2}{\gamma_z^2} \tag{2.8}$$

$$\mu_4 = (b_x)^2 + \mu_2(b_y)^2 + \mu_3(b_z)^2 - \mu_1 \tag{2.9}$$

Finally, recognizing that Equation (2.3) is of the common state-space format $\vec{y} = H\vec{x}$, the state vector can be obtained via the least-squares solution in Equation (2.10) and the error parameters subsequently solved for via Equations (2.6)-(2.9).

$$\vec{x} = (H^T H)^{-1} H^T \vec{y} \tag{2.10}$$

This portion of the algorithm only produces initial estimates of the error parameters. However, as noted in [29], often the initial estimates produced are accurate enough for calibration on their own. In this thesis, this was often the case, and as a result there were many situations during data collection and processing where only the first part of the calibration algorithm was utilized. It should also be noted that the state vector shown in Equation (2.3) has been slightly modified from the vector originally published in [29]. The fourth term of the vector has been modified from what was originally a μ_1 to a μ_2. Careful examination of Equation (2.2) will reveal that a matrix representation of Equation (2.2) may be more accurately described when μ_2 replaces μ_1 as the fourth element of the state vector. This change ultimately led to more accurate calibration results in this thesis, which will be presented and discussed later in this chapter.

2.1.2 Iterative Least-Squares Solution

Once a set of initial estimates for the error parameters has been established, more accurate values can be estimated using an iterative least-squares process. This begins by linearizing Equation (2.1) as in Equation (2.11).

$$
\begin{aligned}
-\delta h = &\left(\frac{\tilde{h_x^b} - b_x}{h\gamma_x^2}\right)\delta b_x + \left(\frac{\tilde{h_x^b} - b_x}{\sqrt{h\gamma_x^3}}\right)^2 \delta\gamma_x \\
&+\left(\frac{\tilde{h_y^b} - b_y}{h\gamma_y^2}\right)\delta b_y + \left(\frac{\tilde{h_y^b} - b_y}{\sqrt{h\gamma_y^3}}\right)^2 \delta\gamma_y \\
&+\left(\frac{\tilde{h_z^b} - b_z}{h\gamma_z^2}\right)\delta b_z + \left(\frac{\tilde{h_z^b} - b_z}{\sqrt{h\gamma_z^3}}\right)^2 \delta\gamma_z
\end{aligned} \tag{2.11}
$$

Here, δh is the difference in magnitude between a given magnetometer measurement and an earth magnetic field model such that

$$
\delta h = h - \tilde{h^b}(t_k) \tag{2.12}
$$

Allowing ζ and η to represent the various quantities in Equation (2.11), it may be simplified into

$$
-\delta h = \zeta_x \delta b_x + \eta_x \delta\gamma_x + \zeta_y \delta b_y + \eta_y \delta\gamma_y + \zeta_z \delta b_z + \eta_z \delta\gamma_z \tag{2.13}
$$

and then, recognizing the vector quantities in Equation (2.11), Equation (2.12) is rewritten in matrix form.

$$
\begin{bmatrix} \delta h_1 \\ \delta h_2 \\ \delta h_3 \\ \vdots \\ \delta h_{k-1} \\ \delta h_k \end{bmatrix} = \begin{bmatrix} \zeta_{x_1} & \eta_{x_1} & \zeta_{y_1} & \eta_{y_1} & \zeta_{z_1} & \eta_{z_1} \\ \zeta_{x_2} & \eta_{x_2} & \zeta_{y_2} & \eta_{y_2} & \zeta_{z_2} & \eta_{z_2} \\ \zeta_{x_3} & \eta_{x_3} & \zeta_{y_3} & \eta_{y_3} & \zeta_{z_3} & \eta_{z_3} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \zeta_{x_{k-1}} & \eta_{x_{k-1}} & \zeta_{y_{k-1}} & \eta_{y_{k-1}} & \zeta_{z_{k-1}} & \eta_{z_{k-1}} \\ \zeta_{x_k} & \eta_{x_k} & \zeta_{y_k} & \eta_{y_k} & \zeta_{z_k} & \eta_{z_k} \end{bmatrix} \begin{bmatrix} \delta b_x \\ \delta\gamma_x \\ \delta b_y \\ \delta\gamma_y \\ \delta b_z \\ \delta\gamma_z \end{bmatrix} \tag{2.14}
$$

Finally, recognizing that Equation (2.14) is once again of state-space form, the error-state vector $\delta\vec{x}$ can be computed via a least squares solution, and used to update the most current estimate of the magnetometer error parameters.

$$\delta\vec{x} = (H^T H)^{-1} H^T \delta\vec{h} \tag{2.15}$$

$$b_x(+) = b_x(-) + \delta\vec{x}(1) \tag{2.16}$$

$$\gamma_x(+) = \gamma_x(-) + \delta\vec{x}(2) \tag{2.17}$$

$$b_y(+) = b_y(-) + \delta\vec{x}(3) \tag{2.18}$$

$$\gamma_y(+) = \gamma_y(-) + \delta\vec{x}(4) \tag{2.19}$$

$$b_z(+) = b_z(-) + \delta\vec{x}(5) \tag{2.20}$$

$$\gamma_z(+) = \gamma_z(-) + \delta\vec{x}(6) \tag{2.21}$$

This process is repeated until the algorithm reaches convergence and the error state vector $\delta\vec{x}$ approaches zero.

2.2 Simulated and Experimental Calibration Results

In order to test and validate the chosen calibration algorithm, the algorithm was implemented in MATLAB and tested on both simulated and real magnetometer data. The case of simulated data will be discussed first, and then the real data experiments will be presented. From the real data experiments, the ability of the algorithm will also be shown through heading calculations prior to calibration and post calibration. These results will show that post calibration, the magnetometer can much more accurately estimate the heading of a vehicle.

2.2.1 Simulation Results

It is important to recognize that Gebre-Egziabher's algorithm is, first, a batch least-squares algorithm and then an iterative batch least-squares algorithm. In both parts, the algorithm seeks to fit the entire batch of data to a perfect ellipsoid locus of data centered at

the origin. This ultimately means the algorithm is limited by the set of data which is utilized. The authors of [29] show that having at least 360 degrees of rotation about the vertical axis is essential for use of the algorithm. They also show that variation of this rotation from the planar case will assist the algorithm in identifying the correct error parameters. Without variation from planar rotation, the estimates of the error parameters may be incorrect during initialization, and ultimately cause the overall algorithm to diverge.

In order to simulate magnetometer measurements, a model similar to the model used in [29] was utilized. Each axis of magnetometer measurements was simulated using Equation (2.22)

$$\tilde{h^b} = (h^b + b)\gamma \tag{2.22}$$

where h^b represents the true magnetic field value in each axis. In order to generate true values for the Earth's magnetic field, an International Geomagnetic Reference Field (IGRF) model [50] was used to calculate the true field vector in Auburn, AL. This true field vector was then rotated 360 degrees about the vertical axis, in order to simulate the measurements a magnetometer would collect if it were rotated in a circular manner. In addition to this, the true field vector was also rolled 5 degrees and rotated 360 degrees about the vertical axis, and then pitched 10 degrees and rotated 360 degrees about the vertical axis. This was to ensure variation in the simulated measurements from planar rotation. All of these simulated, perfect measurements were then compiled into one data structure and then corrupted using set bias values (b) and scale-factor errors (γ) as in Equation (2.22). These values are listed in Table 2.1.

The full data structure was then calibrated using the algorithm presented in Section 2.1, and the resulting corrected data was compared with the corrupted data. Figures 2.1 - 2.3 show this comparison. Figure 2.1a displays a perfect set of magnetometer data with a three-axis view, and Figure 2.1b shows the same data set in the XY plane. For simplicity, the data sets pictured in Figures 2.1 - 2.3 are just the data from the planar rotation about the

	Applied Parameters	Estimated Parameters	Percent Error
b_x (mGauss)	100	80.08	19.92%
b_y (mGauss)	20	24.04	20.2%
b_z (mGauss)	200	209.7	4.85%
γ_x	0.8	0.807	0.875%
γ_y	1.2	1.210	0.833%
γ_z	1.1	1.121	1.9%

Table 2.1: Some Table of data

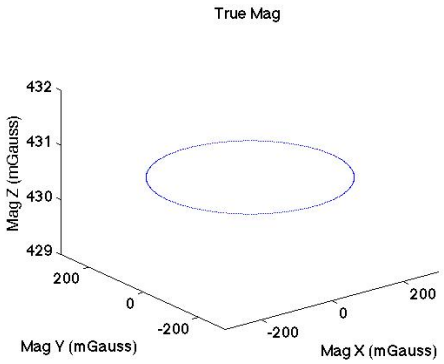

(a) Set of perfect magnetometer measurements over a 360 degree rotation about the vertical axis

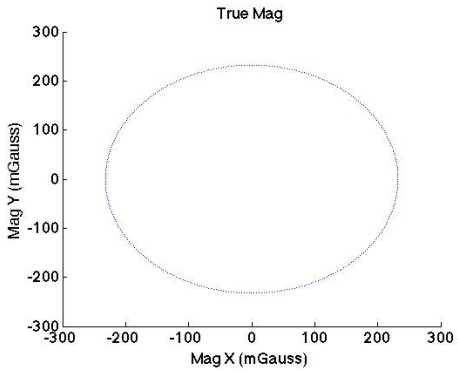

(b) Set of perfect magnetometer measurements viewed in the XY plane

Figure 2.1: Simulated perfect magnetometer measurements

vertical axis, although data from all three simulated rotations was used for the calibration. The uniformity of the data in the absence of noise should be noted in 2.1a, as well as it's centering at the origin in 2.1b. Figure 2.2 shows the data after corruption. As can be seen the uniformity of the data is now disrupted by Gaussian white noise. In addition to this, Figure 2.2b shows bias has shifted the data from the origin, and scale factor errors have stretched the y-axis measurements and shrunk the x-axis measurements. In contrast, viewing the calibrated data from Figure 2.3, it can be seen that while the data is still corrupted by noise, its center has now been shifted much closer to the origin and its axes scaled back to their proper values. The exact estimates of the error parameters, along with a percent error analysis, is given in Table 2.1.

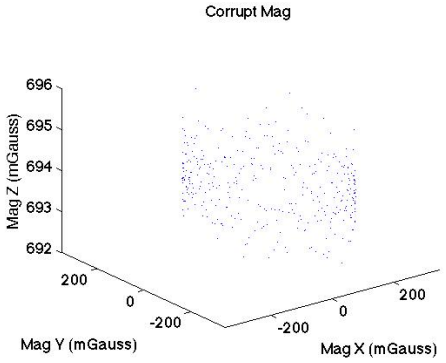

(a) Set of corrupt magnetometer measurements over a 360 degree rotation about the vertical axis

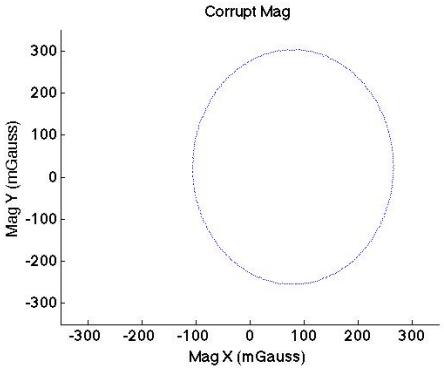

(b) Set of corrupt magnetometer measurements viewed in the XY plane

Figure 2.2: Simulated corrupted magnetometer measurements

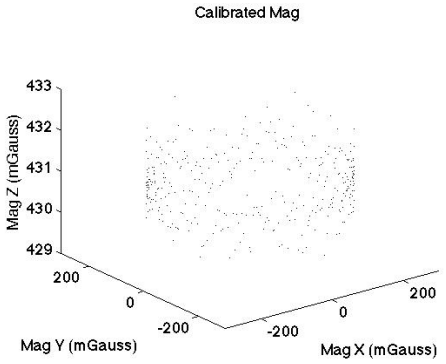

(a) Set of calibrated magnetometer measurements over a 360 degree rotation about the vertical axis

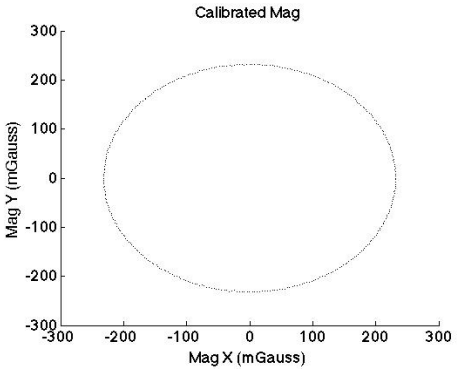

(b) Set of calibrated magnetometer measurements viewed in the XY plane

Figure 2.3: Simulated magnetometer measurements after calibration

2.2.2 Experimental Results

After successful validation of the algorithm in simulation, the algorithm was tested via experimentation. An XSens MTi inertial measurement unit (IMU) mounted inside of a modified Infiniti G35 vehicle was used for data collection. The vehicle was outfitted with a three antenna Septentrio GPS system for attitude truth validation. The vehicle was then driven such that the magnetometer experienced several full rotations about the vertical axis. The corrupt magnetometer data was then calibrated in post processing using the algorithm presented in Section 2.1. Figure 2.4 displays three-axis and XY plane views of the data before calibration. It can be seen that in the x-axis, the data suffers greatly from bias which has shifted the locus away from the origin. A closer view will also reveal a small amount of bias in the y-axis, and some scaling errors which stretch the ellipse along the x-axis. In contrast, Figure 2.5 shows that the calibration algorithm has accounted for the bias which originally shifted the data from the origin. In addition to this, the effects of scale-factor errors in the XY plane have also been removed.

One error which the algorithm suffers from during this data run is its inability to estimate bias in the z-axis. As can be seen, the post-calibration magnetic data has been shifted very close to zero along the z-axis, when in fact the true magnitude of the field in the z-axis is much closer to what was measured before calibration (around -0.4 Gauss). This discrepancy in the calibrated data occurs because the collected data provides no variation from rotation in the horizontal plane. Therefore, because the collected data does not vary from the horizontal plane, the algorithm has no variation in z-axis measurements and ultimately seeks to shift the locus of data closer to the origin along the z-axis as well as the y and x axes. This presents a significant problem for ground vehicles since variation from rotation about the horizontal plane requires dynamics which are often difficult to accomplish in ground vehicles. As a result, for ground vehicle navigation, z-axis bias which effects the magnetometer measurements should be closely monitored. These findings will play a much more important role in the attitude determination work that will be described later in this thesis.

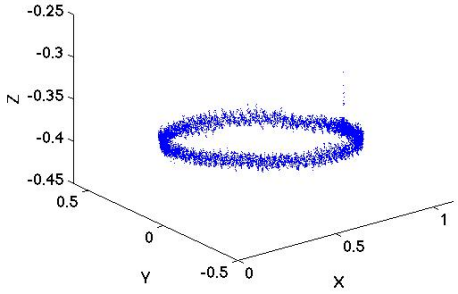

(a) Set of raw magnetometer data

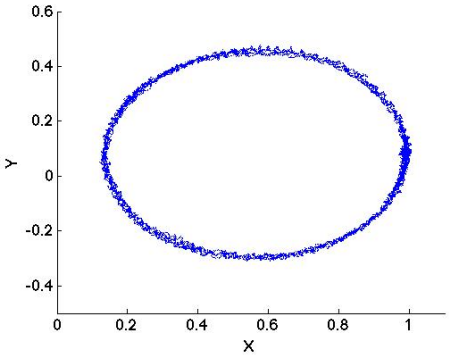

(b) Set of raw magnetometer measurements viewed in the XY plane

Figure 2.4: Raw magnetometer data collected on the modified Infiniti G35

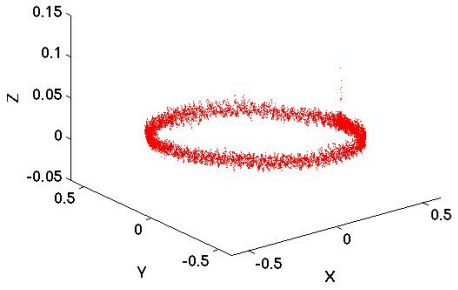

(a) Set of post process calibrated magnetometer measurements.

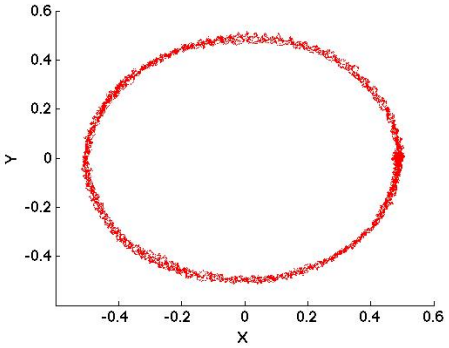

(b) Set of calibrated magnetometer measurements viewed in the XY plane

Figure 2.5: Post process calibrated data from tests run on the modified Infiniti G35 ground vehicle

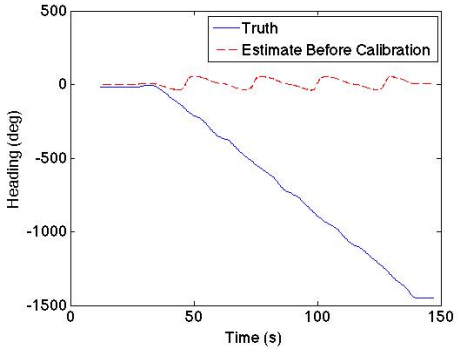

(a) Ground vehicle heading calculated from raw magnetometer measurements vs. the true ground vehicle heading

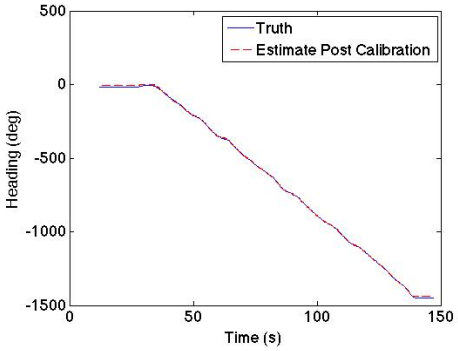

(b) Ground vehicle heading calculated from calibrated magnetometer measurements vs. the true ground vehicle heading

Figure 2.6: Ground vehicle heading calculations

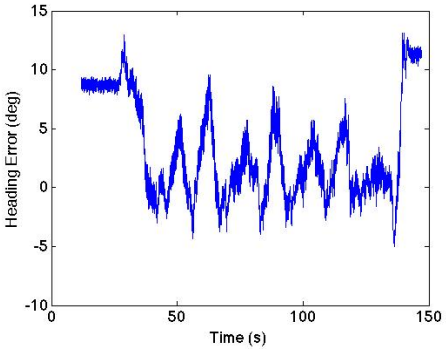

(a) Heading error from calibrated magnetometer calculations

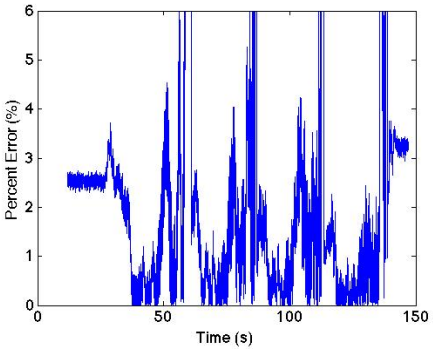

(b) Percent error from calibrated magnetometer calculations

Figure 2.7: Error associated with ground vehicle heading calculations

Fortunately, for ground vehicle heading calculation, if the tilt angles experienced by the vehicle are small, any errors in the z-axis will not effect the final calculation. This allows for a comparison of heading calculation between the raw measurements and calibrated measurements, which is shown in Figure 2.6. Figure 2.6a displays the heading calculated from raw measurements by the magnetometer. As can be seen, the measurement errors greatly corrupt the heading calculation. Referencing Figure 2.4b, it is evident that the bias experienced by the magnetometer in the x-axis keeps it from identifying how the vehicle is turning. The fact that the corrupt x-axis measurements never reach zero, or go below zero, confines the heading calculation to a right-half plane of possible yaw angles. In sharp contrast, it can be seen in Figure 2.6b that the calibrated measurements are much more capable of calculating the true heading of the vehicle. Figure 2.7 depicts the error associated with the heading calculated by the calibrated magnetometer. It is clear from Figure 2.7a that the calibration provides higher accuracy in its heading estimate, however is still subject to several spikes in error. Preliminary analysis of the data from this run has shown that the spikes in heading error are correlated with the velocity of the vehicle. This may mean that the magnetometer is experiencing magnetic interference due to some part of the ground vehicle itself during the run. For example, it has been shown in the past that steel-belted radial tires are capable of producing their own magnetic fields during rotation. [49] This theory leads to interesting implications for the calibration of magnetometers on ground vehicles. However, further analysis on this subject is not pursued in this thesis, and as a result is left to future work.

2.3 Conclusion

There has been a wealth of research to develop faster and more capable methods of magnetometer calibration. For this thesis, a specific algorithm was selected from among the many capable and available algorithms. The algorithm is essentially a two-part least squares algorithm which seeks to fit a batch of data to an ideal ellipse locus of data centered at the

origin. Simulated magnetometer data was corrupted and then calibrated using a MATLAB implementation of the chosen algorithm. Results from these simulated data experiments showed that the algorithm is capable of estimating the bias and scale factor errors experienced by the corrupt data.

Finally, experimental data was collected from an XSens MTi IMU, which was also equipped with a built in 3-axis magnetometer. This sensor was mounted on a modified Infiniti G35 ground vehicle, and the vehicle was driven in a circular trajectory on level ground. The collected data was then calibrated using the chosen algorithm, and results showed that the algorithm was capable of correcting for bias and scale factor errors in the x and y axes. The calibrated data showed errors in the determination of any z-axis bias, and as a result it was noted that, for ground vehicle magnetometer calibration, the effects of calibration in the z-axis should be kept in check in the event that the vehicle only experiences rotation in the horizontal plane. In addition to observing correct bias estimation, it was also observed that the corrected measurements were much more capable of estimating the heading of the ground vehicle than the corrupt measurements.

Chapter 3

Attitude Determination via Magnetic Vector Observations

While magnetometers have frequently been used for heading calculation and stabilization on ground vehicles, they are also very useful for performing attitude determination on many types of vehicles. A brief discussion on the use of vector measurements for attitude determination was given in Section 1.2.2. In this section, several algorithms that perform the calculation described in Section 1.2.2 will be discussed, and the specific algorithm that was used in this thesis will be presented. After a thorough explanation of the algorithm has been given, magnetometer measurements, inertial measurements, and velocity vector measurements will all be applied to estimate the attitude of an aerial trajectory example. The advantages and disadvantages of these various measurements will all be explored in their application with the chosen algorithm, and results will show the capability of using vector measurements for attitude determination on the aerial platform.

3.1 Background Information and Previous Work

The problem of determining the attitude of a system given vector measurements was first proposed by Grace Wahba in 1965 [32]. As described in Section 1.2.2, the problem entails determining the rotation between two vector sets. Each vector set must contain at least two vectors, but may have more than two. These vector sets will reside in two different coordinate frames, and each vector in a set will have a corresponding vector in the opposing coordinate frame that points to the same point as itself. This essentially simplifies to the rotation between the two vector sets ultimately being the same rotation as that between the two coordinate frames they reside in.

Wahba first solved this problem when he considered it in 1965, but since then many have developed faster and more efficient methods for solving the same problem. Davenport formed his own q-method for solving the problem, which seeks a quaternion solution to the problem by determining the largest eigenvector of a 4x4 symmetric matrix [33]. Since Davenport, the QUaternion ESTimator (QUEST) and EStimators of the Optimal Quaternion (ESOQ and ESOQ2) algorithms have also been developed [35, 36]. These algorithms provide much faster solutions to Wahba's problem; however, as Markley and Mortari pointed out, they are ultimately less robust methods than Davenport's q-method [37].

Finally, in [14, 15], a solution to Wahba's problem is presented that estimates a rotation quaternion through an iterative Gauss-Newton method. This method was chosen for use in this thesis, and a brief discussion of it is presented in the following section.

3.2 Attitude Determination Algorithm

Given two vectors $\underline{y_b^1}$ and $\underline{y_b^2}$ in a given vehicle body coordinate frame and two corresponding vectors $\underline{y_e^1}$ and $\underline{y_e^2}$ in an Earth-Centered-Earth-Fixed (ECEF) frame, it is possible to estimate the rotation matrix R_b^e such that

$$\underline{y_e^1} = R_b^e \underline{y_b^1} \tag{3.1}$$

$$\underline{y_e^2} = R_b^e \underline{y_b^2} \tag{3.2}$$

Specifically, vector measurement devices may be used to estimate R_b^e such that $\underline{y_b^1}$ and $\underline{y_b^2}$ might be measurements from a magnetometer, star sensor, sun sensor, or any other kind of vector measurement device, and $\underline{y_e^1}$ and $\underline{y_e^2}$ would be corresponding measurements or model based calculations in the ECEF frame. For this thesis, two different sets of measurement devices were used to perform attitude determination calculations. First, magnetic measurements were paired with acceleration measurements such that the magnetic measurements

collected observations of the Earth's magnetic field and the acceleration measurements collected observations of the Earth's gravity vector. In the second set of measurements, magnetic measurements were paired with vehicle velocity measurements such that the magnetic measurements collected observations of the Earth's magnetic field and the velocity measurements collected observations of the vehicle's velocity vector. Both of these approaches have varying advantages and disadvantages which will be explored going forward.

Having defined the vector sets that will be used to calculate an attitude solution, the process begins by defining an error function H that the algorithm will seek to minimize.

$$H = \epsilon^T \epsilon = (\underline{y_e} - M\underline{y_b})^T (\underline{y_e} - M\underline{y_b}) \tag{3.3}$$

Letting magnetic measurements be indicated by $\underline{m}$, and accelerometer/gravity measurements be indicated by $\underline{\alpha}$, $\underline{y_b}$ and $\underline{y_e}$ can be defined as

$$\underline{y_b} = \begin{bmatrix} \underline{m_b} & \underline{\alpha_b} \end{bmatrix}^T \tag{3.4}$$

$$\underline{y_e} = \begin{bmatrix} \underline{m_e} & \underline{\alpha_e} \end{bmatrix}^T \tag{3.5}$$

and M may be defined as

$$M = \begin{bmatrix} R_b^e & 0 \\ 0 & R_b^e \end{bmatrix} \tag{3.6}$$

When velocity measurements $\underline{v}$ are used rather than accelerometer/gravity measurements $\underline{\alpha}$, $\underline{\alpha_b}$ and $\underline{\alpha_b}$ are simply replaced by $\underline{v_b}$ and $\underline{v_e}$ in Equations (3.4) and (3.5).

The algorithm will seek to minimize H by finding the quaternion which best characterizes the rotation between the body frame and the ECEF frame. In order for this to happen, R_b^e and ultimately M must be rewritten in a quaternion format. Defining a given quaternion n such that

$$n = a\hat{i} + b\hat{j} + c\hat{k} + d \tag{3.7}$$

a given rotation matrix R can be rewritten as

$$R = \begin{bmatrix} d^2 + a^2 - b^2 - c^2 & 2(ab - cd) & 2(ac + bd) \\ 2(ab + cd) & d^2 + b^2 - a^2 - c^2 & 2(bc - ad) \\ 2(ac - bd) & 2(bc + ad) & d^2 + c^2 - b^2 - a^2 \end{bmatrix} \tag{3.8}$$

Rewriting R_b^e in terms of quaternion components, allows for a Gauss-Newton iterative solution to be formed in terms of n and the jacobian matrix J such that,

$$\hat{n}_{k+1} = \hat{n}_k - [J^T(\hat{n}_k)J(\hat{n}_k)]^{-1}J^T(\hat{n}_k)\epsilon(\hat{n}_k) \tag{3.9}$$

$$J = \begin{bmatrix} \frac{\delta M}{\delta a}\underline{y_b} & \frac{\delta M}{\delta b}\underline{y_b} & \frac{\delta M}{\delta c}\underline{y_b} & \frac{\delta M}{\delta d}\underline{y_b} \end{bmatrix} \tag{3.10}$$

$$\epsilon(\hat{n}_k) = \underline{y_e} - M(\hat{n}_k)\underline{y_b} \tag{3.11}$$

Iterating Equation (3.9) until convergence will give the quaternion rotation which minimizes the error function H in Equation (3.3). Once this quaternion has been determined, it can be converted back to a rotation matrix using Equation (3.8) and then converted to euler angles from there if desired.

3.3 Attitude Determination for an Aerial Trajectory

In order to test the attitude determination algorithm presented in Section 3.2, an aerial trajectory was simulated in MATLAB. A representation of the trajectory in the local North-East-Down (NED) frame is shown in Figure 3.1. To simulate accelerometer measurements

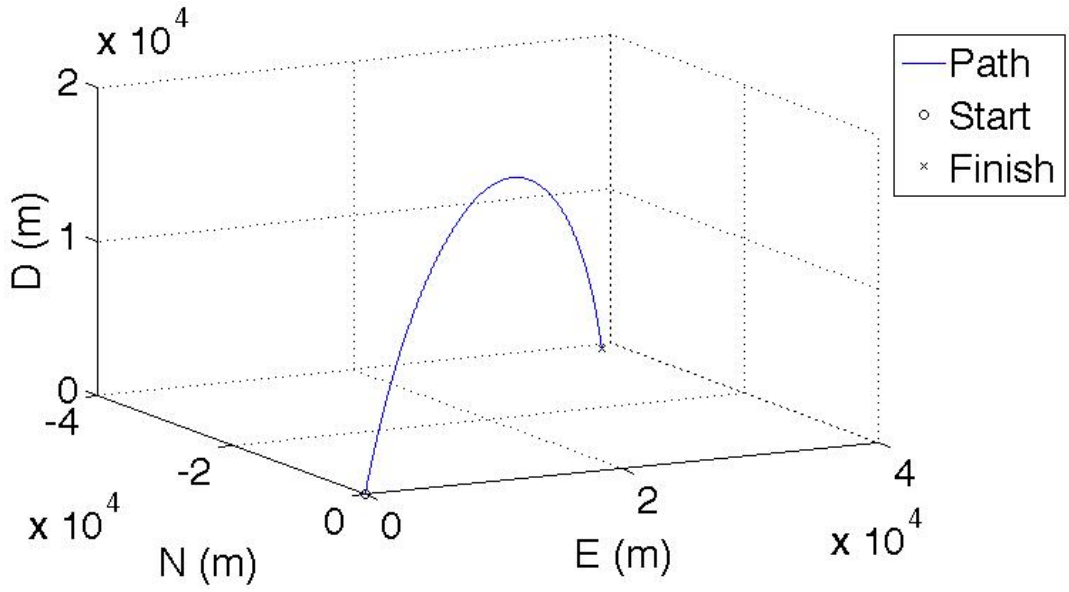

Figure 3.1: Simulated aerial trajectory

for the aerial vehicle, perfect measurements were defined as the combination of accelerations due to gravity $\underline{g_b}$ and accelerations due to vehicle trajectory $\underline{\ddot{x}_b}$ such that

$$\underline{\alpha_b} = \underline{g_b} + \underline{\ddot{x}_b} \tag{3.12}$$

Here $\underline{\ddot{x}_b}$ was defined using standard 6 DoF dynamic equations for particle motion such that

$$\underline{\ddot{x}_b} = \underline{\ddot{x}_{lin}} + \underline{\dot{\omega}} \times \underline{x_b} + \underline{\omega} \times (\underline{\omega} \times \underline{x_b}) + 2\underline{\omega} \times \underline{\dot{x}_b} \tag{3.13}$$

where $\underline{\dot{x}_b}$ and $\underline{x_b}$ are the velocity and position of the vehicle defined in the body frame coordinate system, $\underline{\dot{\omega}}$ and $\underline{\omega}$ are the angular acceleration and angular velocity the vehicle experiences, and $\underline{\ddot{x}_{lin}}$ is the linear acceleration that the vehicle experiences.

Magnetometer measurements and velocity measurements were also simulated in both the vehicle body frame and the global frame, and extracted for testing. Contrary to the formulation shown in Section 3.2, a local NED frame was used as a global frame rather than an ECEF frame. Other than this change, the same equations and formulation described

in Section 3.2 were used to generate attitude estimates at each time step. The process of utilizing the formulation was as follows:

1. The $\underline{y_b}$ and $\underline{y_e}$ vectors were defined by substituting the simulated magnetometer and accelerometer/velocity measurements into Equations (3.4) and (3.5).

2. The $\hat{n}_k$ quaternion was initialized with arbitrary guesses.

3. The M matrix and jacobian matrix J were defined using Equations (3.6), (3.8), and (3.10).

4. The error function ϵ was evaluated using Equation (3.11)

5. A new quaternion estimate $\hat{n}_{k+1}$ was calculated using Equation (3.9)

6. Steps 3-5 were repeated until convergence of the quaternion estimate $\hat{n}$

7. Equation (3.8) was used to convert $\hat{n}$ to a rotation matrix, which was then transformed into euler angles roll ϕ, pitch θ, and yaw ψ.

Figure 3.2 displays the results from the algorithm using simulated accelerometer measurements and magnetometer measurements. As can be seen, the performance of the algorithm with this pair of inputs does not provide a high amount of accuracy. While the roll angle is estimated reasonably well, the yaw estimate suffers from considerable errors during the first thirty seconds of the simulation, and the pitch angle estimate never truly converges to the true pitch of the aerial vehicle. The reason for the poor performance of the algorithm is found in the accelerometer inputs used. The purpose of using accelerometer inputs in the algorithm is to supply an accurate measurement of the gravity vector as seen by the vehicle body. However, as can be seen in Figure 3.3, the accelerometer will measure all of the accelerations experienced by the body. As a result, if the body experiences a large amount of acceleration in its trajectory, the measured gravity vector will be buried under many other

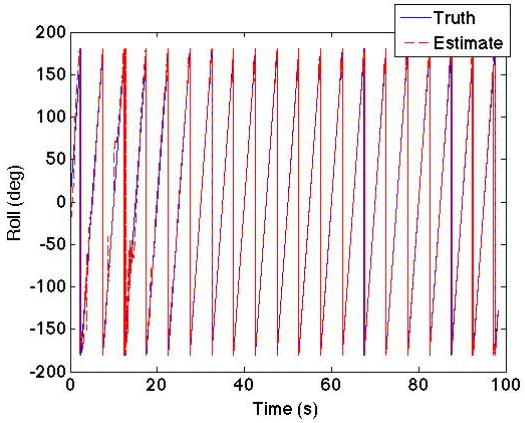

(a) Roll estimate from attitude determination algorithm vs. true roll angle

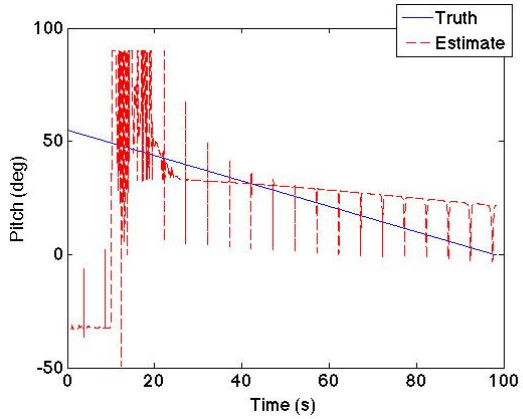

(b) Pitch estimate from attitude determination algorithm vs. true roll angle

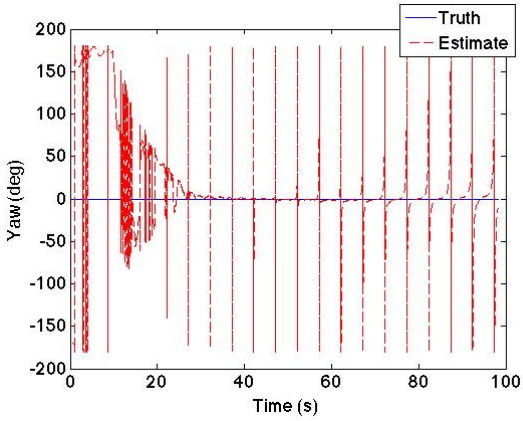

(c) Yaw estimate from attitude determination algorithm vs. true roll angle

Figure 3.2: Attitude estimate from the attitude determination algorithm vs. true attitude using magnetometer and accelerometer measurements

accelerations which corrupt the algorithm solution. Specifically, the simulated aerial trajectory experiences a large amount of initial thrust, which explains the erratic behavior that is initially seen in the pitch and yaw estimates of the system. As the simulation continues, accelerations on the body due to rotation continue to distort measurement of the gravity vector, further corrupting the estimate. One possible solution to this problem may be to use an accurate dynamic model of the vehicle system in order to calculate the non-gravitational accelerations experienced by the body. In doing so, it may be possible to separate the gravity vector measured by the accelerometer from the overall measurement. However, this is a study that is left for future work and will not be explored in this thesis.

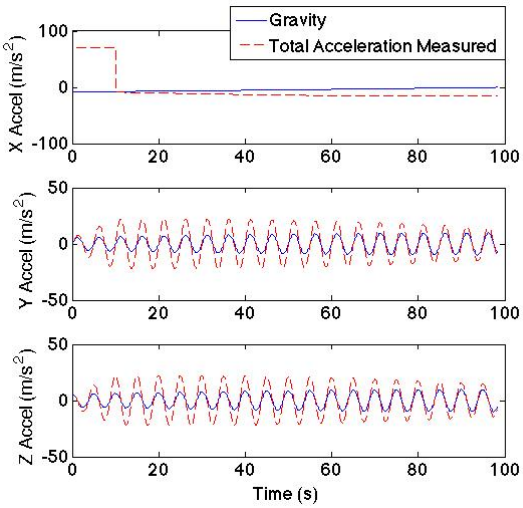

Figure 3.3: True gravity experienced by the aerial vehcile vs. total acceleration measured by the accelerometer

In sharp contrast, pairing velocity measurements with magnetometer measurements proved to be a much more reliable method of attitude determination for the aerial simulation. As seen in Figure 3.4, the use of velocity measurements provides much more accurate results. The attitude estimate follows the true attitude much more closely than it did using accelerometer measurements. There are several periodic spikes throughout the attitude estimate which are produced from brief moments where the algorithm is unable to identify a specific rotation between the vector sets. In other words, there are periodic moments as the aerial vehicle rotates where there is more than one rotation which will solve the attitude determination problem. Errors such as these are dangerous in an estimation scheme because they increase the probability of a gimbal locked solution. In order to explain gimbal lock, it is useful to picture a gyroscope with three gimbals (one for each rotation angle), as in Figure 3.5. Looking at the image, a gimbal locked solution is one in which the rotation angles of the solution cause two or more of the gimbals to become parallel. When this occurs, it causes dimesnions of the solution to be "lost". This problem may be prevented by using more than two vectors in each set, in order to avoid a trivial solution in the algorithm. However, the use of more vector sets ultimately depends on the sensor suite available with any given vehicle.

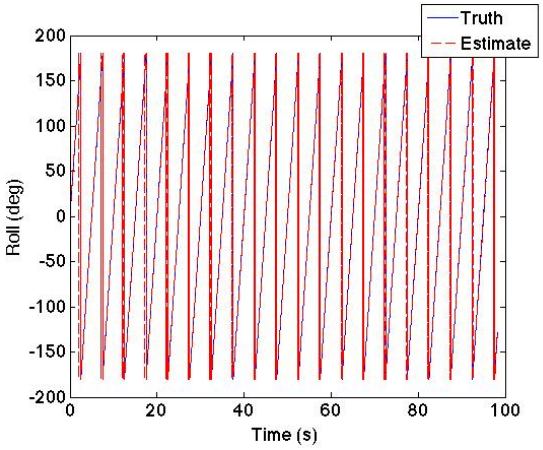

(a) Roll estimate from attitude determination algorithm vs. true roll angle

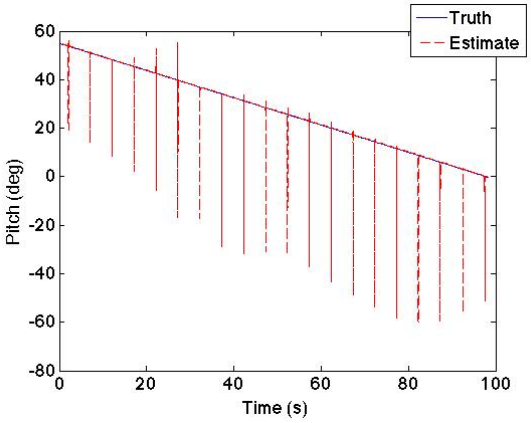

(b) Pitch estimate from attitude determination algorithm vs. true roll angle

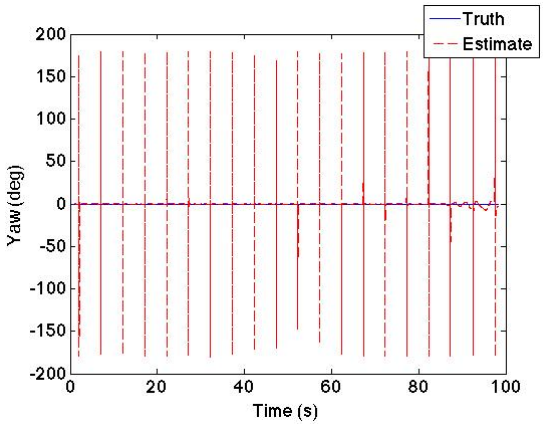

(c) Yaw estimate from attitude determination algorithm vs. true roll angle

Figure 3.4: Attitude estimate from the attitude determination algorithm vs. true attitude using magnetometer and velocity vector measurements

(a) Normal operation (b) Gimbal lock

Figure 3.5: Gyroscope with three gimbals

Another solution is to utilize the attitude determination algorithm in parallel with another estimation scheme such as a Kalman filter. The idea of uniting this algorithm with another estimator has been explored before [14] and will be explored later in this thesis. It should also be noted that in order to produce the results shown in Figure 3.4, very high update-rate sensors were simulated. The reason for this is due to the high dynamics of the aerial vehicle. Specifically, the high roll rate of the vehicle will produce high acceleration, velocity, angular, and magnetic rates of change. As a result, if high update-rate sensors are not used, missing information will not only cause integration errors in most navigation algorithms, but ultimately cause divergence of the estimates from the attitude determination algorithm chosen for this thesis.

3.4 Conclusion

There has been a significant amount of previous work done in the area of attitude determination from vector observations. In this thesis it was chosen to use the attitude determination algorithm presented in [14]. In this chapter, an exposition of that algorithm was given, followed by a presentation of simulation results that was meant to validate the

algorithm. In order to validate the use of this algorithm with the proposed methodologies, an aerial simulation was created which simulated accelerometer, magnetometer, and velocity measurements in both a vehicle body frame and global frame. These measurements were then implemented on the chosen attitude determination algorithm and the results were presented.

It was also shown that utilizing accelerometer and magnetometer measurements with the algorithm produces high amounts of estimate error for trajectories with high dynamics. In contrast, using velocity measurements rather than accelerometer measurements produced much lower amounts of estimate error. As a result, it was shown that the algorithm is capable of performing attitude determination in a high dynamic scenario given velocity and magnetometer measurements. It is also hypothesized that in scenarios of low dynamics, it may be possible to use accelerometer measurements rather than velocity measurements.

Chapter 4

GPS/INS Navigation Algorithms

While magnetometers prove to be useful, stand-alone sensors for heading calculation and attitude determination, it is often desirable to incorporate various useful sensors into one optimal navigation scheme. Because of that, this thesis has also sought to improve the performance of standard navigation algorithms with modified approaches to those algorithms that incorporate magnetometers.

Among the various navigation algorithms available, one of the most commonly known schemes is the Extended Kalman Filter (EKF). The EKF is essentially a navigation algorithm which unites the benefits from several different measurement sources to produce an optimal position, velocity, and attitude (PVA) solution. EKFs may be very broadly separated into model-based and navigation-based EKFs. Model-based EKFs utilize measurements of vehicle parameters in order to supply predictions of the PVA solution. Navigation based EKFs utilize inertial navigation system (INS) sensors such as accelerometers and gyroscopes along with standard INS mechanization equations in order to produce a PVA solution. Either type of EKF will commonly use GPS measurements to improve the predictions supplied by either a vehicle model or INS solution. It is important to differentiate the EKF from its precursor the Kalman Filter (KF). The KF is another commonly known estimation scheme, however it is not ideal for vehicle navigation because it assumes linearity of its system equations. An EKF is a KF that linearizes a nonlinear system such that the KFs assumption of linear equations is met. In this thesis, it was chosen to begin with a standard navigation-based EKF, and see if modifications could be made using magnetometers to improve the performance of the EKF. For this reason, the rest of this chapter will focus on navigation-based EKFs and the modified EKFs proposed in this thesis.

4.1 Standard Extended Kalman Filter Algorithms

The advantage of using the standard EKF is found in its ability to utilize the benefits of both INS sensors and GPS. INS sensors are often high update-rate devices that suffer from several different common error sources. Conversely, GPS has the ability to provide high fidelity PVA measurements at a much slower update-rate. The EKF therefore, works to provide high update-rate, high accuracy PVA solutions, often for use in real time systems.

Before an exposition of the standard EKF architecture can be given, a basic understanding of the error sources GPS/INS systems suffer from must be grasped. INS sensors will often suffer from typical sensor errors such as stochastic wide-band noise, as well as sensor bias and drift. For most cases, the noise experienced by these sensors can be modeled as white Gaussian noise, while the sensor bias and drift must be modeled by a random walk. Several different random walk models are usable, but for this thesis, sensor bias and drift were modeled as a first order Markov process such that

$$\dot{b} = -\frac{1}{\tau}b + \mu \tag{4.1}$$

where τ is the time constant of the process and μ is a driving white-noise term.

GPS PVA measurements themselves can be modeled as being corrupted by white Gaussian noise. However, GPS pseudorange measurements will suffer from the effects of satellite clock bias and reciever clock bias.

$$\tilde{\rho} = \rho + b_{sat} + b_{rcvr} \tag{4.2}$$

GPS receivers will receive information from GPS satellites that will allow them to compensate for satellite clock bias. However, the receiver clock bias and drift must still be accounted for whenever pseudorange and pseudorange rate measurements are used. The receiver clock bias and drift may be modeled linearly such that

$$b_{clk} = \dot{b}_{clk} t_k + w_{clk} \tag{4.3}$$

$$\ddot{b}_{clk} = 0 \tag{4.4}$$

4.1.1 Loosely-Coupled Algorithm

The first navigation filter that will be discussed is the loosely-coupled integration of the EKF. The loosely-coupled EKF is capable of providing estimates of position $\underline{P}$, velocity $\underline{V}$, attitude $\underline{A}$, accelerometer bias $\underline{b_\alpha}$, and gyroscope bias $\underline{b_\omega}$. These parameters together form the state vector $\hat{X}$ of the system. The reason why the filter is called a "loosely-coupled" EKF is because the filter utilizes GPS position $\underline{P_{GPS}}$ and velocity $\underline{V_{GPS}}$ measurements to provide correction to the filter state estimates. This occurs in the portion of the filter which is typically referred to as the measurement update. A further discussion of the filter's measurement update will be given later. The algorithm begins by propagating the state vector forward in time using the system inputs (accelerometer and gyroscope measurements $\underline{\alpha}$ and $\underline{\omega}$).

$$\dot{\underline{\hat{P}}} = \underline{\hat{V}} \tag{4.5}$$

$$\dot{\underline{\hat{V}}} = R_b^e(\underline{\alpha} - \underline{b_\alpha}) - \underline{g} \tag{4.6}$$

$$\dot{\underline{\hat{A}}} = R_b^e(\underline{\omega} - \underline{b_\omega}) \tag{4.7}$$

Here $\underline{g}$ represents the local gravity vector in the ECEF frame. Also, Equation (4.1) is used to propagate estimates of $\underline{b_\alpha}$ and $\underline{b_\omega}$. In addition to propagating the state vector $\hat{X}$ it is also necessary to propagate a state estimate error covariance matrix P such that

$$\dot{P} = FP + PF^T + B_w Q B_w^T \tag{4.8}$$

where Q is the process noise covariance matrix and is defined by the variance of the system process noise sources.

$$Q = \begin{bmatrix} \sigma_f^2 & O_{33} & O_{33} & O_{33} \\ O_{33} & \sigma_\omega^2 & O_{33} & O_{33} \\ O_{33} & O_{33} & \sigma_{bf}^2 & O_{33} \\ O_{33} & O_{33} & O_{33} & \sigma_{b\omega}^2 \end{bmatrix} \tag{4.9}$$

Here each σ^2 term represents a 3x3 diagonal matrix of variances corresponding to each process noise source, and O_{33} indicates a 3x3 matrix of zeros.

B_w is a matrix which introduces the system noise sources into the system model. It is defined as

$$B_w = \begin{bmatrix} O_{33} & O_{33} & O_{33} & O_{33} \\ R_b^e & O_{33} & O_{33} & O_{33} \\ O_{33} & R_b^e & O_{33} & O_{33} \\ O_{33} & O_{33} & R_b^e & O_{33} \\ O_{33} & O_{33} & O_{33} & R_b^e \end{bmatrix} \tag{4.10}$$

Finally, F is the state-transition matrix which is defined by a linearization of Equations (4.1) and (4.5) to (4.7) around the current state estimate such that

$$F = \begin{bmatrix} O_{33} & I_{33} & O_{33} & O_{33} & O_{33} \\ O_{33} & O_{33} & [(R_b^e \underline{\alpha})\wedge] & -R_b^e & O_{33} \\ O_{33} & O_{33} & O_{33} & O_{33} & -R_b^e \\ O_{33} & O_{33} & O_{33} & -\frac{1}{\tau_f} I_{33} & O_{33} \\ O_{33} & O_{33} & O_{33} & O_{33} & -\frac{1}{\tau_\omega} I_{33} \end{bmatrix} \tag{4.11}$$

where I_{33} represents a 3x3 identity matrix.

Propagation of the state $\hat{X}$ and the state estimate error covariance matrix P makes up what is commonly called the time-update portion of the EKF. Many methods of integration are available for propagation; for this thesis, Equations (4.12)-(4.16) (which are based on the equations given in [1]) were used for propagation.

$$R^e_{bk} = R^e_{bk-1}(I_{33} + \Omega^b_{ib}\Delta t) \tag{4.12}$$

$$\Omega^b_{ib} = [(\underline{\omega}_{k-1} - \underline{b_{\omega}}_{k-1})\wedge] \tag{4.13}$$

$$\underline{f^e_{ib}} = \frac{1}{2}(R^e_{bk} + R^e_{bk-1})(\underline{\alpha}_{k-1} - \underline{b_{\alpha}}_{k-1}) \tag{4.14}$$

$$\underline{V}_k = \underline{V}_{k-1} + (\underline{f^e_{ib}} + \underline{g_e})\Delta t \tag{4.15}$$

$$\underline{P}_k = \underline{P}_{k-1} + (\underline{V}_k + \underline{V}_{k-1})\frac{\Delta t}{2} \tag{4.16}$$

where $\underline{g_e}$ in Equation (4.15) is the local gravity vector in the ECEF frame. Finally, the P matrix may also be propagated using Equation (4.17)

$$P_k = FP_{k-1}F^T + B_wQB^T_w \tag{4.17}$$

Following the time update is a process called the measurement update, which was briefly mentioned before. The measurement update begins by defining the measurements that will be used to correct the state estimates. As previously stated, the loosely-coupled filter uses GPS position measurements $\underline{P_{GPS}}$ and velocity measurements $\underline{V_{GPS}}$ to correct the state estimates. In this way, the measurement vector $\underline{z}$ is defined as

$$\underline{z} = \begin{bmatrix} \underline{P_{GPS}} & \underline{V_{GPS}} \end{bmatrix}^T \tag{4.18}$$

The measurements must be correlated to the state vector by a measurement matrix H such that

$$z = H\hat{X} \tag{4.19}$$

Examining $\hat{X}$ and z reveals that H is

$$H = \begin{bmatrix} I_{33} & O_{33} & O_{33} & O_{33} & O_{33} \\ O_{33} & I_{33} & O_{33} & O_{33} & O_{33} \end{bmatrix} \tag{4.20}$$

Having defined a model for measurement updates, the filter gain or Kalman gain L can be defined by Equation (4.21)

$$L = PH^T(HPH^T + R)^{-1} \tag{4.21}$$

where R is a measurement noise covariance matrix such that

$$R = \begin{bmatrix} \sigma_P^2 & O_{33} \\ O_{33} & \sigma_V^2 \end{bmatrix} \tag{4.22}$$

Once the Kalman gain for the measurement epoch has been calculated, the state $\hat{X}$ and state estimate error covariance matrix P can be corrected such that

$$\hat{X}_k = \hat{X}_k + L(z - H\hat{X}_k) \tag{4.23}$$

$$P_k = (I_{33} - LH)P_k \tag{4.24}$$

4.1.2 Closely-Coupled Algorithm

While the loosely-coupled Kalman filter architecture has been widely used, many have also chosen to use a closely-coupled Kalman filter architecture instead [45]- [47]. The distinction between the loosely-coupled architecture and the closely-coupled architecture is found in their different uses of GPS measurements. While the loosely-coupled algorithm uses GPS

position $\underline{P_{GPS}}$ and velocity measurements $\underline{V_{GPS}}$, the closely-coupled algorithm goes a step deeper and uses the actual pseudorange $\underline{\rho}$ and pseudorange rate measurements $\underline{\dot{\rho}}$ that a GPS receiver will produce.

The closely-coupled approach is advantageous because during GPS outages or GPS degradation, the closely-coupled algorithm will outperform the loosely-coupled algorithm. This is because, in situations where less than four satellites are in view, a GPS receiver can no longer compute a PVA solution. However, the receiver may continue to generate pseudorange and pseudorange rate measurements to whatever satellites are in view. For this reason, in situations where less than four satellites are in view, the closely-coupled algorithm is still able to provide some aiding during its measurement update.

In order to use pseudorange and pseudorange rate measurements as measurement updates, several changes must be made to the loosely-coupled architecture in both the time update and measurement update. The first change that must be made is the addition of receiver clock bias and drift terms to the state vector of the system such that

$$\hat{X} = \begin{bmatrix} \underline{P} & \underline{V} & \underline{A} & \underline{b_\alpha} & \underline{b_\omega} & b_{clk} & \dot{b}_{clk} \end{bmatrix}^T \tag{4.25}$$

The new terms b_{clk} and $\dot{b}_{clk}$ may be propagated forward using Equations (4.3) and (4.4). The addition of these states in $\hat{X}$ will also lead to changes in Q, B_w, and F as shown in Equations (4.26)-(4.28),

$$Q = \begin{bmatrix} \sigma_f^2 & O_{33} & O_{33} & O_{33} & O_{31} & O_{31} \\ O_{33} & \sigma_\omega^2 & O_{33} & O_{33} & O_{31} & O_{31} \\ O_{33} & O_{33} & \sigma_{bf}^2 & O_{33} & O_{31} & O_{31} \\ O_{33} & O_{33} & O_{33} & \sigma_{b\omega}^2 & O_{31} & O_{31} \\ O_{13} & O_{13} & O_{13} & O_{13} & \sigma_{bclk}^2 & 0 \\ O_{13} & O_{13} & O_{13} & O_{13} & 0 & \sigma_{\dot{b}clk}^2 \end{bmatrix} \tag{4.26}$$

$$B_w = \begin{bmatrix} O_{33} & O_{33} & O_{33} & O_{33} & O_{32} \\ R_b^e & O_{33} & O_{33} & O_{33} & O_{32} \\ O_{33} & R_b^e & O_{33} & O_{33} & O_{32} \\ O_{33} & O_{33} & R_b^e & O_{33} & O_{32} \\ O_{33} & O_{33} & O_{33} & R_b^e & O_{32} \\ O_{23} & O_{23} & O_{23} & O_{23} & I_{22} \end{bmatrix} \tag{4.27}$$

$$F = \begin{bmatrix} O_{33} & I_{33} & O_{33} & O_{33} & O_{33} & O_{31} & O_{31} \\ O_{33} & O_{33} & [(R_b^e \underline{\alpha})\wedge] & R_b^e & O_{33} & O_{31} & O_{31} \\ O_{33} & O_{33} & O_{33} & O_{33} & R_b^e & O_{31} & O_{31} \\ O_{33} & O_{33} & O_{33} & -\frac{1}{\tau_f} I_{33} & O_{33} & O_{31} & O_{31} \\ O_{33} & O_{33} & O_{33} & O_{33} & -\frac{1}{\tau_\omega} I_{33} & O_{31} & O_{31} \\ O_{13} & O_{13} & O_{13} & O_{13} & O_{13} & 0 & 1 \\ O_{13} & O_{13} & O_{13} & O_{13} & O_{13} & 0 & 0 \end{bmatrix} \tag{4.28}$$

Since the closely-coupled algorithm replaces GPS position and velocity measurements with pseudorange and pseudorange rate measurements, the measurement vector z becomes

$$z = \begin{bmatrix} \rho_1 & \dots & \rho_n & \dot{\rho}_1 & \dots & \dot{\rho}_n \end{bmatrix} \tag{4.29}$$

where ρ_n and $\dot{\rho}_n$ may be modeled as

$$\rho_n = |\underline{P_{sat,n}} - \underline{\hat{P}}| + c b_{clk} \tag{4.30}$$

$$\dot{\rho}_n = \underline{u}^T |\underline{\dot{P}_{sat,n}} - \underline{\hat{V}}| + c \dot{b}_{clk} \tag{4.31}$$

Here $\underline{P}_{sat,n}$ and $\underline{\dot{P}_{sat,n}}$ represent the n^th satellite position and velocity vectors, respectively, c is the speed of light, and u is the line-of-sight vector from receiver to satellite given as

$$u = \frac{\underline{P_{sat,n}} - \underline{\hat{P}}}{|\underline{P_{sat,n}} - \underline{\hat{P}}|} \tag{4.32}$$

Contrary to the loosely-coupled algorithm, there is not a linear relationship between the measurement vector z and the state vector $\hat{X}$ in the closely-coupled algorithm. As a result, H must be computed as the Jacobian matrix of z with respect to $\hat{X}$ such that

$$H = \begin{bmatrix} u_1^T & O_{13} & O_{19} & 1 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ u_n^T & O_{13} & O_{19} & 1 & 0 \\ O_{13} & u_1^T & O_{19} & 0 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ O_{13} & u_n^T & O_{19} & 0 & 1 \end{bmatrix} \tag{4.33}$$

Finally, the R matrix of the system will also be changed such that

$$R = \begin{bmatrix} \sigma_\rho^2 & O_{33} \\ O_{33} & \sigma_{\dot{\rho}}^2 \end{bmatrix} \tag{4.34}$$

4.2 MGVA and MVVA Filter Algorithms

Due to the interesting and beneficial nature of the closely-coupled architecture, it was chosen as a basis for the modified algorithms developed in this thesis. As a result, the modified algorithms are in large part the same as that of the standard closely-coupled algorithm. The modifications made to the algorithms were all for determining the benefits of fusing magnetic sensors with the other sensors already used in the standard architecture. In specific, it was believed that the fusion performed in the modified algorithms would ultimately lead to improved performance when compared with the standard algorithm.

It has been discussed in previous chapters that magnetometers are very useful for both heading stabilization and six DoF attitude determination. Since magnetometers have often

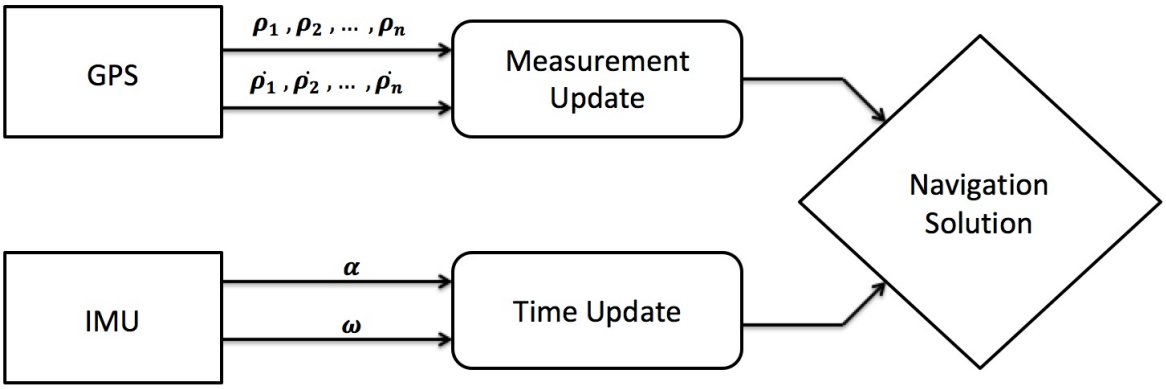

Figure 4.1: Block diagram of the standard closely-coupled EKF algorithm

been used for six DoF attitude determination, it was chosen in this thesis to use magnetometers to provide an attitude measurement update to the standard closely-coupled EKF. Figures 4.1 - 4.3 help to explain how these attitude measurement updates were implemented. Figure 4.1 displays a block diagram of the standard closely-coupled algorithm. As can be seen, INS sensors are mechanized and used to propagate the current state estimate in the time update. These state estimates are then corrected using pseudorange and pseudorange rate measurements in the measurement update.

The change in the algorithm can be seen in Figures 4.2 and 4.3. In the modified algorithms, it is proposed that measurements of the earth's magnetic field could be paired with gravity vector or velocity vector measurements (whenever available) and used to calculate a "measurement" of the current system attitude using the algorithm described in Chapter 3. As implied by the names of the filters, the MGVA filter would utilize magnetometer and accelerometer measurements (the accelerometer measuring the gravity vector), along with models of the earth's magnetic field and gravity vector as in Figure 4.2. The MVVA filter would utilize magnetometer and body frame velocity vector measurements (taken from vehicle CAN wheel speed measurements), along with a model of the earth's magnetic field and

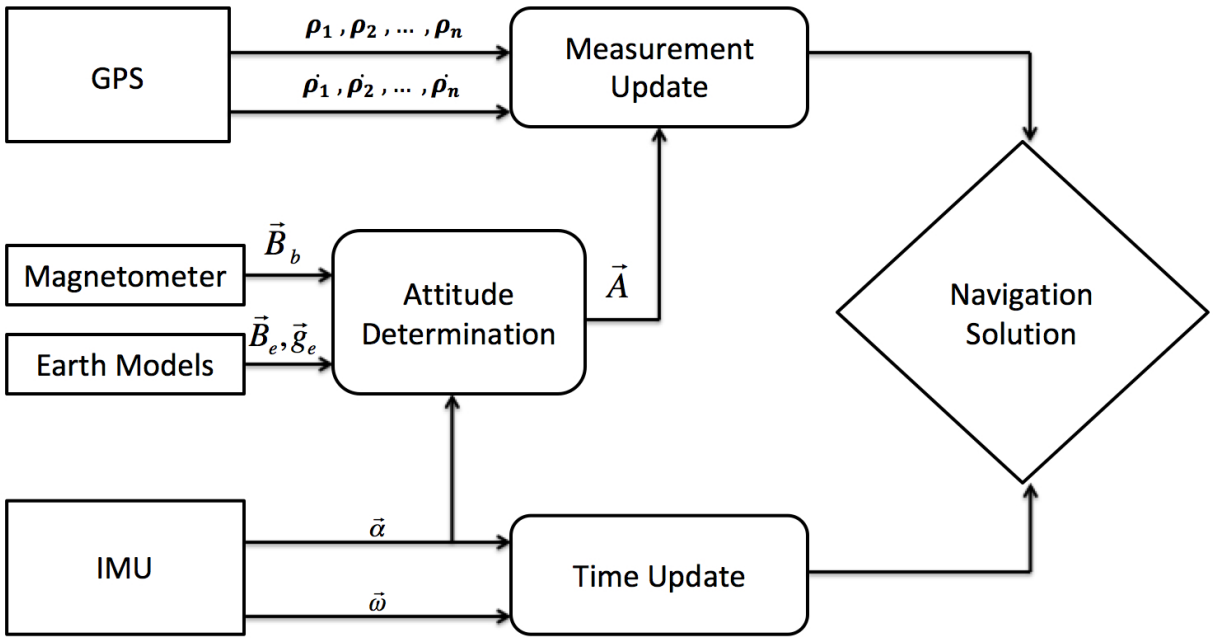

Figure 4.2: Block diagram of the MGVA filter

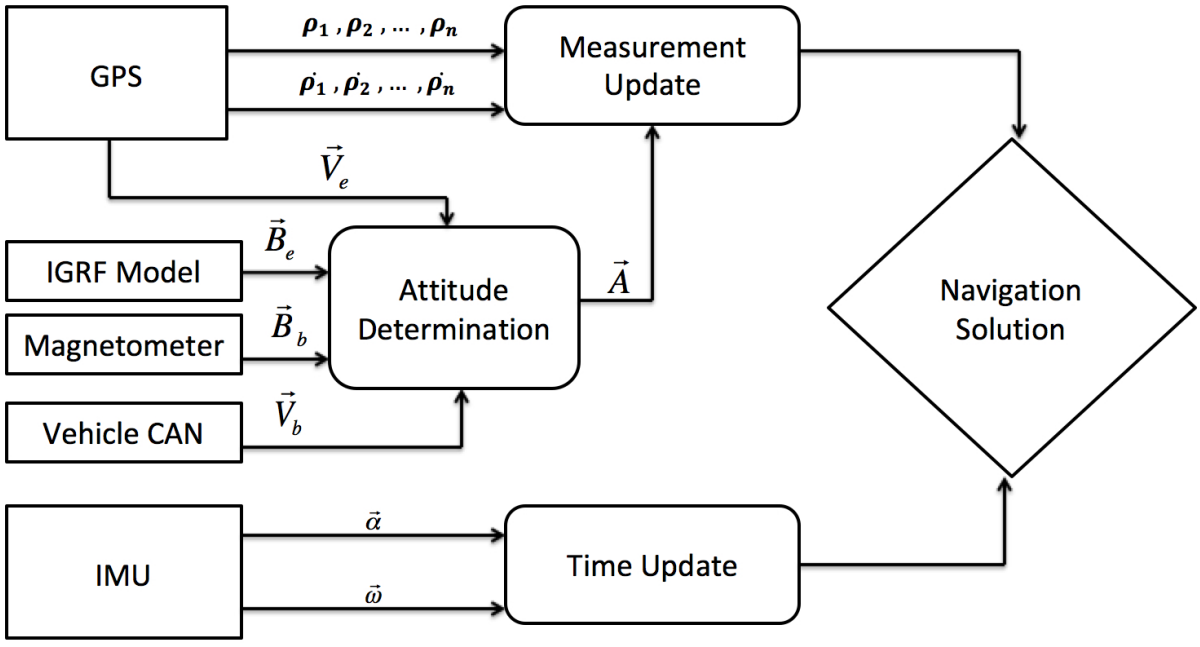

Figure 4.3: Block diagram of the MVVA filter

GPS velocity vector measurements as in Figure 4.3. The attitude measurement update created from these processes could then be added to the measurement vector and used together with pseudorange and pseudorange rate measurements to correct the current state estimate.

This modification made to the standard algorithm will not have an effect on the time update portion of the algorithm. However, it will affect the measurement update. The measurement vector z will be modified to include the attitude calculation $\underline{A}$ as in Equations (4.35) and (4.36)

$$\underline{A} = \begin{bmatrix} \phi & \theta & \psi \end{bmatrix}^T \tag{4.35}$$

$$z = \begin{bmatrix} \rho_1 & \dots & \rho_n & \dot{\rho}_1 & \dots & \dot{\rho}_n & \underline{A} \end{bmatrix} \tag{4.36}$$

where ϕ, θ, and ψ are roll, pitch and yaw angles respectively. This will result in a new H matrix which relates the new attitude measurement updates to the attitude estimates in the state vector. It will also lead to a new R matrix as shown below.

$$H = \begin{bmatrix} u_1^T & O_{13} & O_{13} & O_{16} & 1 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ u_n^T & O_{13} & O_{13} & O_{16} & 1 & 0 \\ O_{13} & u_1^T & O_{13} & O_{16} & 0 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ O_{13} & u_n^T & O_{13} & O_{16} & 0 & 1 \\ O_{33} & O_{33} & I_{33} & O_{36} & O_{31} & O_{31} \end{bmatrix} \tag{4.37}$$

$$R = \begin{bmatrix} \sigma_\rho^2 & O_{33} & O_{33} \\ O_{33} & \sigma_{\dot{\rho}}^2 & O_{33} \\ O_{33} & O_{33} & \sigma_A^2 \end{bmatrix} \tag{4.38}$$

The addition of the attitude measurements results in the addition of an identity matrix to the last row of the modified H matrix. This is important to note because it creates a direct

relationship between the observed attitude and the estimated attitude. This information will become essential later when the observability of the MGVA and MVVA filters is discussed.

It is also important to note that when accelerometer measurements are paired with magnetometer measurements for attitude determination, the attitude determination algorithm is able to provide attitude updates at a much faster rate than GPS is provided to the algorithm. For this reason, when attitude updates are available but GPS is not, z, H, and R may be redefined as in Equations (4.39)-(4.41).

$$z = \begin{bmatrix} \vec{A} \end{bmatrix} \tag{4.39}$$

$$H = \begin{bmatrix} O_{33} & O_{33} & I_{33} & O_{36} & O_{31} & O_{31} \end{bmatrix} \tag{4.40}$$

$$R = \begin{bmatrix} \sigma_A^2 \end{bmatrix} \tag{4.41}$$

This ability of the MGVA filter to operate in the absence of GPS measurements carries interesting implications with it. Because the MGVA filter is able to operate in the absence of GPS measurements, it is significant to note that the MGVA filter has the potential to improve navigation performance during GPS signal outages. If the MGVA filter can provide good attitude and accelerometer bias estimation during GPS signal outages, then it is possible that the MGVA filter can ultimately allow the propagated state estimates from an IMU to be corrected even during GPS outages. An exploration of this capability has not be performed in this thesis, however, and as a result is left for future work.

4.3 Conclusion

Many navigation algorithms have been utilized in the past and are suitable for various types of vehicle navigation. From among these, one of the most common navigation algorithms is the Extended Kalman Filter. In specific, the loosely-coupled architecture of

the EKF seeks to estimate an accurate PVA solution as well as inherent accelerometer and gyroscope biases. The loosely-coupled EKF forms predictions of these states by propagating its estimates forward using accelerometer and gyroscope observations, then corrects these state estimates using GPS position and velocity observations when available.

In contrast to the loosely-coupled EKF, the closely-coupled EKF utilizes GPS pseudorange and pseudorange rate observations, rather than position and velocity observations. In doing so, it improves the performance of the EKF in GPS degraded environments. The effect of this is that GPS receiver clock bias and drift terms are added to the state vector of the system, and GPS pseudorange and pseudorange rate measurements replace GPS position and velocity observations in the system measurement vector. This also results in a new measurement matrix which is formed as the jacobian matrix of the measurement vector with respect to the state vector.

The architecture of the MGVA and MVVA filters was presented which utilizes magnetometers in order to provide the standard loosely and closely-coupled EKFs with attitude measurement updates. It was chosen to base the MGVA and MVVA filters on the closely-coupled filter due to the ability of the closely-coupled filter to work with limited functionality in GPS degraded environments. This ability is an improvement over the loosely-coupled architecture, which is often unable to function in GPS degraded environments. The MGVA filter pairs magnetometer measurements with accelerometer measurements, while the MVVA filter pairs magnetometer measurements with velocity measurements. In doing so the MGVA and MVVA filters form calculations of the current system attitude. This calculated attitude is then applied as an attitude update during state correction. The effect of applying an attitude update to the system is such that a direct relationship between observed and estimated attitude is created. This will become essential later on when the observability of the MGVA and MVVA filters is considered.

Chapter 5

Observability of the MGVA and MVVA Filters

5.1 Observability Theory

Observability is an integral part of determining the capabilities of any given navigation or estimation system. In general, the observability of a system can be defined as a system's ability to determine state variables based on system dynamics and measurements. For example, if it is desired to know the distance between two points, a device or tool that measures length (like a ruler or a tape measure) may be able to "observe" the distance between the two points. In this way, it can be said that the system is observable. However, if such a tool is not available, or if the only tools available cannot measure length (e.g., a level or a drafting compass), then the system is said to be unobservable.

Because of this, the unobservability or observability of a system holds serious implications for a system. As was discussed in Section 1.1, there has been an extensive amount of work done in the area of observability. Goshen-Meskin and Bar-Itzhack were at the forefront of this thesis when they put forth a theoretical approach to the observability analysis of linear time-invariant (LTI) systems [18, 19]. In [20], Hong was able to show that unobservable position, velocity, and accelerometer bias can arise from errors in attitude, gyro bias, and GPS lever arm estimates.

In [21], Rhee investigated the observability of the standard loosely-coupled EKF for various aerial vehicle maneuvers. Rhee was able to show that for low excitation trajectories of constant linear acceleration, the attitude and x,y accelerometer biases of the loosely-coupled filter are unobservable. It was also shown that for trajectories with changes in acceleration, or steady-state turning, various modes of the filter experience a type of 'excitation' that causes them to become observable. Ryan also examined these trajectories for a modified

loosely-coupled filter which relied on sensors commonly found in ground vehicle electronic stability control (ESC) systems [23]. Ryan's findings were consistent with Rhee's in that he showed that his modified filter was only observable during vehicle maneuvers which excite various modes of the filter. Specifically, both Rhee and Ryan showed that for low excitation trajectories, the attitude as well as the x and y accelerometer bias states are unobservable.

Examining a model of accelerometer inputs, the reasoning behind this becomes evident.

$$\underline{\alpha_{meas}} = \underline{\alpha_{true}} + C_e^b \underline{g_e} + \underline{b_\alpha} + \underline{\epsilon_\alpha} \tag{5.1}$$

As Ryan shows in [22], accelerometer measurements $\underline{\alpha_{meas}}$ are a function of the true acceleration $\underline{\alpha_{true}}$, acceleration due to gravity $\underline{g_e}$, accelerometer bias b_α, and noise ϵ_α. In Equation (5.1) it is evident that if the estimated attitude of the system C_e^b is incorrect, the filter will not be able to detect the true effects of gravity $C_e^b g_e$. Because of this, if the attitude is incorrect, the filter will lump errors due to attitude together with errors due to accelerometer bias causing the combination of states to become unobservable. In contrast, if the system experiences enough excitation in the various axes of the accelerometer, the accelerometer bias b_α and the true attitude C_b^e will become observable together. This ultimately leads to two essential facts. First, the combination of attitude states and accelerometer bias states are observable together. Specifically, the system equations of the navigation EKF are such that very strong couplings exist between the x-axis accelerometer bias and the pitch angle, and between the y-axis accelerometer bias and the roll angle. Second, if the system does not experience enough dynamics, errors due to both attitude and bias estimates will become lumped together, and cause the combination of estimates to become unobservable.

This lack of observability ultimately leads to the reasoning behind the modified filters proposed in this thesis. Since accelerometer bias and attitude are states that are observable together, it is proposed that if the attitude of the system can be observed during low excitation trajectories, then the accelerometer bias should also become observable. Because of this, the MGVA and MVVA filters should become observable even during low excitation

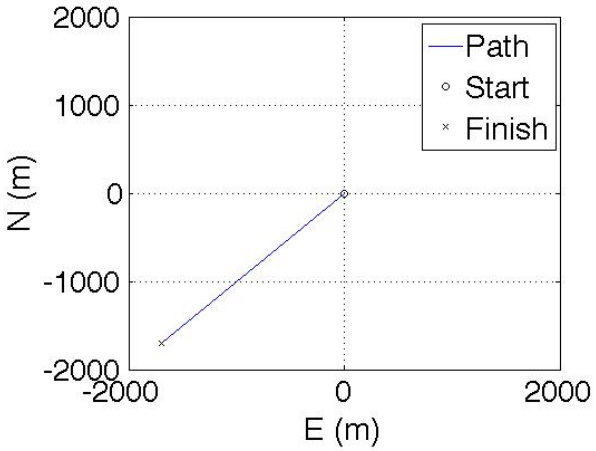

Figure 5.1: Trajectory of a simulated ground vehicle driving straight with constant velocity

trajectories, because they both provide attitude updates to the system (via the aiding from the magnetometer and acceleromter/velocity measurement) in all scenarios.

5.2 Simulation Results

In order to test the observability of the MGVA and MVVA filters, it was chosen to first examine the filter for a simulated low excitation trajectory. A ground vehicle, traveling at constant velocity and heading for a time period of 120 seconds was simulated. Figure 5.1 shows the trajectory of the simulated vehicle in the North-East plane. Accelerometer, gyroscope, magnetometer, velocity, GPS pseudorange, and GPS pseudorange rate measurements were all simulated, corrupted and then implemented on both the standard closely-coupled EKF as well as the MGVA and MVVA filters. The error parameters used to corrupt the simulated magnetometer were the same as those used in Chapter 2. These parameters can be found in Table 2.1. With regard to the accelerometer bias errors, all biases were modeled as first order Markov processes (see Equation 4.1). To simulate these processes, a time constant τ of 10, and an initial bias of 0.5 $(\frac{m}{s^2})$ were used. Additionally, the standard deviation of the process driving white-noise μ was 0.0333 $(\frac{m}{s^2})$.

As a metric for evaluating the observability, Equation (5.2) given in [48] was used

$$O_{LTV}(t_f, t_0) = \int_{t_0}^{t_f} F^T(\tau, t_0) H^T H F(\tau, t_0) d\tau \tag{5.2}$$

Here, $O_{LTV}(t_f, t_0)$ is the observability grammian, evaluated on the interval from t_0 to t_f, and observability of a system is indicated if the grammian is non-singular. In order to tell whether or not O_{LTV} is non-singular, the rank of the matrix is analyzed. If the grammian matrix is full rank, this indicates that the matrix is non-singular, and therefore the system is observable. It is important to note that the grammian is only valid on the interval from t_0 to t_f, and that in general, if the grammian reaches full rank at any point on that interval, it will remain full rank for the rest of the interval. Because of this, the time period that it takes the grammian rank to stabilize will be important to consider moving forward.

Figures 5.2 and 5.3 display the grammian rank for the first twenty seconds of the simulation with both the standard filter and the modified filters (MGVA and MVVA) respectively. Because both modified filters performed identically during observability testing, the results from both filters are represented by Figure 5.3. Also, only the first twenty seconds of the simulation are displayed because, for these tests, once the rank reached a steady state value it remained there for the remainder of the simulation. As can be seen, the standard filter grammian reaches a rank of 14, and then is rank deficient by three for the remainder of the simulation. Whereas, with the modified filters, the grammian rank is increased by three, making the system full rank, and then remains full rank for the remainder of the simulation. This does not come as a surprise because, referring to Equation (4.37) it can be seen that the addition of the attitude measurement update to the modified filters serves to add a 3x3 identity matrix to the last row of the H measurement matrix in both filters. This effectively creates three direct paths between the filter states and the filter measurements, thus making the attitude states observable and increasing the rank of both modified filters by three. It is also important to note the rise time of the grammian rank. The grammian rank is only valid inside time interval it is tested for (t_0 to t_f). Changing the window of time over which the grammian is examined will change the results seen in testing. If for example, a system is

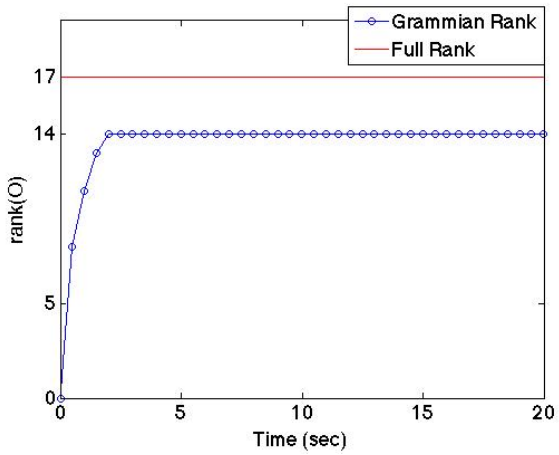

Figure 5.2: Rank of the observability grammian, evaluated for the standard (17 state) closely-coupled filter, over the first twenty seconds of the simulation

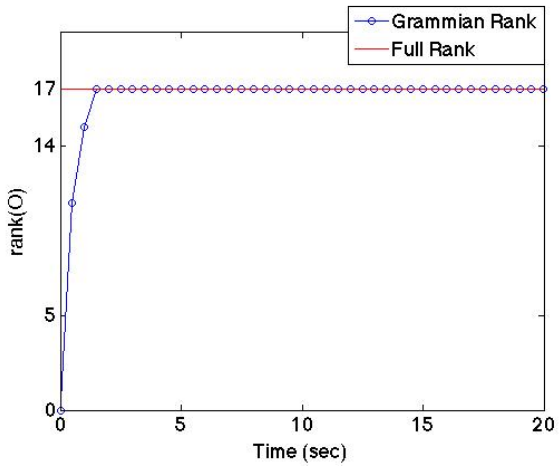

Figure 5.3: Rank of the observability grammian, evaluated for the MGVA and MVVA filters, over the first twenty seconds of the simulation

examined over a time interval of 5 seconds, and the grammian is seen to rise to full rank, this does not imply that the system is full rank over a shorter interval (e.g., two seconds). When tested over a two second interval, the grammian may still be rank deficient. As a result, the rise time of the grammian rank is important because it tells the shortest interval over which the rank is valid. For these tests, the rise time for the standard filter was 2 seconds, while the rise time of the modified filters was 1.5 seconds.

Examining Figures 5.4 - 5.9, the effectiveness of the standard filter, the MGVA filter, and the MVVA filter during the simulation can be seen. Figure 5.4 displays the attitude

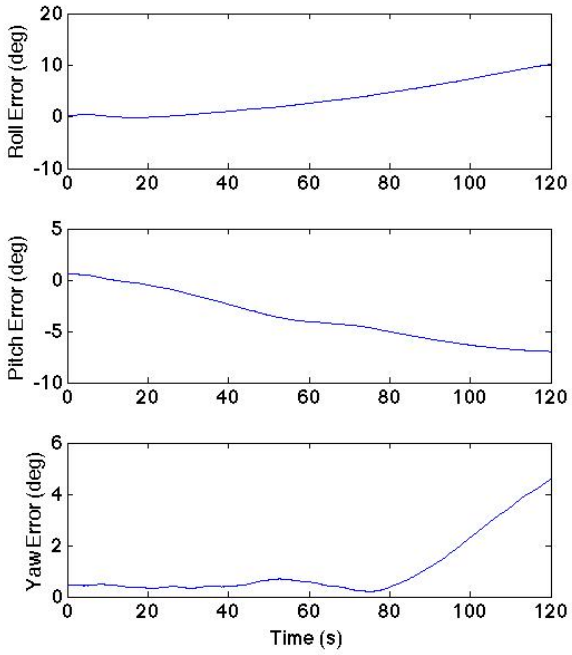

Figure 5.4: Attitude error of the standard closely-coupled filter

estimate errors over the course of the simulation for the standard filter. As can be seen, a combination of initial attitude errors and drifting estimates due to a lack of observability, cause the estimate error to grow unbounded. Figure 5.5 shows that this ultimately leads to an inability of the filter to estimate the x and y accelerometer biases. The estimates begin near the true bias, but drift off just as the attitude estimate does. In contrast, the z accelerometer bias is still observable. This lines up with the findings of Ryan in [23], since the z accelerometer bias should be observable always, regardless of dynamics.

In contrast, Figure 5.6 shows that the MGVA filter is able to accurately estimate the attitude of the vehicle. In all three euler angles, sub-degree levels of accuracy are achieved by providing attitude updates from the magnetometer and accelerometer measurements. Additionally, Figure 5.7 shows that good estimation of the vehicle attitude allows the filter to accurately observe the accelerometer bias in the x and y axes. As the attitude estimate errors of the simulation begin to converge to zero near 30 seconds, the bias estimates of the filter can also be seen to converge to their true values.

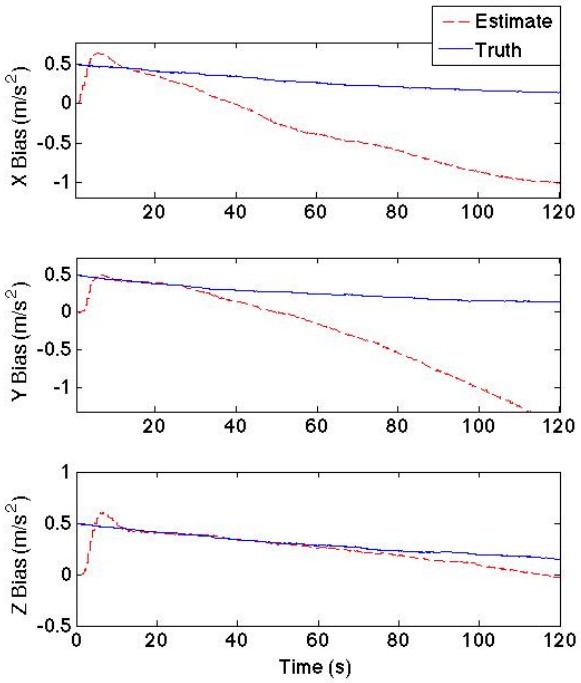

Figure 5.5: Accelerometer bias estimates vs. truth for the standard closely-coupled filter

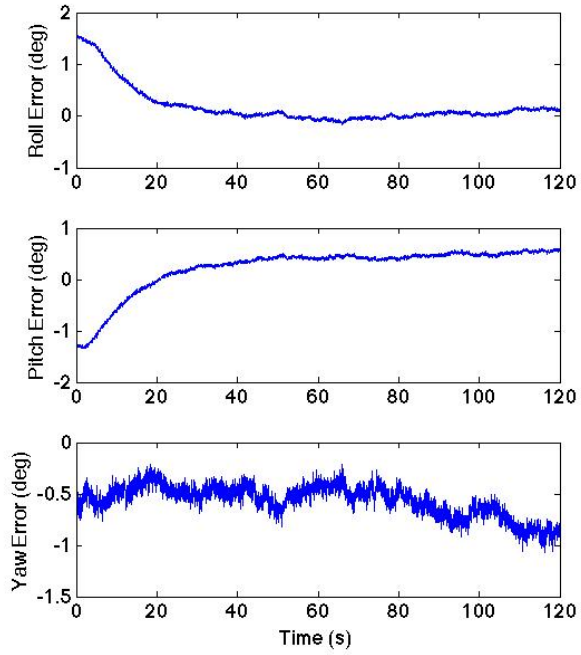

Figure 5.6: Attitude error of the MGVA filter

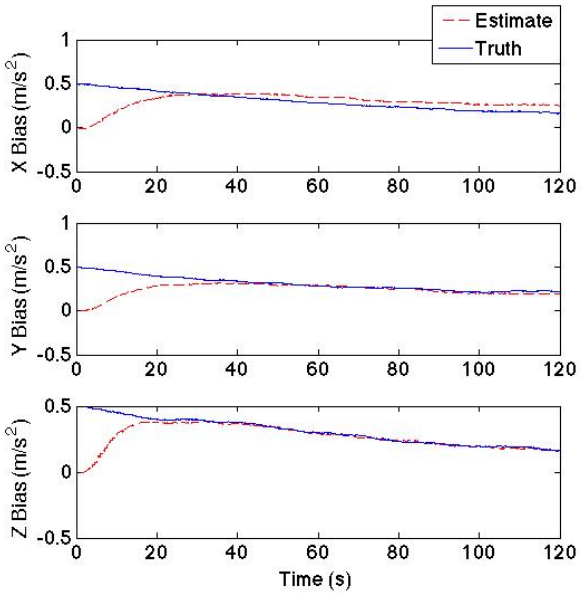

Figure 5.7: Accelerometer bias estimates vs. truth for the MGVA filter

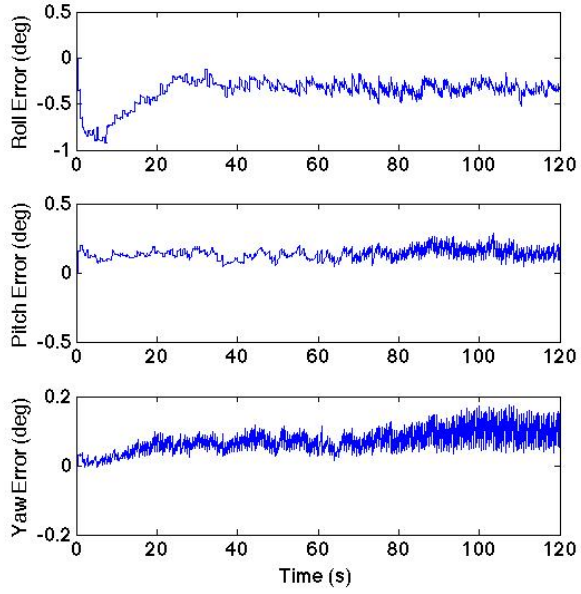

Figure 5.8: Attitude error of the MVVA filter

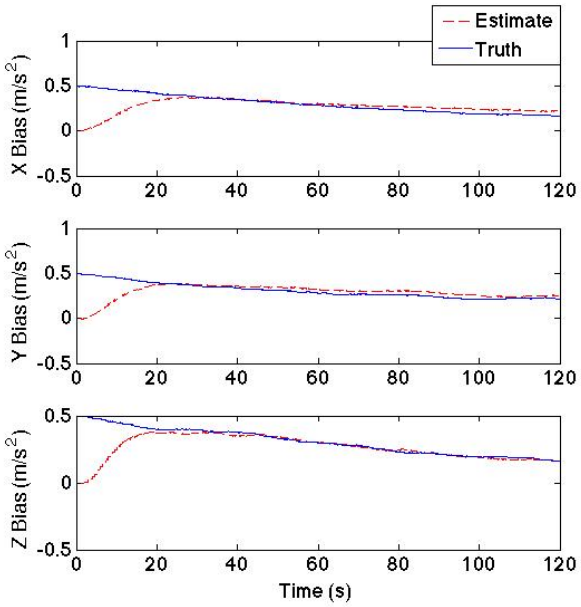

Figure 5.9: Accelerometer bias estimates vs. truth for the MVVA filter

Similar results are seen using the MVVA filter. It can be seen in Figure 5.8 that the MVVA filter outperforms the MGVA filter in several regards during the simulation. First, no part of the attitude estimate ever exceeds more than a degree of error using the MVVA filter, while the MGVA filter experiences initial errors greater than 1 degree in roll and pitch, and must converge to the true attitude over time. Second, it can be seen that the yaw error in the MGVA filter slowly grows as the simulation runs. However, with the MVVA filter, the yaw error stabilizes near 0.2 degrees. Finally, examining Figure 5.9, it can be seen that with the MVVA filter, the accelerometer bias estimates are able to converge to the true values once both the pitch and roll estimates stabilize around 30 seconds.

5.3 Experimental Results

Having validated the MGVA and MVVA filters in simulation, the filters were further tested using real data collected on a ground vehicle. The modified Infinity G35 test vehicle that was described in Chapter 2 was once again used here. The vehicle was outfitted with a three-antenna Septentrio GPS system, and an XSens MTi IMU was mounted inside the

Figure 5.10: NCAT's 1.7 mile long test track where the test vehicle was driven

vehicle, aligned with the vehicle body axes. In addition to these sensors, the vehicle CAN messages were also collected and used to calculate the velocity vector of the vehicle at any moment in the body frame. The vehicle was taken to the National Center for Asphalt Technology (NCAT) in Auburn, Alabama. NCAT's facilities posses a 1.7 mile long test track (shown in Figure 5.10) where ground vehicle's may travel at high speeds. The track is made up two long straight portions which are connected by 8 degree banked turns. As can be seen in Figure 5.11, the vehicle was driven with constant velocity for several full revolutions around the track collecting data from the XSens, vehicle CAN system, and the GPS Septentrio system.

The XSens collected accelerometer, gyroscope, and magnetometer data, while the CAN system collected wheel speed measurements, and the Septentrio collected GPS pseudorange, pseudorange rate, position, velocity, and attitude measurements. The accelerometer, gyroscope, pseudorange, and pseudorange rate measurements were all implemented on the standard closely-coupled filter. The MGVA filter used accelerometer and magnetometer measurements from the XSens, along with an IGRF model and a WGS84 gravity model to provide attitude updates to the system. The MVVA filter used wheel speed measurements from the CAN system to calculate the velocity vector of the vehicle in the body frame, and then this measurement was used along with the magnetometer measurements to provide body frame vector measurements for the attitude determination algorithm. In order to provide global frame vector measurements an IGRF model was used once again, but this

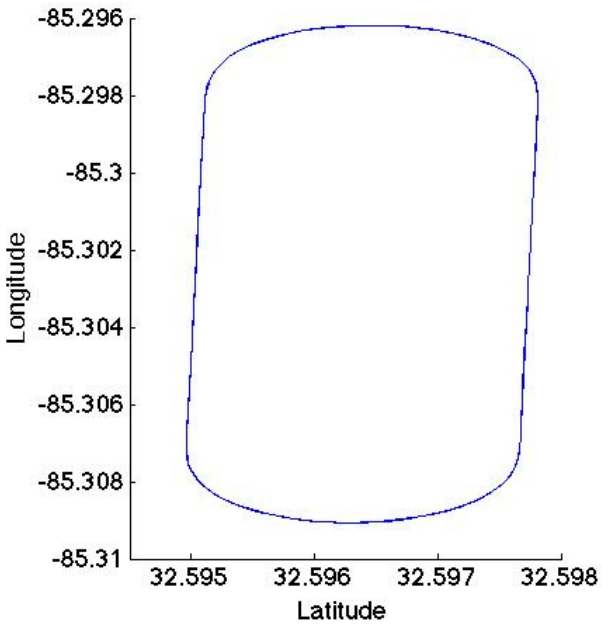

Figure 5.11: Trajectory driven by the test vehicle at NCAT

time with GPS velocity measurements. Position, velocity and attitude measurements from the Septentrio were all used as a truth reference for the vehicle trajectory. Finally, in order to have a reference for the actual accelerometer bias corrupting the sensor measurements, attitude updates from the Septentrio were run with the standard filter to obtain estimates of the accelerometer bias. The mean value of the steady-state bias estimates was used to define what the 'true' accelerometer bias. However, for the tested algorithms (the standard filter, the MGVA filter, and the MVVA filter), no GPS based attitude updates were utilized to improve estimation. Only the methods defined in Chapter 4 were used during testing. Results from all three filters are discussed below.

Figure 5.12 displays the accelerations and yaw rates measured by the XSens IMU during the run. As can be seen, the only periods where the vehicle experiences acceleration is during turning. Other than this, accelerations in both the x and y axes are kept near zero, and the yaw rate of the vehicle is also kept near zero.

Figures 5.13 and 5.14 show the attitude estimate and the accelerometer bias estimates for the standard closely-coupled filter during this run. As can be seen in Figure 5.13, the

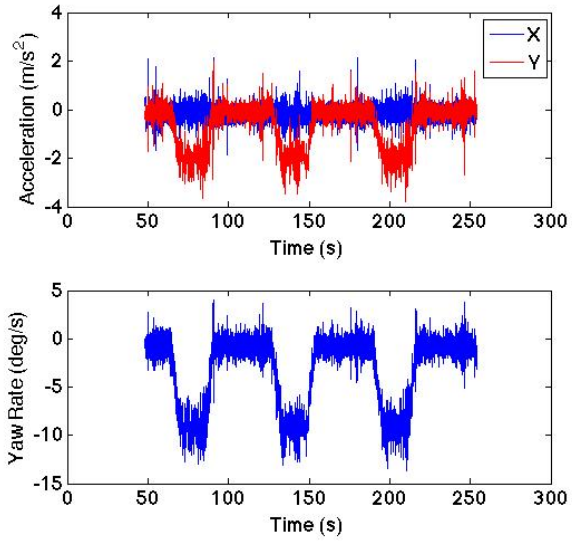

Figure 5.12: Accelerations and yaw rates experienced by the vehicle during testing at NCAT

filter has a reasonable amount of success in estimating the roll of the vehicle. It is able to achieve estimates within several degrees of the true roll. The errors that the filter experiences with regard to the roll angle are most evident during turning of the vehicle, and during the second straight driving portion of the trajectory. It can be seen that in these moments the estimates struggle due to small amounts of drift that pull them away from the true angle. However, during the final straight portion of the track, the roll estimate once again converges to the true value. Examining the trajectory of the vehicle, it seems likely that the reason behind the standard filter's success in roll estimation lies in the lateral excitation that the vehicle experiences on the turns. While the estimate still struggles from drift, it is quickly able to reconverge to the true roll angle when the vehicle experiences a change in acceleration (moments where the y axis acceleration changes in correspondence with the changing yaw rate). In contrast, the filter is not able to estimate the pitch angle of the vehicle well. In particular, during turning the pitch estimate diverges from the true pitch angle by about 6 degrees. Finally, the filter appears to be successful in estimating the yaw of the vehicle,

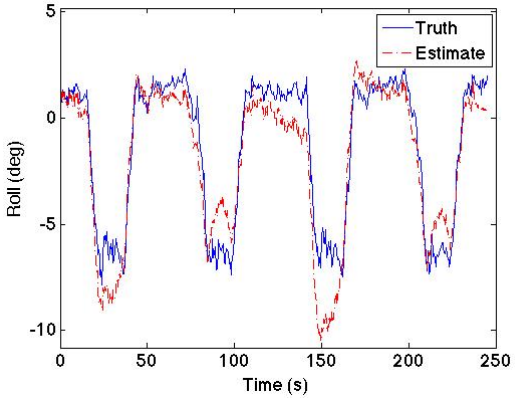

(a) Roll estimate vs. truth

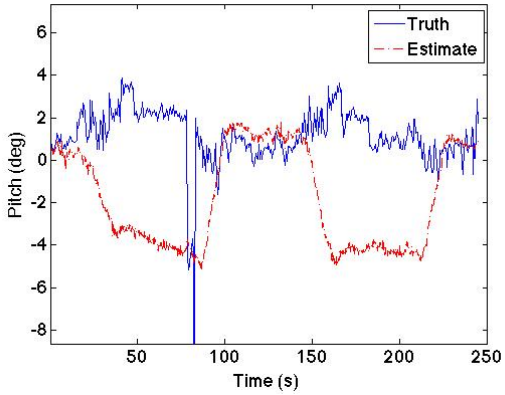

(b) Pitch estimate vs. truth

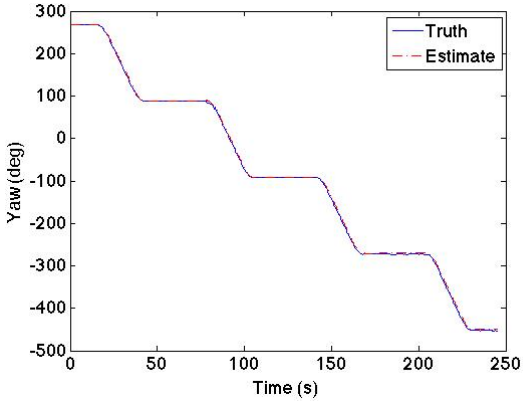

(c) Yaw estimate vs. truth

Figure 5.13: Attitude estimate vs. truth for the standard closely-coupled filter during testing at NCAT

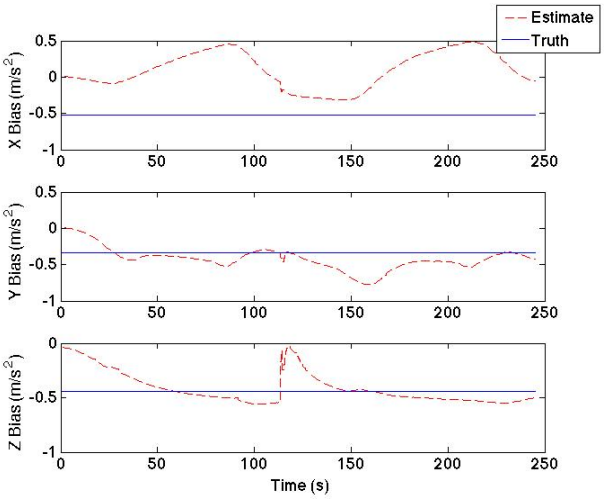

Figure 5.14: Bias estimates vs. true bias for the standard closely-coupled filter during experiments at NCAT

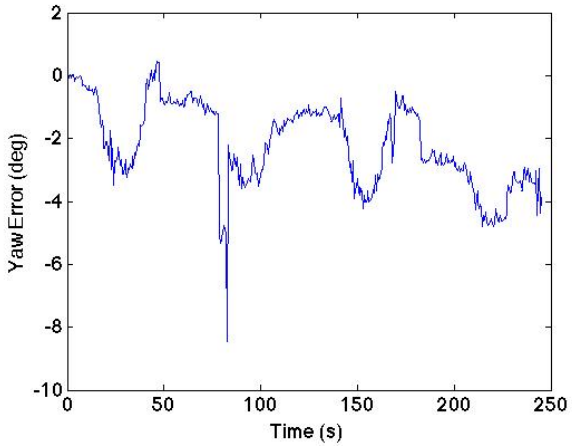

Figure 5.15: Yaw estimate error for the standard closely-coupled filter during experiments at NCAT

however a closer look at the data reveals that the estimate actually suffers from a small amount of heading drift.

Figure 5.15 displays the yaw estimate error over the course of the run. As can be seen, the estimate error grows over time because the filter is unable to observe the true heading of the system. Examining Figure 5.14 shows that the system is unable to observe the x axis accelerometer bias. Rather than converging toward the true bias, the estimate oscilates around an incorrect estimate. Studying Figure 5.13, it can be seen that this is largely due to the filter's incorrect estimate of the vehicle pitch angle. As the pitch estimate varies and diverges from the true pitch angle, so the bias estimate varies and is never able to observe the true x axis bias. However, the filter has a reasonable amount of success in estimating the y axis accelerometer bias. Again, looking back at Figure 5.13, the ability of the filter to estimate the roll angle of the vehicle is directly tied to it's ability to estimate the y axis accelerometer bias. As the roll angle converges and diverges from the true roll angle, the bias estimate also converges and diverges from the true bias. Finally, it can be seen that the standard filter is able to accurately estimate the z-axis accelerometer bias. This does not come as a surprise since the z-axis accelerometer bias should always be observable.

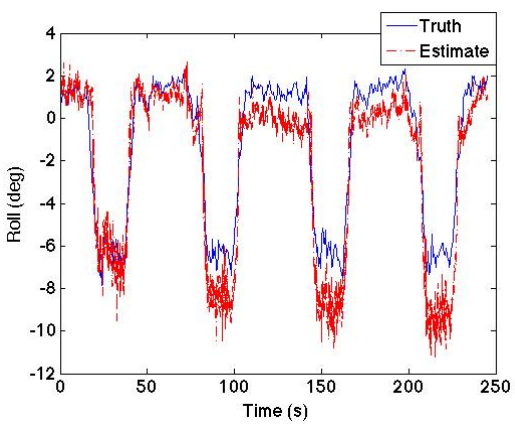

(a) Roll estimate vs. truth

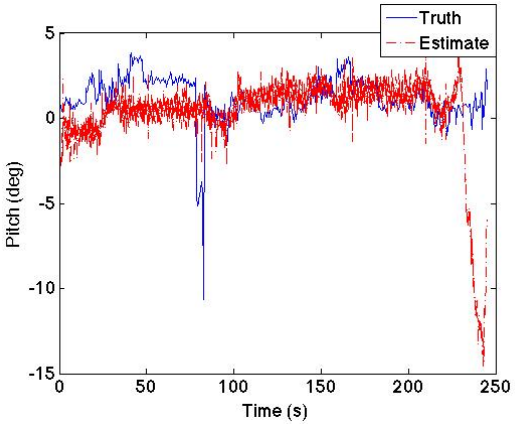

(b) Pitch estimate vs. truth

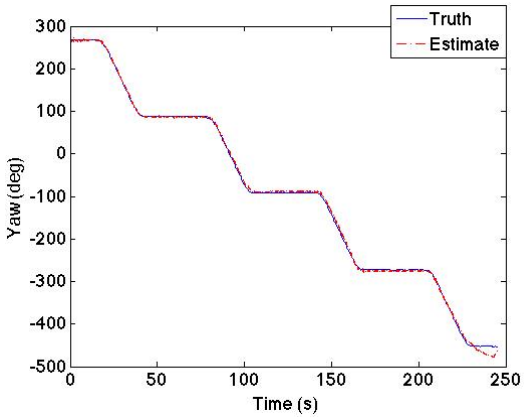

(c) Yaw estimate vs. truth

Figure 5.16: Attitude estimate vs. truth for the MGVA filter during testing at NCAT

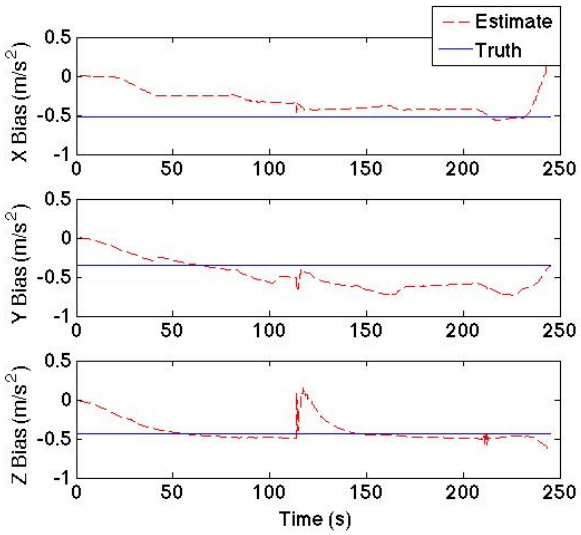

Figure 5.17: Bias estimates vs. true bias for the MGVA filter during experiments at NCAT

Figures 5.16 and 5.17 show results from the same test, but for the MGVA filter. It can be seen that the MGVA filter succeeds in estimating the attitude of the vehicle. The filter is able to estimate the attitude within several degrees of the reference for all three euler angles. Correspondingly, Figure 5.17 shows that the filter also succeeds in estimating the accelerometer bias effecting the system. With regard to both the x and y axis accelerometer bias, the filter produces estimates which converge to values that are near the reference values of bias. Specifically, comparing Figures 5.14 and 5.17, it can be seen that the MGVA filter shows a large amount of improvement in the estimation of x axis accelerometer bias.

It should be noted that in order for the MGVA to function during this experimental run, it was necessary to account for centripetal acceleration experienced by the vehicle. Examining Figure 5.12, it can be seen that while turning, the vehicle experiences lateral acceleration due partially to centripetal acceleration and partially to gravitational acceleration acting in the y axis as the vehicle rolls on the banked turns of the track. The MGVA filter utilizes accelerometer measurementes to provide the attitude determination algorithm with an observation of the Earth's gravity vector. However, an accelerometer mounted to a turning vehicle, will measure not just the effects of gravity, but also any centripetal accelerations encountered by the vehicle. If this raw measurement is used in the attitude determination algorithm, it will inevitably cause the algorithm to calculate an incorrect attitude solution. This would lead to an inability of the MGVA filter to accurately estimate the attitude of the vehicle. As a result, the centripetal acceleration experienced by the vehicle must be accounted for. In order to do this during experimental testing, the vehicle yaw rate measured by the gyroscope and the vehicle velocity calculated from vehicle CAN measurements, were both used to calculate the centripetal acceleration encountered by the vehicle as in Equation 5.3

$$\alpha_{cent} = \dot{\psi} V_t \tag{5.3}$$

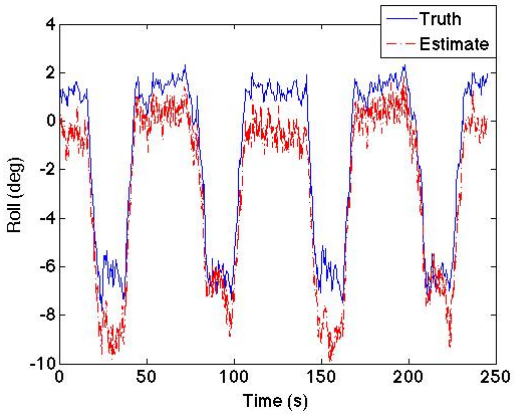

(a) Roll estimate vs. truth

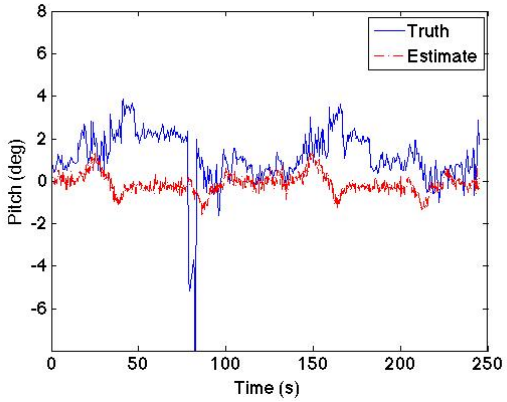

(b) Pitch estimate vs. truth

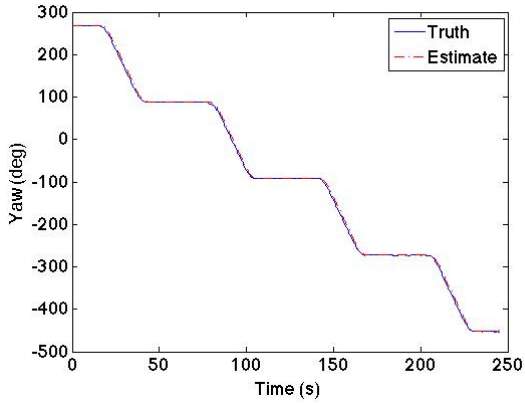

(c) Yaw estimate vs. truth

Figure 5.18: Attitude estimate vs. truth for the MVVA filter during testing at NCAT

where α_{cent} is the calculated centripetal acceleration, $\dot{\psi}$ is the measured yaw rate, and V_t is the tangential velocity of the vehicle. Subtracting α_{cent} from the y axis acceleromter measurement ultimately accounts for the centripetal acceleration experienced by the vehicle, and allows the MGVA filter to continue functioning while turning.

Finally, Figures 5.18 and 5.19 display the attitude and accelerometer bias estimates for the MVVA filter. The MVVA filter is able to estimate the roll, pitch, and yaw of the vehicle to within several degrees of the truth during the run. Additionally, Figure 5.19 shows that the MVVA filter is able to observe the accelerometer bias in all three axes. Corresponding to it's pitch estimate, the MVVA filter's x axis bias estimate is able to converge toward the

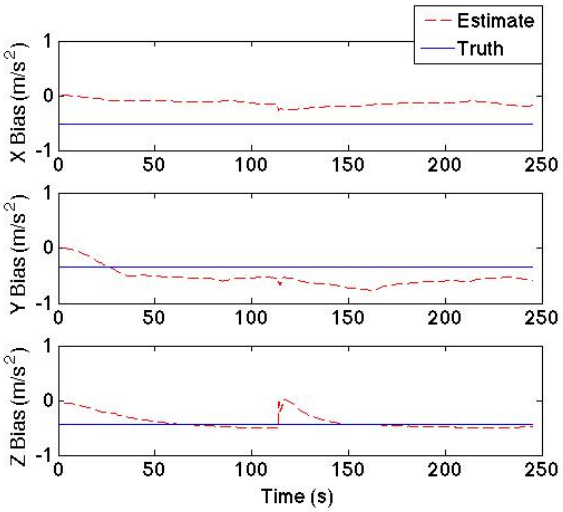

Figure 5.19: Bias estimates vs. true bias for the MVVA filter during experiments at NCAT

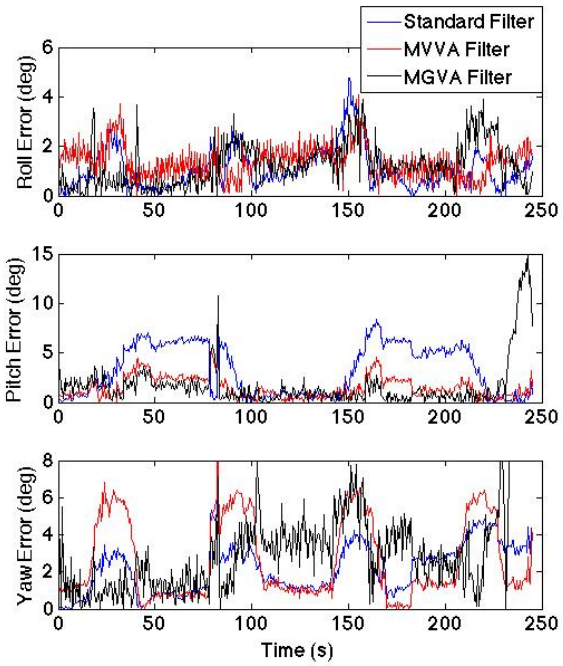

Figure 5.20: Absolute attitude error of the standard filter and the MVVA filter

true bias, while the standard filter's x axis estimate was not. Similarly, the y axis and z axis bias estimates of the MVVA filter are both able to converge to their true values.

Figure 5.20 displays a comparison of the roll, pitch, and yaw estimate errors between the standard filter, the MGVA filter, and the MVVA filter. It can be seen that, the MGVA and MVVA filters are able to provide roll estimates on par with the standard filter. With regard to the pitch angle, the MVVA filter is more accurate than the standard filter by several degrees during the run, and the MGVA filter is more accurate than the MVVA filter through out. With regard to yaw, the MGVA filter is initially more accurate than both the standard filter, and the MVVA filter. However, as the experiment progresses, the MVVA filter yaw error grows until it is larger than both the standard filter and the MVVA filter. Comparing, the MVVA filter and the standard filter, the standard filter is at first able to provide more accuracy than the MVVA filter. However, closer examination of the end of the experiment shows that the standard filter's estimate error during striaght driving grows during the experiment (reaching a maximum of 4 degrees), while the MVVA filter's estimate is bounded (remaining near 1.5 degrees of error). Because of this, it appears that the MVVA filter is superior in its ability to prevent the drift of the yaw estimate over time.

5.4 Conclusion

The topic of navigation filter observability is essential for fully determining the capabilities of any navigation scheme. Those who have studied this topic in the past have found that for the EKF, trajectories of constant acceleration and straight driving cause the attitude and x, y accelerometer bias of a system to become unobservable. In this chapter, the modified filters introduced in Chapter 4 have been proposed as a solution to the problem of unobservable trajectories. It was hypothesized that if the MGVA and MVVA filters could successfully apply accurate attitude updates from magnetometer measurements, then perhaps the filters would be able to not only observe the attitude of a system, but also any inherent x, y accelerometer biases as well.

A vehicle driving straight, at constant speed, for a prolonged period of time, was first simulated. The inputs, and measurements from this simulation were then implemented on the standard filter, the MGVA filter, and the MVVA filter. The MGVA filter utilized magnetometer and accelerometer measurements to provide the filter with attitude updates, while the MVVA filter used velocity vector and magnetometer measurements to do so. The observability of all three filters was tested in simulation, and it was shown that the standard filter, for the given trajectory, was only able to reach an observability rank of 14, with full rank being 17. In contrast, the MGVA and MVVA filters, for the case of straight and steady driving, were able to achieve full rank observability over the course of the simulation. Additionally, it was shown that the MGVA and MVVA filters were able to provide much more accurate estimates of attitude and accelerometer bias for the trajectory.

To further validate the MGVA and MVVA filters, real data was collected at the National Center for Asphalt Technology's test track. A modified Infinity G35 test vehicle was driven around the track several times at a constant speed, and all the sensor measurements collected were then applied in post-process to the standard filter, the MGVA filter, and the MVVA filter. It was shown that, similar to what was seen in simulation, the standard filter was unable to accurately estimate the pitch and yaw of the vehicle. However, in contradiction to the simulation results, the filter was able to accurately estimate the roll of the vehicle. It was seen that this may have been due to the lateral excitation that the vehicle experienced while turning on the track at NCAT. In simulation, the trajectory of the vehicle was such that no lateral or longitudinal accelerations were experienced by the vehicle. However, at the NCAT test track, the avoidance of lateral excitation by the vehicle was impossible due to the turns of the track. As a result, it is likely that the lateral excitation experienced by the vehicle improved the system's ability estimate the roll angle. In correspondence to these results, the filter was also able to estimate the y accelerometer bias of the system well, however the x axis accelerometer bias estimate was not able to converge to the true bias.

Similar to what was seen in simulation, the MGVA filter was able to estimate the attitude and accelerometer bias of the system correctly. It was noted, however, that in order for the filter to function, centripetal acceleration had to be accounted for using gyroscope yaw rate measurements and vehicle velocity calculations obtained from vehicle CAN system measurements. In agreement with what was seen in simulation, the MVVA filter was also successful in estimating the attitude of the vehicle during it's test drive. This ultimately led to the filter's ability to estimate the true accelerometer bias acting on the system in all three axes.

Comparing filter attitude estimate errors, it was shown that the MGVA and MVVA filters were both able to achieve roll estimates on par with the standard filter. With regard to pitch, the MVVA filter produced estimate errors that were lower than the standard filter by several degrees. Additionally, the MGVA filter's pitch estimate errors were lower than the MVVA filter's estimates throughout the experiment. Finally, the MGVA filter produced low initial yaw estimate errors that grew large as the experiment progressed. Comparing the MVVA filter and the standard filter, it was seen that the standard filter at first produced lower yaw estimate errors than the MVVA filter. However, as the experiment progressed the standard filter's yaw estimate errors grew while the MVVA filter's yaw estimate errors remained bounded. Because of this, it was concluded that the MVVA filter, in this regard, provided an improvement to the standard filter, and demonstrated greater yaw estimation capabilities than the MGVA filter during this run.

Chapter 6

Utilization of the Heading Constraint for Ground Vehicle Applications

While Chapter 5 showed that the MGVA and MVVA filters were able to provide good results for the low excitation trajectories, it is important to compare them to other methods which have been used in the past to supplement the standard filter. Therefore, this chapter compares the MGVA and MVVA filters to a method tested by Ryan in [23] that uses velocity measurements, under known conditions of straight driving, to present the standard filter with a heading constraint for the vehicle. An exposition of this method will be given in the sections below, followed by several tests which compare the performance of the heading constraint method to that of the MGVA and MVVA filters.

6.1 Heading Constraint Formulation

For ground vehicle estimation, there are several distinctions that must be made with regard to vehicle heading. As Ryan explains in [23], a ground vehicle will have both a course, and a heading (or yaw angle). Heading is defined as the angular direction a vehicle is facing with respect to north, while course is defined as the angular direction of the vehicle's velocity vector with respect to north. The difference between these two angles is referred to as the side-slip that a vehicle experiences, and can be defined by Equations (6.1) and (6.2)

$$\beta = \nu - \psi \tag{6.1}$$

$$\beta = atan(\frac{V_y}{V_x}) \tag{6.2}$$

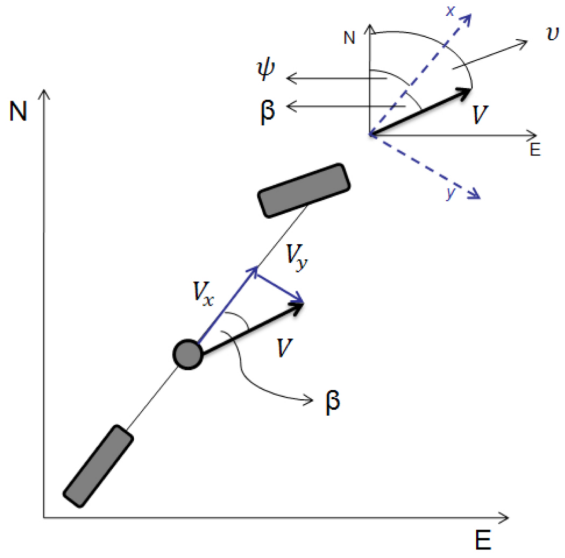

Figure 6.1: Vehicle heading, course and side-slip defined

where β is the side-slip angle, ν is course, ψ is the heading of the vehicle, and V_x and V_y represent x and y body frame velocities, respectively. For the case of a vehicle driving straight, with no lateral velocities, it can be seen that the expression in Equation (6.2) will go to zero. When this happens, the sideslip of the vehicle is zero, and ultimately the course of the vehicle ν is equal to the heading of the vehicle ψ.

Leveraging this information, it is possible to improve the performance of the standard loosely-coupled EKF by setting heading equal to course in the filter. This process is implemented by first calculating the course of the vehicle based on velocity measurements in the local navigation frame such that

$$\nu = atan(\frac{V_{east}}{V_{north}})^{GPS} \tag{6.3}$$

Next, the course calculation can be supplied to the filter as a measurement update, such that the z measurement vector of the filter becomes

$$z = \begin{bmatrix} \underline{P_{GPS}} & \underline{V_{GPS}} & \nu \end{bmatrix}^T \tag{6.4}$$

the H measurement matrix becomes

$$H = \begin{bmatrix} I_{33} & O_{33} & O_{31} & O_{31} & O_{31} & O_{33} & O_{33} \\ O_{33} & I_{33} & O_{31} & O_{31} & O_{31} & O_{33} & O_{33} \\ O_{13} & O_{13} & 0 & 0 & 1 & O_{13} & O_{13} \end{bmatrix} \tag{6.5}$$

and the R covariance matrix becomes

$$R = \begin{bmatrix} \sigma_P^2 & O_{33} & O_{31} \\ O_{33} & \sigma_V^2 & O_{31} \\ O_{13} & O_{13} & \sigma_\psi^2 \end{bmatrix} \tag{6.6}$$

This formulation leads to several important considerations that must be made. First, the filter may only apply the heading constraint when it holds true that the vehicle is not experiencing any side-slip. In order to account for this, Ryan configured the EKF he studied to only apply the constraint when the gyroscope of the system had measured very low amounts of z-axis rotation (less than 3 deg/s) over the course of some time period. It is also important to note that this assumptionof zeros sideslip does not always hold for straight driving. As a result, the use of the heading constraint can be a potential source of error in the estimation scheme. Second it should be noted that the filter used by Ryan in [23], as well as this formulation, rely upon measurement updates of the yaw angle between the local NED frame and the body frame. Because of this, the simplest loosely-coupled filter to utilize the heading constraint is one which operates in a local NED frame rather than an ECEF frame (as was done in Chapter 4). This is the type of filter that Ryan analyzed, and following from Ryan's work, it is the type of filter that the heading constraint was implemented on in this thesis. Next, it should be noted that in addition to the zero side-slip assumption, the constraint also assumes zero velocity in the down direction. Because of this, if it is so desired, the V_{down} component of the $\underline{V_{GPS}}$ vector in Equation (6.4) may be set equal to zero, further constraining the filter. Finally, Ryan showed that the heading constraint affects the

observability of the filter such that the heading and z-axis gyro bias are made observable. The benefit of heading and z-axis gyro bias observability should not be underestimated, since the yaw angle is the critical rotation angle for navigation.

6.2 Comparison of the MGVA and MVVA Filters to the Heading Constraint

In order to compare the MGVA and MVVA filters to the heading constrained filter, the loosely-coupled EKF implementing the heading constraint was evaluated using the simulated and experimental trajectories that the MGVA and MVVA filters were tested on in Chapter 5. Results from these tests can be seen in Figures 6.2-6.8.

Figures 6.2 and 6.3 show how the EKF with the heading constraint performed in simulation. Comparing these results with those from the standard filter in Chapter 5 (Figure 5.4) reveals that, with respect to the tilt angles, the heading constrained filter experiences drifting errors similar to the standard filter. However, it can be seen that contrary to the performance of the standard filter, the heading constrained filter achieves sub-degree levels of accuracy with it's yaw estimate.

Results from experimentally testing the EKF, with data collected at NCAT, can be seen in Figures 6.4-6.6. The EKF is able to estimate the roll with small amounts of drift effecting the algorithm throughout. The filter's estimate of pitch diverges from the true value almost immediately, and suffers from drifting errors through the run. However, the yaw estimate is successfully corrected by the heading constraint. As Figure 6.5 shows, as soon as the vehicle leaves the turn and begins driving straight, the heading constraint is activated, and the yaw estimate errors are reduced to the sub-degree level.

Figure 6.7 gives a comparison of attitude estimate errors between the heading constrained EKF, MGVA filter, and the MVVA filter. The MGVA and MVVA filters are able to obtain more accurate estimates of the tilt angles compared to the heading constrained filter. It is clear that the MGVA and MVVA filters are able to mitigate some of the difficulties seen in estimating these angles during the trajectory. In contrast, the heading constrained

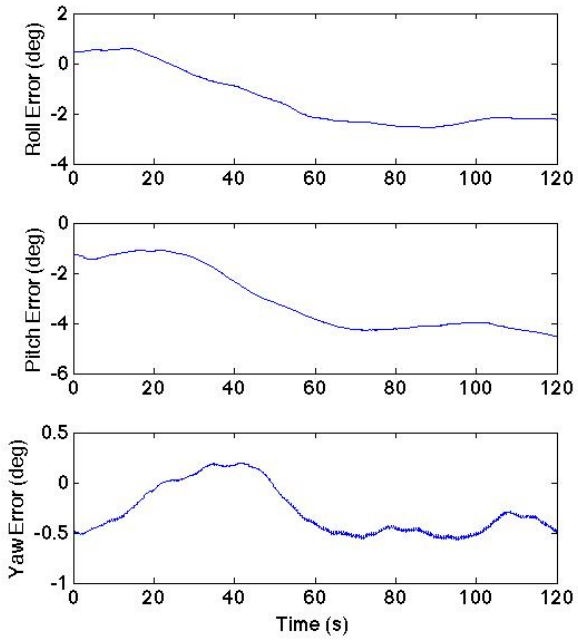

Figure 6.2: Attitude estimate error for the loosely-coupled EKF with the heading constraint over the course of the dynamic simulation

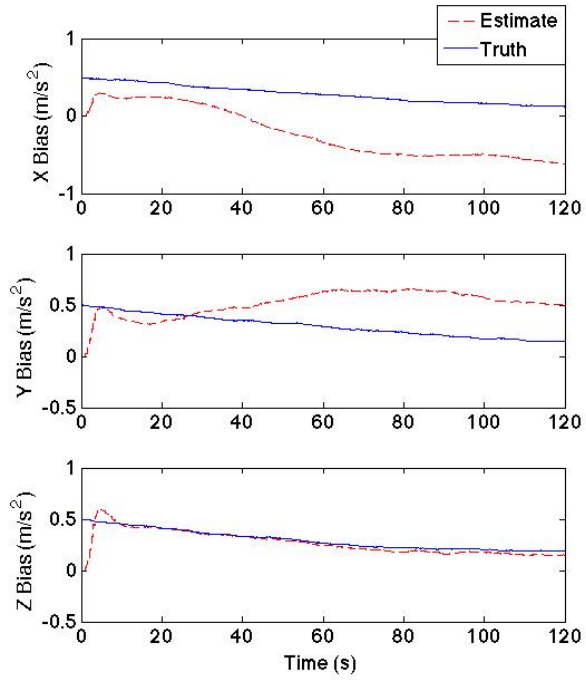

Figure 6.3: Accelerometer bias estimates vs. truth for the loosely-coupled EKF with the heading constraint over the course of the dynamic simulation

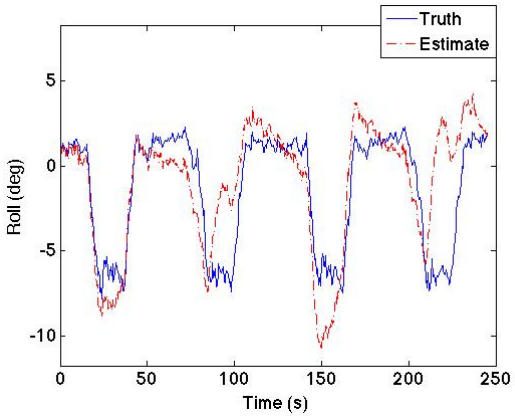

(a) Roll estimate vs. truth

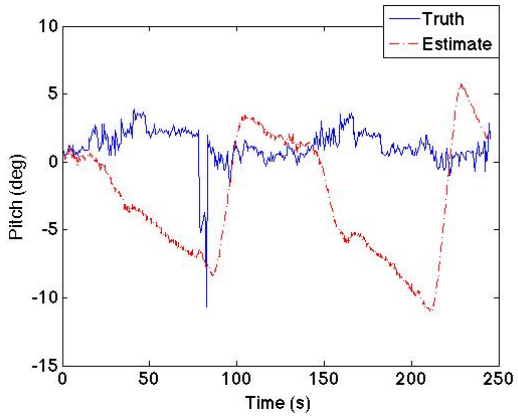

(b) Pitch estimate vs. truth

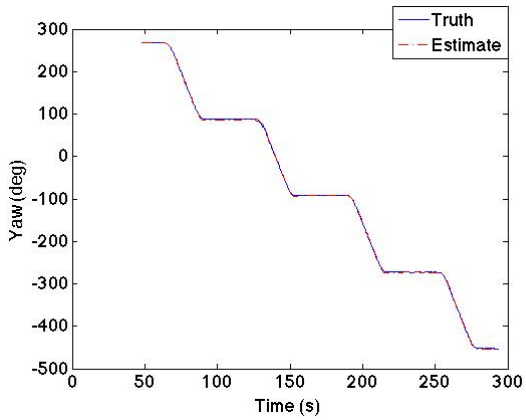

(c) Yaw estimate vs. truth

Figure 6.4: Attitude estimate vs. truth for the loosely-coupled EKF with heading constrained during dynamic trajectory at NCAT

filter is clearly able to provide a more accurate heading estimate during portions of straight driving.

These results lead to an important consideration. It is possible that the MGVA and MVVA filters may be improved using the heading constraint in their design rather than yaw updates from the attitude determination algorithm. In other words, it is possible that more accurate filters could be produced if the MGVA and MVVA filters used the attitude determination algorithm to update only roll and pitch during portions of straight driving. They could then supplement their yaw updates with information from a heading constraint, thus creating greater accuracy in heading estimation. Figure 6.8 depicts the effects of such an implementation. As can be seen, the MGVA filter's yaw estimate greatly benefits from the use

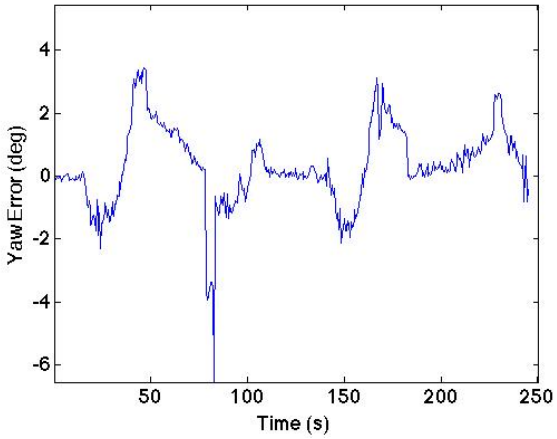

Figure 6.5: Yaw estimate error for the loosely-coupled EKF with heading constrained during dynamic trajectory at NCAT

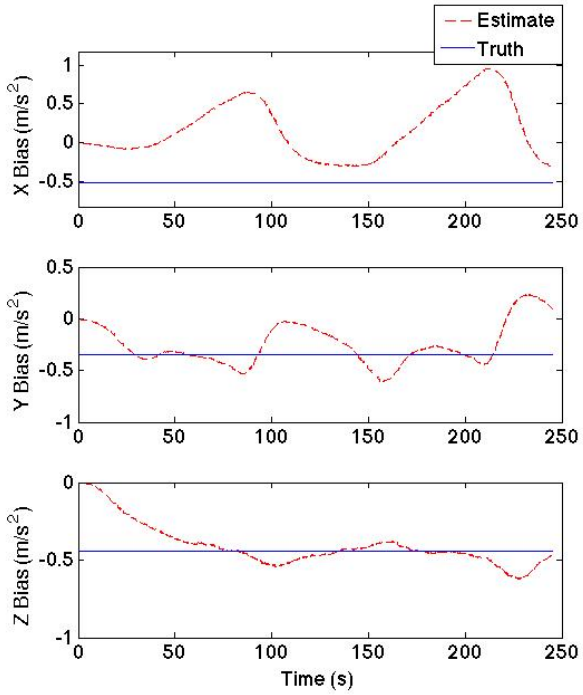

Figure 6.6: Accelerometer bias estimates vs. truth for the loosely-coupled EKF with heading constrained during dynamic trajectory at NCAT

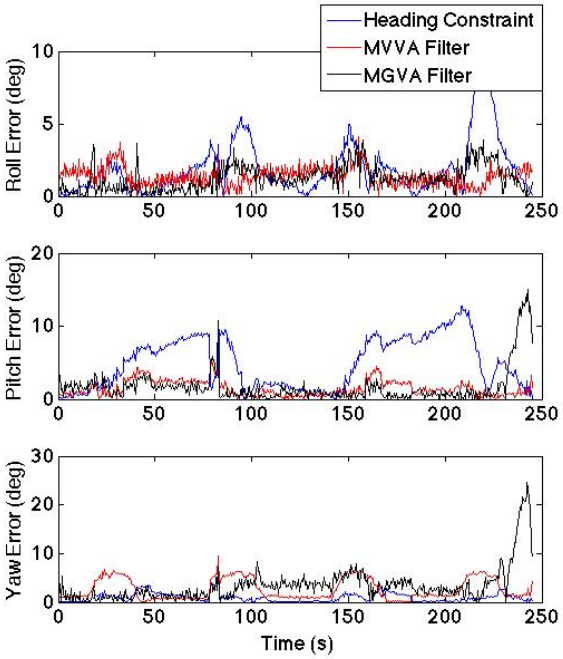

Figure 6.7: Attitude error comparison between the heading constrained EKF and the MVVA filter

of the heading constraint for portions of straight driving while testing. The use of the heading constraint results in several degrees of less error compared to the MGVA filter without the heading constraint. In contrast, the MVVA filter does not show improvement from the use of the heading constraint. Closer analysis of the yaw measurement updates coming from the attitude determination algorithm reveals that for portions of straight driving the attitude determination algorithm supplies the MVVA filter with a yaw updates that differ from the heading constraint update by a mean value of 0.07 degrees. As a result, it is evident that the accuracy of the heading constrained filter compared to the MVVA filter without the heading constraint is due first to the errors the MVVA filter experiences while turning and next to other aspects of the heading constrained filters architecture (e.g., its loosely-coupled implementation, or its use of just a heading update rather than a three Euler angle update).

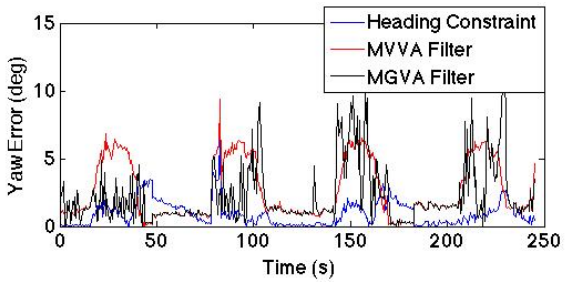

Figure 6.8: Yaw estimate error comparison between the heading constrained EKF, and the MGVA and MVVA filters using the heading constraint

6.3 Conclusion

In this chapter, a comparison of the MGVA and MVVA filters was made to a loosely-coupled filter with constrained heading that was studied by Ryan [23]. A presentation of this filter was given detailing how the filter utilizes velocity measurements to calculate the course of the vehicle, and then set the heading equal to the course during periods of straight driving. Having defined the heading constrained filter, it was then implemented on the data that the MGVA and MVVA filters were previously tested on in Chapter 5.

In simulation, the heading constrained filter performed similar to the standard filter in Chapter 5, suffering from drifting tilt angle estimates which led to drifting accelerometer bias estimates. In contrast to the standard filter however, it was seen that the heading constrained filter was able to maintain sub-degree levels of accuracy with respect to it's yaw estimate. Testing with experimental data revealed that the filter once again performed similar to the standard filter, obtaining reasonable roll estimates that suffered from some drifting errors, and poor pitch estimates which never converged to the true pitch. In contrast to this, the heading estimate of the filter was successfully corrected by the heading constraint, leading to an increased amount of accuracy in the yaw estimate.

Comparison of the heading constrained filter with the MGVA and MVVA filters showed that the MGVA and MVVA filters were able to provide better estimates of the tilt angles, while the heading constrained filter displayed better estimates of the yaw angle. It was concluded from these results that a better version of the MGVA and MVVA filters might be

possible if the MGVA and MVVA filters were to utilize heading constrained measurements during straight driving, rather than their own yaw measurement update. Implementation of such filters showed that the MGVA filter was able to benefit significantly from the use of the heading constraint during straight driving. In contrast, the MVVA filter did not benefit from the use of the heading constraint. Further analysis revealed that the updates provided by the heading constraint during straight driving were almost identical to those provided by the attitude determination algorithm in the MVVA filter. Because of this, it was concluded that the heading constrained filters success over the MVVA filter (without the heading constraint) was not simply due to the use of the heading constraint, but rather due other differences between the heading constrained architecture studied by Ryan [23] and the MVVA filter.

Chapter 7

Performance of the MGVA and MVVA Filters Under Dynamic Trajectories

Having shown the capabilities of the MGVA and MVVA filters under trajectories of low dynamics (straight and steady driving), the question arises as to how the filters will perform under normal vehicle operating conditions. In other words, the case of driving with either constant velocity, acceleration, or heading is fairly rare. It follows that if the MGVA and MVVA filters are to be considered truly useful, they cannot function only under these rare conditions but must function under any kind of vehicle trajectory. Because of this, several simulations and experiments were performed in order to validate the use of the MGVA and MVVA filters for a more dynamic trajectory. In the following sections, the results from these simulations and experiments will be presented, in order to show the usefulness of the MGVA and MVVA filters under more dynamic conditions.

7.1 Dynamic Trajectory Simulation Results

In order to test the effectiveness of the MGVA and MVVA filters, a simulated trajectory was created which consisted of several seconds of rapidly changes linear acceleration, followed by several seconds of no acceleration. The simulated trajectory consisted of straight driving without any changes in roll or pitch, for simplicity. This type of trajectory represents what a ground vehicle would experience during rapid acceleration and braking, followed by driving at constant velocity, during straight and level driving. Both the MGVA filter and the MVVA filter were tested with this trajectory, and results from these simulations are shown below. Figure 7.1 displays the true accelerations and yaw rates of the simulated vehicle. As can be seen the vehicle experiences rapidly changing accelerations, in the longitudinal direction, for several seconds beginning around the 15 second mark, followed by a period of no acceleration.

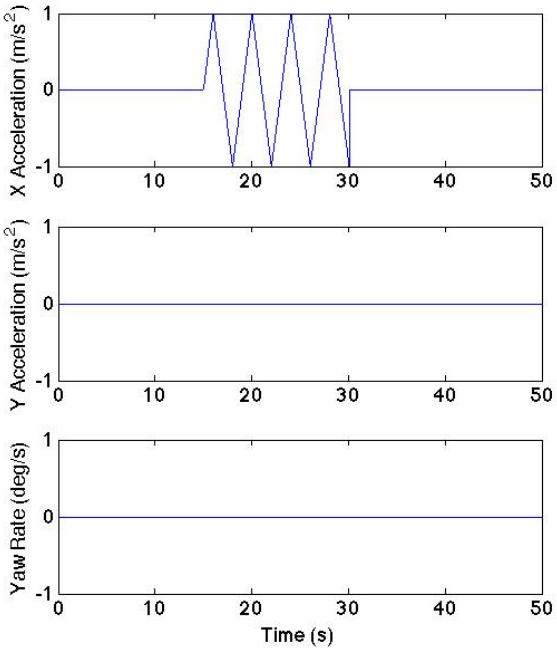

Figure 7.1: Simulated ground vehicle accelerations and yaw rates

Figures 7.2 and 7.3 display how the MGVA filter performed during the simulation. As can be seen, the variations in acceleration had a large effect on the filter's attitude estimate. The roll, pitch, and yaw estimates all follow trends which appear to come from the variation in acceleration shown in Figure 7.1. The result of these variations is that all three euler angles experience errors on the order of several degrees during acceleration changes. It is possible, as has been shown before in this thesis, the MGVA filter may not be suitable for trajectories which experience increased dynamics. The dynamics experienced by the system make the accelerometer's measurements a function of both the gravity vector and vehicle accelerations, and ultimately corrupt the attitude calculations made by the attitude determination algorithm.

In spite of the poor attitude performance during the initial portion of the trajectory, Figure 7.3 shows that the filter is still able to observe the accelerometer bias. This is simply due to the constant velocity the system experiences after rapid acceleration and deceleration. Taking a closer look at Figure 7.3 shows that the x and y accelerometer bias estimates suffer

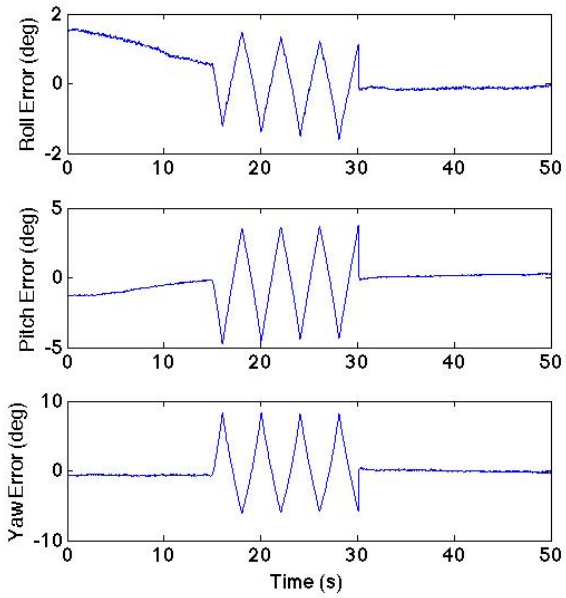

Figure 7.2: Attitude estimate error for the MGVA filter over the course of the dynamic simulation

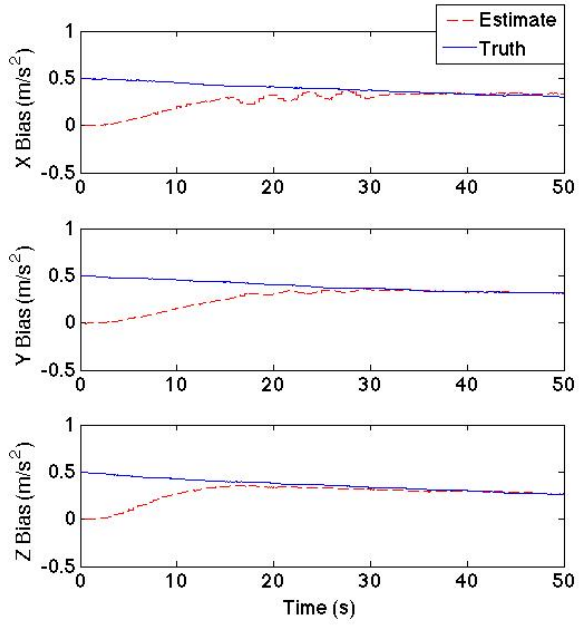

Figure 7.3: Accelerometer bias estimates vs. truth for the MGVA filter over the course of the dynamic simulation

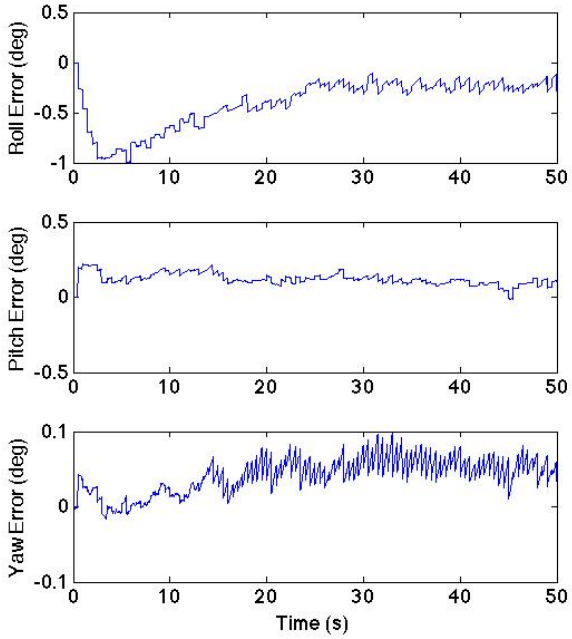

Figure 7.4: Attitude estimate error for the MVVA filter over the course of the dynamic simulation

from a slight oscillation right before convergence, due to the corrupt attitude estimates. It is evident that oscillations in attitude estimates from the MGVA filter ultimately lead to oscillating accelerometer bias estimates. If the attitude estimates oscillate about the true attitude of the system, then the bias estimates will also oscillate about the true bias effecting the system. However, if the oscillations are large enough, the filter will still experience error due to accelerometer bias. Because of this, it is evident that accelerations experienced by the vehicle ultimately lead to a degradation in performance of the MGVA filter.

Figures 7.4 and 7.5 show the attitude estimate error and accelerometer bias estimates for the MVVA filter. In Figure 7.4, despite the extreme changes in vehicle acceleration, the MVVA filter is still able to achieve sub-degree levels of accuracy with all three euler angles. This is because accurate measurements of the vehicle's velocity allow the attitude determination algorithm to form an accurate estimate of the vehicle attitude, and supply the filter with an update. In Figure 7.4 it can be seen that this ultimately leads to accurate estimation of accelerometer bias. Therefore, this demonstrates that even during increased

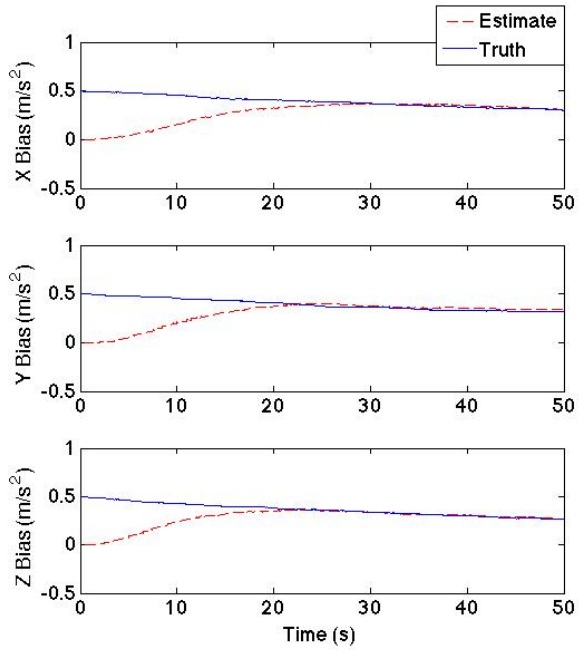

Figure 7.5: Accelerometer bias estimates vs. truth for the MVVA filter over the course of the dynamic simulation

vehicle dynamics, the MVVA filter is still able to perform well, estimating both the vehicle attitude and accelerometer biases accurately.

7.2 Dynamic Trajectory Experimental Results

Having tested the filters on a simulated dynamic trajectory, the results seen in simulation were further validated through experimentation with real data. The same experimental setup used in Chapter 5 was used once again to collect data for two different dynamic trajectories. The first trajectory involved periodically accelerating and deccelerating the vehicle as it drove down the track. This resulted in in varying linear accelerations of the vehicle as it moved down the track. The second trajectory involved the performance of double lane change maneuvers at high speeds on the straight portions of the track. Results from implementing the MGVA and MVVA filters on the two trajectories are given in the following sections.

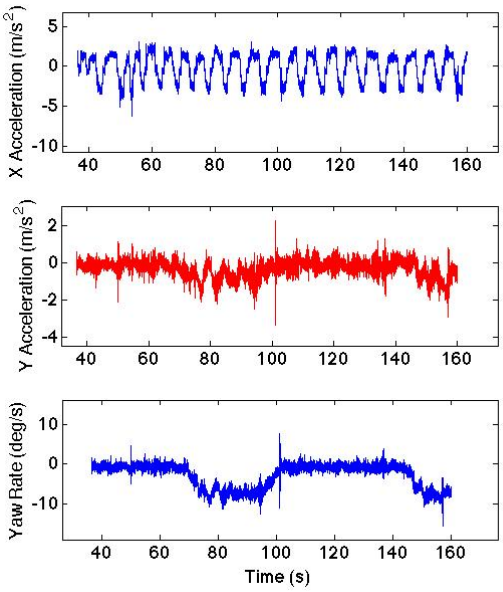

Figure 7.6: Acceleration and yaw rate of the vehicle for the first dynamic trajectory

7.2.1 Dynamic Trajectory with Changes in Linear Acceleration

Figure 7.6 shows the accelerations and yaw rates experienced by the vehicle during the first trajectory. The accelerations in the x-axis of the vehicle change constantly, while the y-axis accelerations of the vehicle experience little acceleration unless the vehicle is turning. During the turns, the y-axis of the vehicle experiences varying amounts of acceleration. With regard to the yaw rate, the vehicle only experiences changes in yaw on the turns of the track; otherwise the vehicle was driven straight.

Figures 7.7 and 7.8 depict the attitude estimate and the accelerometer bias estimates from the MGVA filter during the test. Figure 7.7 shows that during the experiment the MGVA filter struggles to give an accurate estimate of the attitude, producing oscillating estimates, similar to what was seen in simulation. Of the three Euler angle estimates, the filter shows the largest success in roll estimation. Additionally, a comparison of Figures 7.7 and 7.2 reveals that with regard to the tilt angles, longitudinal changes in acceleration cause larger errors in pitch estimation than roll estimation. It is possible that this is the case because, longitudinal changes in the vehicle's acceleration cause the vector measured by the

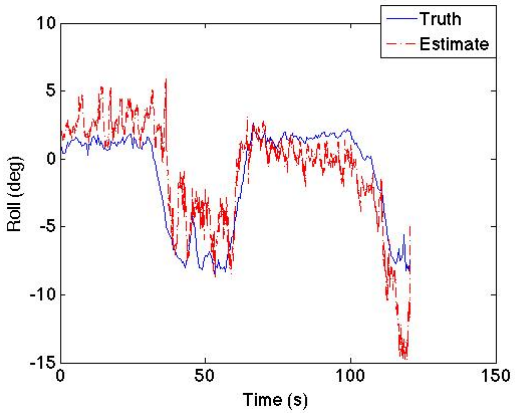

(a) Roll estimate vs. truth

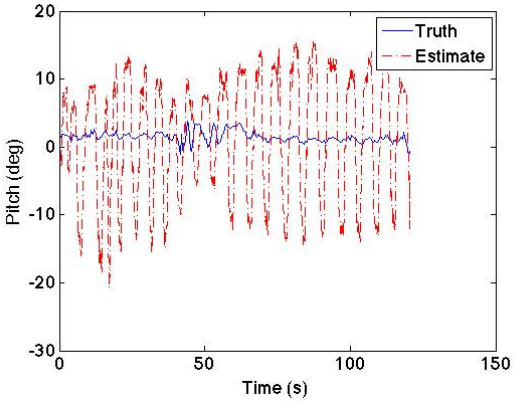

(b) Pitch estimate vs. truth

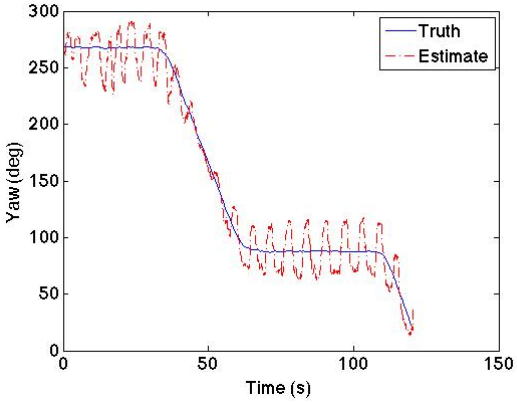

(c) Yaw estimate vs. truth

Figure 7.7: Attitude estimate vs. truth for the MGVA filter during dynamic trajectory at NCAT

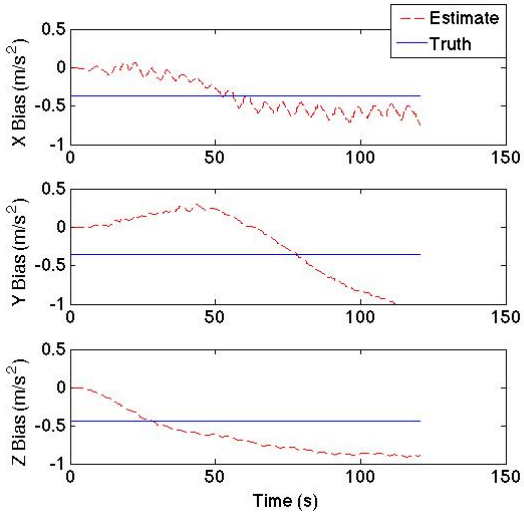

Figure 7.8: Accelerometer bias estimates vs. truth for the MGVA filter over the course of the dynamic trajectory at NCAT

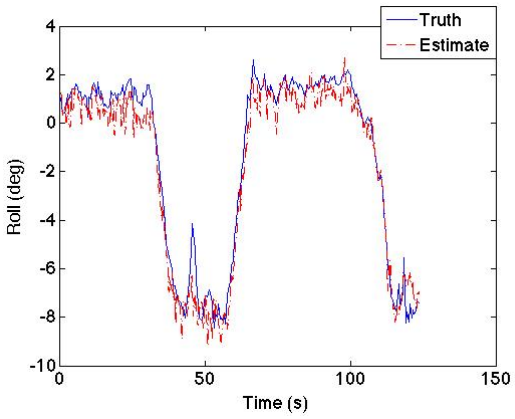

(a) Roll estimate vs. truth

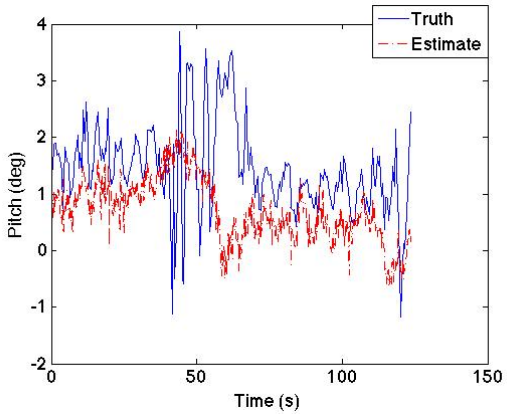

(b) Pitch estimate vs. truth

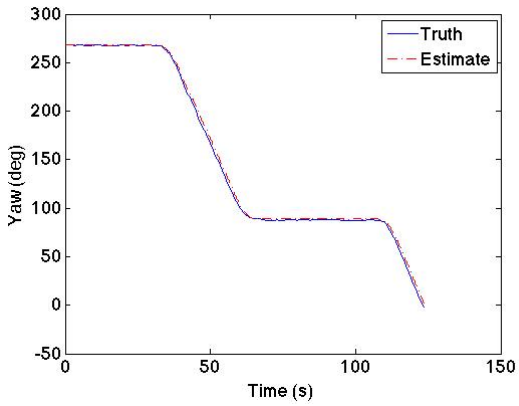

(c) Yaw estimate vs. truth

Figure 7.9: Attitude estimate vs. truth for the MVVA filter during dynamic trajectory at NCAT

accelerometer to pitch away from the vector that would be measured by a static sensor. This ultimately leads the attitude determination algorithm to believe that the vehicle is pitching when in fact the orientation of the vehicle is not changing. As a result, it may be that longitudinal changes in acceleration cause the MGVA filter's largest tilt angle estimation errors to occur in pitch estimation. Figure 7.8 shows that these attitude errors during the trajectory ultimately lead to errors in the filter's accelerometer bias estimates. In the y axis, the estimate experiences small oscillations, and ultimately diverges from the reference value. In the x axis, the estimate experiences even larger oscillations, however it is able to oscillate near the reference.

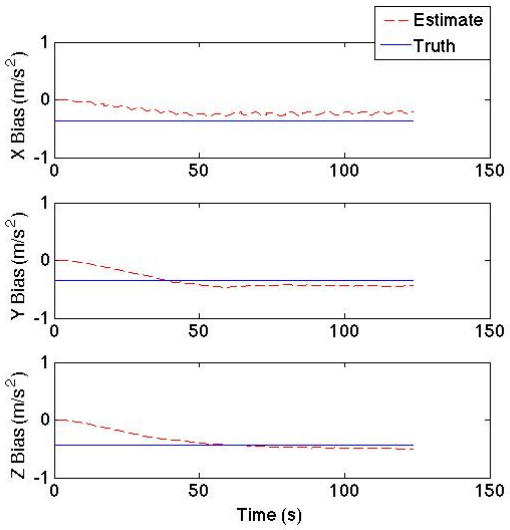

Figure 7.10: Accelerometer bias estimates vs. truth for the MVVA filter over the course of the dynamic trajectory at NCAT

In contrast, Figures 7.9 and 7.10 shows that the MVVA filter has much more success in its estimation of the attidue and accelerometer biases. Figure 7.9 shows that with increased dynamics, the system is actually able to produce a more accurate attitude estimate than what was seen during the low dynamics trajectory (see Figure 5.18). Closer examination of the data reveals that during these tests, the MVVA filter achieves sub-degree levels of accuracy for all three euler angles of the vehicle. Finally, examining Figure 7.10 shows that once again, the ability to observe the attitude of the vehicle leads to the MVVA filter's ability to obtain valid accelerometer bias estimates in both the x and y axes.

7.2.2 Dynamic Trajectory with Double Lane Change Maneuver

The second dynamic trajectory that the filters were tested on involved the vehicle performing double lane change maneuvers. The double lane change maneuvers consisted of the vehicle merging from one lane into an adjacent lane for a prolonged period of time, and then merging back into the original lane. During each straight portion of the track, two double lane change maneuvers were performed. The acceleration and yaw rate of the vehicle for this trajectory are shown in Figure 7.11. In the y axis, the vehicle experiences varying

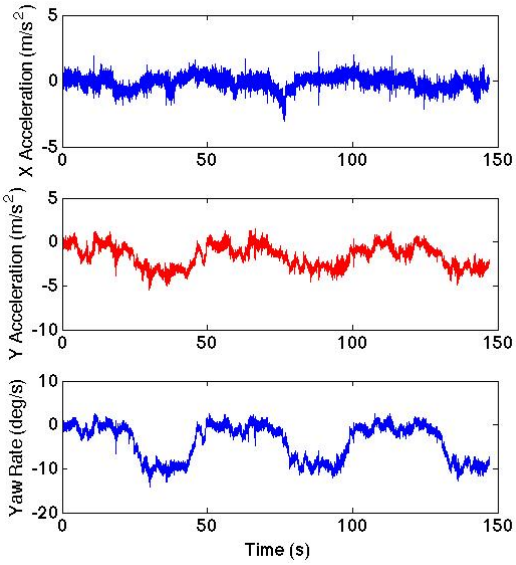

Figure 7.11: Acceleration and yaw rate of the vehicle for the double lane change maneuver

lateral acceleration during merging and turning. Additionally, the x axis of the vehicle also experiences varying accelerations. Specifically, varying acceleration occurs due to moments of breaking and acceleration while turning the vehicle.

Figures 7.12 and 7.13 depict how the MGVA filter handled the attitude and accelerometer bias estimation during the double lane change maneuvers. As can be seen, the filter has some success in estimating all three rotational angles, however, it struggles once again from oscillations in the estimates due to changing accelerations while merging and during acceleration/braking. Oscillations in the estimates can be most clearly recognized in the roll estimates during lane changing, and the effects of acceleration/braking can be seen in the two large dips in the pitch estimate during turning. The effects of the oscillating attitude estimates can also be seen in Figure 7.13. Once again, the filter is able to estimate the accelerometer bias effecting the system, however its estimates suffer from oscillations just as the attitude estimates do.

Figures 7.14 and 7.15 display results from implementing the MVVA filter on the trajectory. The MVVA filter is once again successful in its attitude estimation. It can be seen that the filter is able to provide good estimation of the roll of the vehicle during lane changes.

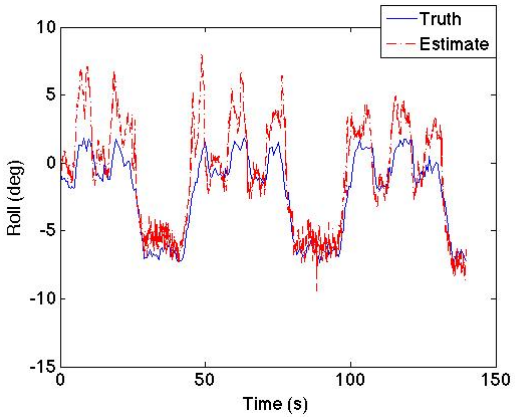

(a) Roll estimate vs. truth

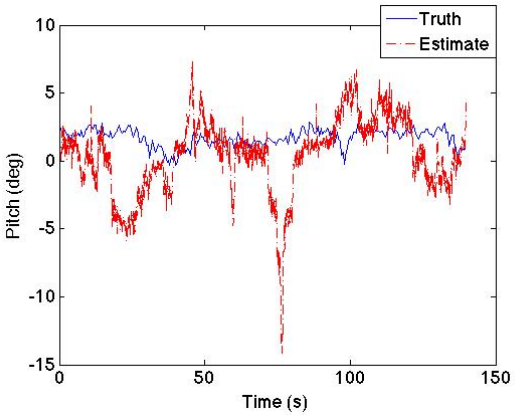

(b) Pitch estimate vs. truth

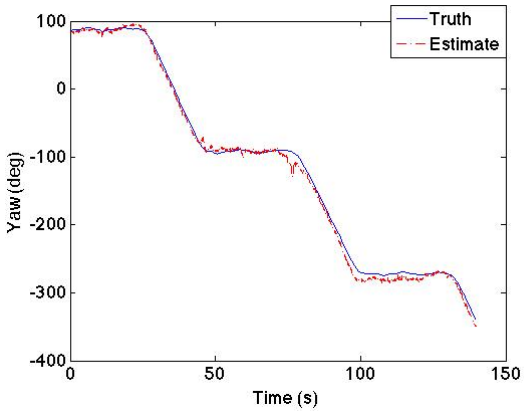

(c) Yaw estimate vs. truth

Figure 7.12: Attitude estimate vs. truth for the MGVA filter during double lane change maneuver

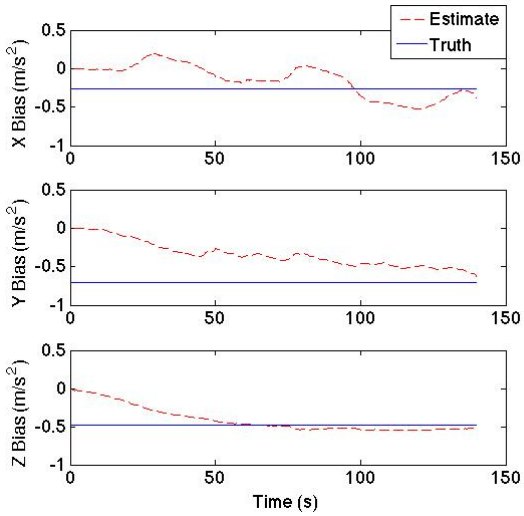

Figure 7.13: Accelerometer bias estimates vs. truth for the MVVA filter over the course of the double lane change maneuver

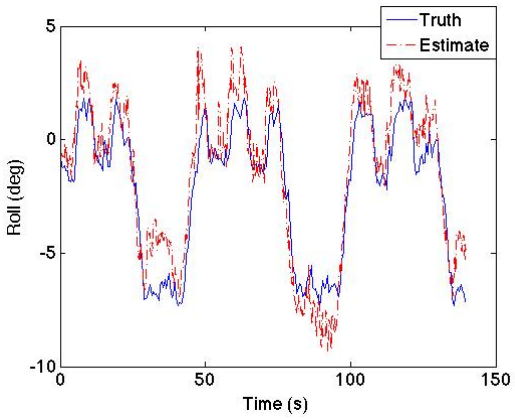

(a) Roll estimate vs. truth

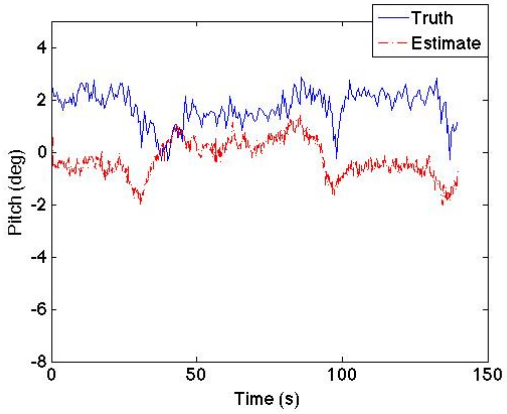

(b) Pitch estimate vs. truth

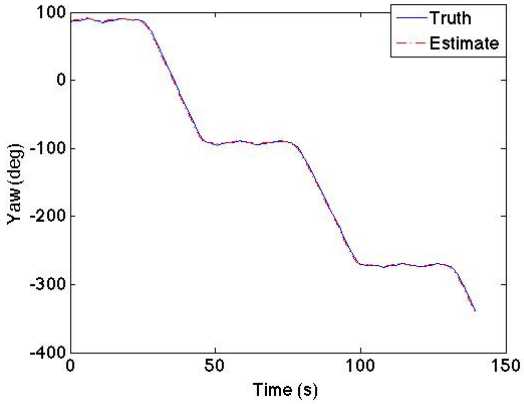

(c) Yaw estimate vs. truth

Figure 7.14: Attitude estimate vs. truth for the MVVA filter during double lane change maneuver

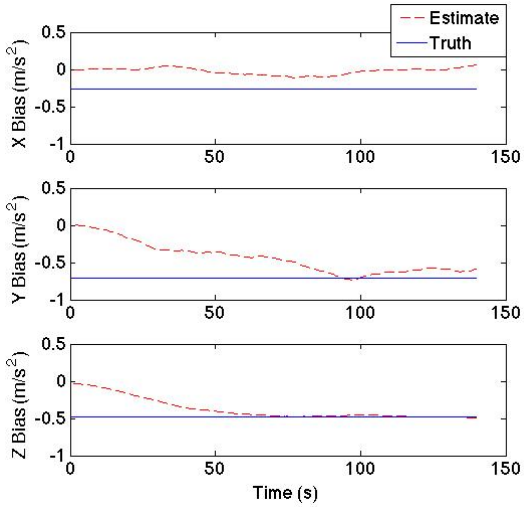

Figure 7.15: Accelerometer bias estimates vs. truth for the MVVA filter over the course of the double lane change maneuver

It can also be seen that the filter produces a pitch estimate that follows the trend of the reference, however is in error throughout by roughly two degrees. It is possible that this two degree offset is due to misalignment errors that were not noticed during testing and data collection. If the sensor was slightly pitched downwards from the vehicle frame by two degrees, then it may account for the errors seen in Figure 7.14. Finally, Figure 7.15 shows that the attitude estimation the filter provides ultimately leads to bias estimates which are stable, and converge near the reference accelerometer bias.

7.3 Conclusion

In order to further validate the MGVA and MVVA filters, the algorithms were tested in simulation and experimentation on a more dynamic trajectory. The filter was tested on a simulated trajectory which contained several seconds of longitudinal acceleration and decceleration, followed by several seconds of no acceleration. It was shown that under dynamic portions of the trajectory, the MGVA filter was unable to obtain a valid estimate of the vehicle attitude. It was not until the simulated vehicle stopped accelerating and deccelerating that the filter was able to accurately estimate the attitude. In spite of the initial attitutde estimate errors, the filter was still able to produce accurate estimates of the accelerometer bias, largely due to the filter's accurate attitude estimate that followed the period of vehicle accleration and decceleration. In contrast, the increased dynamics did not hinder the performance of the MVVA filter. The filter was still able to calculate accurate attitude measurements during the test and update the filter with valuable information. This in turn led to accurate estimates of the accelerometer biases corrupting the filter inputs.

The MGVA and MVVA filters were then tested on two experimental sets of data collected at NCAT. During the first experimental run, the vehicle was driven around the track periodically accelerating and decellerating. For this experimental run, the MGVA filter suffered from oscillating attitude estimates as a result of the changing linear accelerations. It was noted that, with regard to tilt angle estimation, the filter suffered the most in pitch

estimation. It was also noted that these results may imply a correlation between linear acceleration and pitch estimation error in the MGVA filter. Corresponding to these results, the filter suffered from oscillating accelerometer bias estimates, as well as a diverging y axis bias estimate. The conclusion from this experiment was that the MGVA filter may only be suitable for use during periods of steady driving. In contrast the MVVA filter performed well during the experimental run. The MVVA filter was able to produce sub-degree accurate attitude estimates during portions of the run, and ultimately was able to produce accurate estimates of the accelerometer biases.

During the second experimental run, the vehicle performed several double lane change maneuvers as it traveled down the track. For this trajectory, it was seen that the MGVA filter once again suffered from oscillations in the attitude estimates due to changing accelerations on the vehicle during lane changing and turning. Corresponding to the oscillations in attitude estimates, the filter accelerometer bias estimates also suffered from oscillations. However, these estimates were still able to draw near to the reference. With regard to the MVVA filter, the filter was able to provide good roll estimation, and a pitch estimate which followed the trend of the pitch reference. However, it was noted that the pitch estimate did suffer from a small offset which effected the estimate throughout the experiment. It was suggested that the offset that was seen may have been due to small misalignment errors that were not detected during experimentation. Finally, it was seen that the MVVA filter was able to provide accelerometer bias estimates that were stable and converged near the reference values.

Chapter 8

Conclusions and Future Work

8.1 Conclusion

This thesis has shown that magnetometers are useful sensors, with large opportunities for use in the navigation industry. Despite the sensors nature to be prone to error, it has been shown that proper correction can yield high value information for navigation systems. The versatility of these sensors affords them the potential to provide many different kinds of navigation information. Specifically, this thesis has shown that magnetometers can be applied to standard algorithms to provide increased accuracy and observability for ground vehicle navigation filters.

There has been an extensive amount of research done in the areas of sensor fusion and GPS/INS enhancement. In this thesis, it was chosen to explore how magnetometers could be used to enhance the observability and the estimation capabilities of a standard closely-coupled EKF. As a result, this thesis contributes the following. First, it contributes an observability analysis of the MGVA and MVVA filters, and shows that for low excitation ground vehicle trajectories, the two filters are able to increase the observability to full rank. Next, this thesis provides an analysis of the attitue and accelerometer bias estimation capabilities of the two filters, for low-excitation ground vehicle trajectories. This thesis, also contributes a comparison of the two modified filters, to a heading constraint method studied by another author in the past [23]. Finally, this thesis contributes a performance analysis of the MGVA and MVVA filters on a dynamic ground vehicle trajectory.

At the start of this thesis, it was shown that magnetometers struggle from different error sources. Among these the most prominent error sources are due to magnetic interference that is measured by the magnetometer. In this thesis, a two-part least squares estimator was

chosen for magnetometer calibration. The architecture of the chosen algorithm was explained in detail, and the algorithm was tested on both simulated and experimentally collected magnetometer data. The simulations and experiments that were performed showed that the algorithm was successful in accounting for various error sources and ultimately correcting heading calculations from magnetometers.

It was also shown that magnetometers are highly useful sensors for euler angle attitude determination. From among the available literature, a Newton-Gauss based estimator was chosen that was able to calculate the angular rotation between two vector sets in different coordinate frames. This algorithm was explained in detail, and then tested in simulation on an aerial trajectory of interest. The algorithm was first tested using magnetometer measurements and accelerometer measurements. It was shown that the accelerometer-magnetometer pair did not work well on trajectories with high dynamics, largely due to a clouding of the gravity vector in the accelerometer measurements taken. The algorithm was then tested magnetometer measurements and velocity vector measurements. These tests showed that under high dynamics the velocity-vector and magnetometer pair are much more able to determine an accurate attitude solution.

This thesis has also delivered an exposition of the standard loosely-coupled EKF, the standard closely-coupled EKF, and two modified EKFs, the MGVA filter and the MVVA filter. Following this exposition, the stadard closely-coupled EKF, the MGVA filter, and the MVVA filter were both tested in simulation and experimentation on low dynamics trajectories. Results in simulation showed that the standard filter was unable to provide accurate accelerometer bias estimates due to incorrect attitude estimates from the filter. In contrast, the MGVA and MVVA filters were able to provide much more accurate estimates of the vehicle attitude, which led to good estimation of the accelerometer bias. Experimentally, the standard EKF avoided some of the drifting roll estimates found in simulation. It is likely that this was the case due to the lateral changes in acceleration that the vehicle experienced while turning. In contrast, the filter experienced drifting pitch and yaw estimate errors due

to the unobservable trajectory. All of this led to poor estimation of the x-axis accelerometer bias, but good estimation of the y and z axes accelerometer biases. Along the same experimental trajectory, the MGVA and MVVA filters were both successful in estimating the attitude and accelerometer bias during the run. Comparison of the MGVA and MVVA filters with the standard filter showed that the MGVA and MVVA filters were able to achieve better pitch accuracy and roll estimates on par with the standard filter. Additionally, it was seen that the MGVA filter provided higher pitch accuracy than the MVVA filter throughout the experiment. With regard to yaw, the MGVA filter initially produced lower errors than the standard filter and MVVA filter, however as the experiment progressed the yaw estimate error from the MGVA filter grew larger than both the standard filter and the MVVA filter. Comparing the MVVA filter to the standard filter, it was seen that initially the standard filter produced higher accuracy than the MVVA filter. However, toward the end of the experiment it was seen that bounded errors in the MVVA filter's estimate allowed the filter to achieve higher accuracy than the standard filter.

The MGVA and MVVA filters were also compared to a method studied by Ryan in [23] and termed the heading constraint. The heading constraint method was explained in detail, and then tested on both the simulated and experimental trajectories that the MGVA and MVVA filters was tested on in Chapter 5. In simulation, it was shown that the heading constraint was successful in reducing the yaw estimate errors to the sub-degree level. Experimentally, the filter was able to perform the same, reducing yaw errors during portions of straight driving. Comparison of the heading constraint to the MGVA and MVVA filters showed that the MGVA and MVVA filters were superior in providing estimates of the roll and pitch of the vehicle. However, the heading constrained filter was able to provide more accurate estimates of the yaw angle. This led to the conclusion that the MGVA and MVVA filters could be improved upon if they were to utilize a heading constraint during portions of straight driving rather than using their own yaw estimates. Implementation of the heading constraint on the MGVA and MVVA filters revealed that the heading constraint was

successful in improving heading estimation on the MGVA filter, but not on the MVVA filter. Further analysis showed that the yaw updates provided to the MVVA filter from the attitude determination algorithm were almost identical to those that the heading constraint could provide. This led to the conclusion that the increased accuracy of the heading constrained filter compared to the MVVA filter (without the heading constraint) was due in part to errors the MVVA filter encountered while turning, and in part to other aspects of the filter architectures not directly related to the use of the heading constraint.

Finally, in order to validate the MGVA and MVVA filters for typical driving scenarios, the two filters were tested on two trajectories with much higher dynamics. Results from the first trajectory showed that, under increased dynamics the MGVA filter struggled due to oscillating attitude estimates that ultimately led to poor bias estimation. This led to the conclusion that the MGVA filter may only be suitable for use during steady driving. As expected, the MVVA filter provided better estimates of the attitude due to its use of the vehicle velocity vector rather than the acceleration due to gravity. During experimental testing with the MVVA filter, it was shown that sub-degree attitude estimates were achieved during the run. This ultimately also led to good estimation of the accelerometer biases corrupting the system. Results from the second trajectory also showed that under increased dynamics the MGVA filter struggled from oscillating attitude and accelerometer bias estimates. In contrast, the MVVA filter was able to provide good estimation without oscillations. However, it was noted that the MVVA filter's pitch estimate suffered from a slight offset that effected it throughout the experiment. It was concluded that this may have been due to misalignment errors that were not detected during experimentation.

8.2 Future Work

Moving forward from this thesis, there are several questions that have been raised that remain to be answered. With regard to magnetometer calibration and data collection, it has been discussed that if a magnetometer does not rotate about the vertical axis, in a

plane other than the horizontal, the calibration algorithm may not perform optimally. This presents an issue for ground vehicles, since it is difficult to perform 360 degree rotations outside of the horizontal plane while in a ground vehicle. The sensor cannot be rotated apart from the vehicle, because then it would not detect the errors that corrupt it as a result of being in the vehicle. Additionally, a set of data in the horizontal plane cannot simply be rotated out of the plane in post-processing, and then used for calibration. This would not be a true representation of the errors corrupting the sensor; it would only rotate the errors presented in the horizontal plane, rather than measuring the errors that corrupt the magnetometer as it is rotated. In addition to this, only post-process magnetometer calibration was utilized in this thesis, however, several algorithms exist that are capable of estimating calibration parameters in real-time. These would prove beneficial due to their ability to modify calibration parameters during experimentation. This ability could potentially lead to more accurate calibrations tailored to each data set. None of these methods were explored in this thesis, however, and are therefore left as an endeavor for future work.

Although sub-degree attitude accuracy was achieved during some of the experimental testing performed in this thesis, the majority of testing done with both the MGVA and MVVA filters did not yield this type of accuracy. There are several parts of the current algorithm which could be modified in order to make the algorithm more accurate. One area of interest that has already been mentioned is in magnetometer calibration. It is no surprise that a more accurate calibration of the magnetometer could lead to better results on a consistent basis. In addition to this, this thesis has not explored the use of more than 2 vector sensors for enhanced attitude determination. The use of sun sensors, star sensors, and other vector measurement devices could potentially increase the accuracy of the attitude solutions output by the attitude determination algorithm. Using more than two sensors could also lead to specific attitude solutions on a consistent basis, thus removing any error spikes in the solution whenever the algorithm cannot determine a specific solution (see Figure 3.4 and Section 3.3).

Finally, examining the architecture of the MGVA filter, it is significant to note that the attitude estimation capabilities of the filter are independent of GPS measurements. This is a useful quality of the MGVA filter because, it gives the filter potential to improve navigation performance during GPS signal outages. During a GPS outage, the MGVA filter theoretically has the capability to continue attitude and accelerometer bias estimation, and ultimately improve vehicle navigation even when GPS is unavailable. This capability has also been left as future work, however, and was not explored in this thesis.

There are many opportunities for forward movement in the disciplines of sensor fusion and attitude determination. It has been the goal of this thesis to provide a thorough examination of one way in which magnetic sensors can be used in the pursuit of this forward motion. As can be seen, there are still areas of the subjects explored in this thesis that remain to be pioneered. But it is certain that the momentum of the navigation industry can only bring about better and brighter things from the future.

Bibliography

[1] Groves, Paul. Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems. Artech House, 2013.

[2] Eisenbeiss, Henri. A Mini Unmanned Aerial Vehicle (UAV): System Overview and Image Acquisition. International Archives of Photogrammetry. Remote Sensing and Spatial Information Sciences 36.5/W1 (2004).

[3] Bevly, David M. Global Positioning System (GPS): A Low-Cost Velocity Sensor For Correcting Inertial Sensor Errors on Ground Vehicles. Journal of dynamic systems, measurement, and control 126.2 (2004): 255-264.

[4] Yun, X., et al. Testing and Evaluation of an Integrated GPS/INS System for Small AUV Navigation. Oceanic Engineering, IEEE Journal of 24.3 (1999): 396-404.

[5] Misra, Pratap, and Per Enge. Global Positioning System: Signals, Measurements and Performance Second Edition. Lincoln, MA: Ganga-Jamuna Press, 2006.

[6] Lo, Sherman, et al. Alternative Position Navigation & Timing (APNT) Based on Existing DME and UAT Ground Signals. Proceedings of the Institute of Navigation GNSS Conference, Portland, OR. 2011.

[7] Walter, Debbie, et al. Road Navigation Using Multiple Dissimilar Environmental Features to Bridge GNSS Outages. (2015): 2098-2114.

[8] Evennou, Frdric, and Franois Marx. Advanced Integration of WiFi and Inertial Navigation Systems for Indoor Mobile Positioning. Eurasip journal on applied signal processing 2006 (2006): 164-164.

[9] Wenzel, Thomas A., et al. Dual Extended Kalman Filter for Vehicle State and Parameter Estimation. Vehicle System Dynamics 44.2 (2006): 153-171.

[10] Nistr, David, Oleg Naroditsky, and James Bergen. Visual Odometry for Ground Vehicle Applications. Journal of Field Robotics 23.1 (2006): 3-20.

[11] Caruso, Michael J. Applications of Magnetoresistive Sensors in Navigation Systems. SAE transactions 106 (1997): 1092-1098.

[12] Liu, Sheng-Wu, Zhao-nian Zhang, and James C. Hung. A High Accuracy Magnetic Heading System Composed of Fluxgate Magnetometers and a Microcomputer. Aerospace and Electronics Conference, 1989. NAECON 1989., Proceedings of the IEEE 1989 National. IEEE, 1989.

[13] Psiaki, Mark L., Francois Martel, and Parimal K. Pal. Three-Axis Attitude Determination Via Kalman Filtering of Magnetometer Data. Journal of Guidance, Control, and Dynamics 13.3 (1990): 506-514.

[14] Marins, Joo Lus, et al. An Extended Kalman Filter for Quaternion-Based Orientation Estimation Using MARG Sensors. Intelligent Robots and Systems, 2001. Proceedings. 2001 IEEE/RSJ International Conference on. Vol. 4. IEEE, 2001.

[15] Hagan, Martin T., et al. Neural Network Design. Vol. 20. Boston: PWS publishing company, 1996.

[16] El-Diasty, Mohammed, and Spiros Pagiatakis. Calibration and Stochastic Modelling of Inertial Navigation Sensor Errors. Journal of Global Positioning Systems 7.2 (2008): 170-182.

[17] Syed, Z. F., et al. A New Multi-Position Calibration Method for MEMS Inertial Navigation Systems. Measurement Science and Technology 18.7 (2007): 1897.

[18] Goshen-Meskin, D., and Bar-Itzhack, I., 1992, Observability Analysis of Piece-Wise Constant Systems. I. Theory, IEEE Trans. Aerosp. Electron. Syst., 28(4), pp. 10561067.

[19] Goshen-Meskin, D., and Bar-Itzhack, I., 1992, Observability Analysis of Piece-Wise Constant Systems. II. Application to Inertial Navigation In-Flight Alignment [Military Applications], IEEE Trans. Aerosp. Electron. Syst., 28(4), pp. 10681075.

[20] Hong, S., Lee, M. H., Chun, H.-H., Kwon, S.-H., and Speyer, J., 2005, Observability of Error States in GPS/INS Integration, IEEE Trans. Veh. Technol., 54(2), pp. 731743.

[21] Rhee, I., Abdel-Hafez, M. F., and Speyer, J. L., 2004, Observability of an Integrated GPS/INS During Maneuvers, IEEE Trans. Aerosp. Electron. Syst., 40(2), pp. 526535.

[22] Ryan, Jonathan G., and David M. Bevly. On the Observability of Loosely Coupled Global Positioning System/Inertial Navigation System Integrations With Five Degree of Freedom and Four Degree of Freedom Inertial Measurement Units. Journal of Dynamic Systems, Measurement, and Control 136.2 (2014): 021023.

[23] Ryan, Jonathan. A Fully Integrated Sensor Fusion Method Combining a Single Antenna GPS Unit With Electronic Stability Control Sensors. Diss. Auburn University, 2011.

[24] Tumanski, Slawomir. Handbook of Magnetic Measurements. CRC Press, 2011.

[25] Canciani, Aaron J., Raquet, John F., Magnetic Anomaly Navigation Accuracy with Respect to Map Quality and Altitude, Proceedings of the 2016 International Technical Meeting of The Institute of Navigation, Monterey, California, January 2016, pp. 110-116.

[26] Li, You, Niu, Xiaoji, Zhang, Peng, Lan, Haiyu, Zhuang, Yuan, El-Sheimy, N., Smartphone-Based Indoor Navigation Using PDR and Magnetic Matching, Proceedings of the 28th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2015), Tampa, Florida, September 2015, pp. 2060-2066.

[27] Tyrn, Carl. Magnetic Terrain Navigation. Unmanned Untethered Submersible Technology, Proceedings of the 1987 5th International Symposium on. Vol. 5. IEEE, 1987.

[28] Gebre-Egziabher, Demoz, et al. A Gyro-Free Quaternion-Based Attitude Determination System Suitable for Implementation Using Low Cost Sensors. Position Location and Navigation Symposium, IEEE 2000. IEEE, 2000.

[29] Gebre-Egziabher, Demoz, et al. Calibration of Strapdown Magnetometers in Magnetic Field Domain. Journal of Aerospace Engineering 19.2 (2006): 87-102.

[30] Gebre-Egziabher, Demoz. Magnetometer Autocalibration Leveraging Measurement Locus Constraints. Journal of aircraft 44.4 (2007): 1361-1368.

[31] Dunbar, Brian. 2012: Magnetic Pole Reversal Happens All The (Geologic) Time. NASA. NASA, 30 Nov. 2011. Web. 26 May 2016.

[32] Wahba, Grace. A Least Squares Estimate of Satellite Attitude. SIAM review 7.3 (1965): 409-409.

[33] Davenport, Paul B. A Vector Approach to the Algebra of Rotations With Applications. Vol. 4696. National Aeronautics and Space Administration, 1968.

[34] Davenport, P. B., W. M. Rumpl, and G. L. Welter. In-Flight Determination of Spacecraft Magnetic Bias Independent of Attitude. (1988).

[35] Shuster, Malcolm David, and S. D. Oh. Three-Axis Attitude Determination From Vector Observations. Journal of Guidance, Control, and Dynamics 4.1 (1981): 70-77.

[36] Mortari, Daniele. ESOQ: A Closed-Form Solution to the Wahba Problem. Journal of the Astronautical Sciences 45.2 (1997): 195-204.

[37] Markley, F. Landis, and Daniele Mortari. How to Estimate Attitude From Vector Observations. (1999).

[38] Wertz, James R., ed. Spacecraft Attitude Determination and Control. Vol. 73. Springer Science & Business Media, 2012.

[39] Gambhir, B. Determination of Magnetometer Biases Using Module RESIDG. Computer Sciences Corporation, Report 3000-32700 (1975): 01.

[40] Lerner, B. Gambhir. Scalar Checking, Spcecraft Attitude Determination and Control, pp 328-334, 1978.

[41] Thompson, R. H., G. F. Neal, and M. D. Shuster. Magnetometer Bias Determination and Spin-Axis Attitude Estimation for the AMPTE Mission. Journal of Guidance, Control, and Dynamics 7.4 (1984): 505-507.

[42] Alonso, Roberto, and Malcolm D. Shuster. Attitude-Independent Magnetometer-Bias Determination: A survey. Journal of the Astronautical sciences 50.4 (2002): 453-476.

[43] Alonso, Roberto, and Malcolm D. Shuster. TWOSTEP: A Fast Robust Algorithm for Attitude-Independent Magnetometer-Bias Determination. Journal of the Astronautical Sciences 50.4 (2002): 433-452.

[44] Alonso, Roberto, and Malcohn D. Shuster. Complete Linear Attitude-Independent Magnetometer Calibration. Journal of the Astronautical Sciences 50.4 (2002): 477-490.

[45] Martin, Scott M. Closely Coupled GPS/INS Relative Positioning for Automated Vehicle Convoys. Diss. Auburn University, 2011.

[46] Faulkner, N. M., S. J. Cooper, and P. A. Jeary. Integrated MEMS/GPS Navigation Systems. Position Location and Navigation Symposium, 2002 IEEE. IEEE, 2002.

[47] Li, Yong, et al. Low-Cost Tightly Coupled GPS/INS Integration Based On A Nonlinear Kalman Filtering Design. Proceedings of ION National Technical Meeting. 2006.

[48] Stengel, R., 1994, Optimal Control and Estimation, Dover Publications, New York.

[49] Milham, Samuel, James B. Hatfield, and Richard Tell. Magnetic fields From Steel-Belted Radial Tires: Implications for Epidemiologic Studies. Bioelectromagnetics 20.7 (1999): 440-445.

[50] IAGA V-MOD Geomagnetic Field Modeling: International Geomagnetic Reference Field IGRF-12. IAGA V-MOD Geomagnetic Field Modeling: International Geomagnetic Reference Field IGRF-12. National Oceanic and Atmospheric Administration, 22 Dec. 2014. Web. 14 July 2016.