A Fully Integrated Sensor Fusion Method Combining a Single Antenna GPS Unit with Electronic Stability Control Sensors

by

Jonathan Ryan

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, 2011

Keywords: sensor fusion, vehicle dynamics, sideslip estimation, roll estimation, indirect tire pressure monitoring, steering misalignment detection

Approved by

David Bevly, Chair, Associate Professor of Mechanical Engineering
John Y. Hung, Professor of Electrical Engineering
Subhash Sinha, Professor of Mechanical Engineering

Abstract

This work presents a method for incorporating GPS (Global Positioning System) and standard roll stability control (RSC) sensors into the electronic stability control (ESC) and RSC systems. It is an adaptation of the very well known loosely-coupled GPS/INS (Inertial Navigation System) integration strategy which has been modified for the purposes of ESC systems. The first modification is the removal of the pitch rate gyroscope, a sensor which is unavailable on commercial vehicles. The second modification deals with the observability problems of the standard loosely coupled filter by adding heading constraints when the vehicle is not turning. The structure and algorithm of this method is presented. Observability conditions are evaluated, and the convergence of the estimates are analyzed via simulations. The conclusions from these simulations are compared with the expectations from the literature and observability condition checks. An experiment which illustrates the long term performance of the bias estimation was performed, followed by an experiment showing the roll and sideslip estimation performance during dynamic events. It is shown that over the long term the inertial bias estimates will converge if the vehicle experiences adequate dynamics, and that the system is able to accurately estimate sideslip and roll during dynamic maneuvers. The system is also able to estimate slow sideslip buildup, an important capability for ESC systems. The unified system is compared with a less integrated or "modular" approach for both experiments.

Furthermore, a method for using GPS to detect tire pressure changes is presented based on the hypothesis that the tire effective radius varies according to tire pressure. A technique using GPS and wheel speed signals to estimate the effective radius of the tires is discussed

and validated in simulation and experiment. Experiments are given to show how the radius estimate varies according to tire pressure, and a simple pressure loss detection law is discussed. A method to detect steering misalignment is also presented.

Acknowledgments

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-0809382. I would first like to acknowledge my advisor, Dr. Bevly, for his obvious role in my graduate education. In many ways he is the reason I started down this path of grad school, urging me to consider graduate studies in the first place and also challenging me to go after fellowships I never would have otherwise dreamed of. In many ways this exemplifies his main contribution to my education. That is, he believes in his students and expects them to achieve excellence. For this, and for his investment in me as a student, I am grateful. I also acknowledge his contribution to the GAVLAB intramural basketball team as one of our primary offensive producers. I also wish to acknowledge my committee, Dr. John Hung and Dr. Subhash Sinha, for their oversight on this thesis.

I would now like to acknowledge my parents, who have been encouraging and supportive throughout my entire life, and no less in graduate school. There is no doubt that I would not be where I am today without them. Likewise I want to thank my grandparents for their constant love and support, for always being there, and for always encouraging me to do my best, something which they have each modeled with great integrity.

I also owe much to my fellow students in the GAVLAB. Rob, Broderick, and Will were especially helpful when I first began and each frequently offered their own time to help me with this or that. Likewise Hodo, Ben, Lashley, and Wei were always there when I had questions. All of these guys set an example both in excellence in research and also in the importance of working with and investing in coworkers around you. I owe a great deal also to the men of L2: Lowell, Ryan, Scott and Wei. Those were good times, and I dare say we took L2 to new heights and ushered in its golden era (mostly because Ryan and Lowell took

the effort to deep clean the place). I also want to thank Jeremy, Jordan, Chris, and John for their input and help over the last months. Jordan especially deserves special gratitude for taking charge of the new G35 and making it the sleek, *007*-esque data collection machine that it is.

I wish to thank Ford Motor Company for sponsoring this work and giving me the opportunity to come to Dearborn for the summer. It was an invaluable experience, and I enjoyed meeting all of the good people up there. Special thanks goes to Jianbo Lu for his involvement in this work and all of the assistance he provided over the years.

Finally, I wish to acknowledge my wife Ginger Ryan with deepest gratitude for her unwavering love and support during this time. She has always completely supported me and believed in me throughout my studies, even when it meant delaying some of her own dreams. She has truly exemplified to me the love of Christ in the way that she has sacrificed for our family and cared for my interests above her own. Her companionship is sweet, her wisdom is deep, and her baking is delicious. This work would not be what it is without her, for I would not be who I am without her. For all of these people and this opportunity I give thanks to the Lord Jesus Christ.

Table of Contents

List of Figures

Chapter 1

Introduction and Background

Intelligent safety systems are an increasing focal point in today's automotive industry. The motivation for these systems stems from the tragic reality that tens of thousands of people are killed in motor vehicle accidents every year. For example, over 37,000 people were killed in vehicle accidents in 2008. Rollover accounted for 33% of these deaths. In fact, motor vehicle crashes are the leading cause of death among Americans between the ages of 1 and 34 [17]. The problem has received so much attention that Congress has enacted legislation requiring all new vehicles after 2012 to include certain intelligent safety systems as standard features such as electronic stability control (ESC).

Electronic stability control systems and roll stability control (RSC) systems are key elements of the modern effort to improve and increase the safety capabilities of passenger vehicles. ESC systems seek to control unsafe yaw and lateral motions of the vehicle. These motions can occur when the vehicle begins to lose traction and a dangerous over steer situation arises. One extreme example would be the back end of the vehicle "sliding out" or "fish-tailing" during a turn. ESC systems apply control in these situations via differential braking of the individual tires to create the desired control moment. This is also true of RSC systems, which as their name suggests seek to minimize unsafe levels of vehicle roll. Both of these systems utilize feedback control systems, requiring information about particular states of the vehicle. Two of the most critical states for these systems are the sideslip and roll angles, which are also two of the most expensive to directly measure. In order to obtain accurate information about these states without greatly increasing the production cost of the vehicle, state estimation theory must be applied with the sensors which are already on board. This is the goal of this thesis.

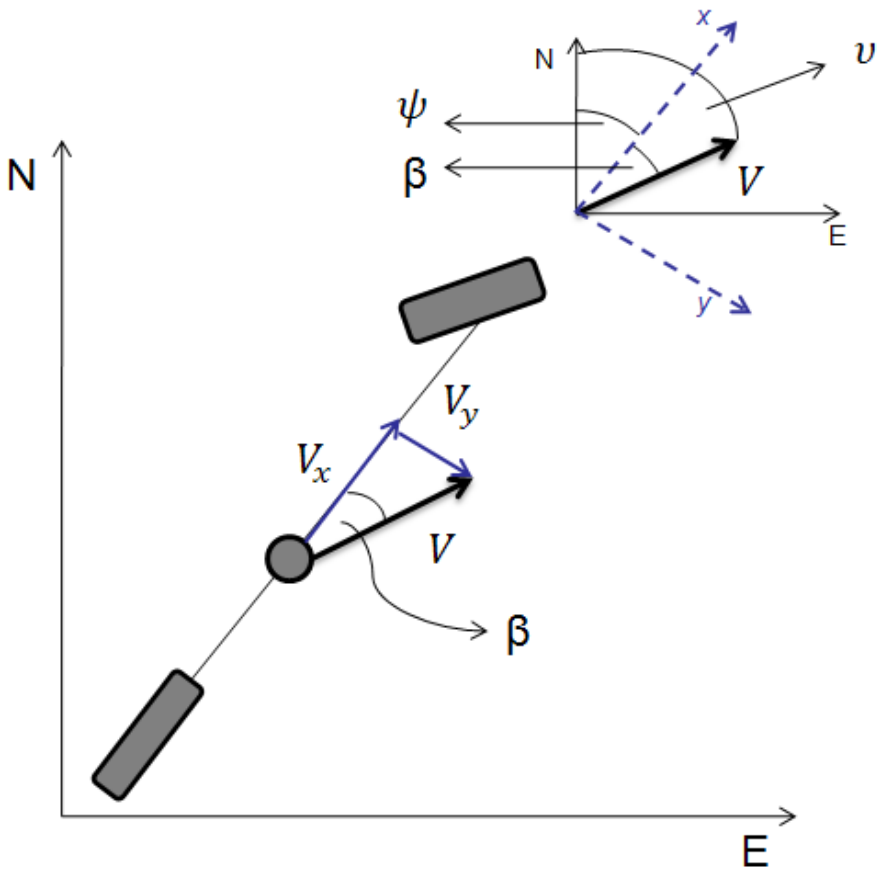

Figure 1.1: Sideslip Definitions in the Navigation and Body Coordinate Frames.

Now consider the two definitions of sideslip followed by a discussion of the vehicle coordinate frame. Figure 1.1 shows a simple diagram of a single track vehicle situated in the North-East plane of the North-East-Down (NED) coordinate frame. Not shown is the z axis of the NED frame, which goes into the page to complete a right handed system. The blue dashed lines show the body coordinate frame, with the x axis aligned with the forward direction of the vehicle, the y axis perpendicular to right side, and the z axis pointing into the page to complete the right handed system. For the purposes of this diagram, the two z axes are collinear, but this is not true generally. In Figure 1.1, the bold vector V represents the velocity vector of the vehicle's center of gravity, while the blue vectors V_x, V_y represent the components resolved into the body frame. The angle ν is defined as the "course" angle, and represents the angle of V from North. The angle ψ is defined as the "heading" angle

or yaw angle and represents the angular direction that the vehicle is facing from North. Both angles are positive in the clockwise direction. The angle β is the sideslip angle, which can now be defined in two ways. First, it can be thought of as the difference between the direction that the vehicle is *moving* and the direction that it is *facing* as described in (1.1).

$$\beta = \nu - \psi \tag{1.1}$$

Sideslip can also be thought of as the ratio of the lateral and forward velocities, as in (1.2).

$$\beta = atan(\frac{V_y}{V_x}) \tag{1.2}$$

In both definitions, the sideslip angle is positive in the clockwise direction. Equation (1.2) is more suitable to intuition, because it basically shows that sideslip is simply an angular representation of the vehicle's lateral velocity, scaled according to forward speed. If the sideslip angle is large, it means that the vehicle is sliding no matter what the forward speed is. While there is always a small amount of sideslip during any turn, it holds true that cars are not intended to go sideways. Therefore the goal for safety systems is to keep sideslip angles low.

Now consider the definitions of the attitude angles and the body fixed coordinate axis. Figure 1.2 shows the SAE coordinate axis for the body frame (original image courtesy of [32]). It is a right handed system, with the X axis pointing in the forward direction, the Y axis to the right of the vehicle, and the Z axis pointing down. The angle ϕ marks the positive roll angle direction, θ the positive pitch angle direction, and ψ the positive yaw angle direction. An interesting side note involving the pitch angle is that if the vehicle is driving on a road with a positive grade (pitch) angle, the Z velocity will actually be negative and vice versa. This is a point of potential confusion which should be remembered when considering road grades.

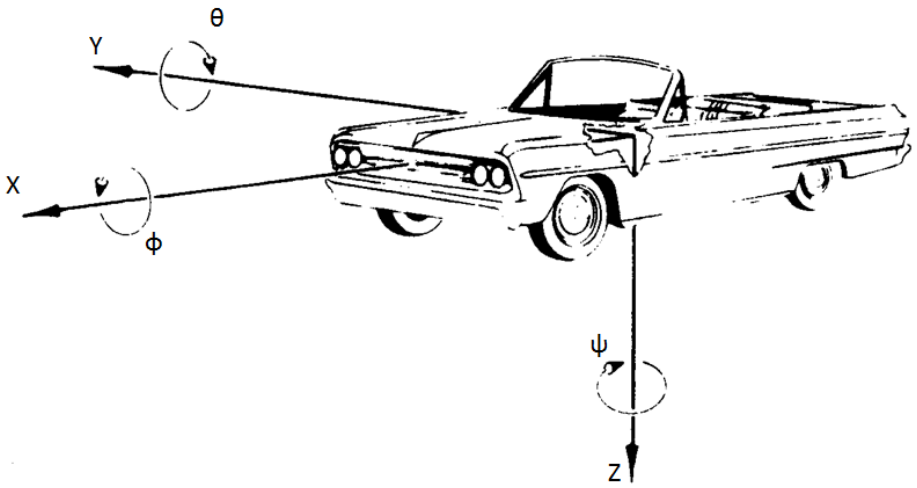

Figure 1.2: SAE Coordinate System, Body Frame [32].

1.1 Literature Review

State estimation approaches for vehicle dynamics can be broadly categorized into two groups: model-based approaches and kinematic approaches. Both approaches use some sort of "model", but the difference lies in what parameterizes the equations which make up the model. In the model-based case, the model parameters correspond to vehicle parameters such as mass, inertia, wheel base, tire stiffness, etc. Strictly speaking, model based estimators might even employ some kinematic equations in the model. What matters, however, is that these equations are inevitably parameterized by vehicle parameters. By contrast "kinematic estimators" do not use any vehicle parameters. Instead all terms in their "model" equations are constituted by sensor signals such as acceleration, rotation rate, velocity, or position. Many of these accelerations and rotation rates are directly measurable with common ESC sensors, while GPS provides information about the velocities and positions. Kinematic estimators do rely on knowledge of sensor location, although this has been shown possible to estimate accurately when it is difficult to measure. This distance known as the "lever arm" does not change, making it a relatively easy problem to overcome [27]. Kinematic estimators also rely on parameters of the *sensors*, such as bias time constants and noise variances, yet these are possible to identify off line. Furthermore, these parameters do not change. While it

is true that some of these sensor parameters are a function of temperature, this relationship can be accounted for as well. The primary advantage, then, of "kinematic estimators" is that they do not rely on changing vehicle parameters. The system is robust to different loadings, after market suspension modifications, tire wear, or even a completely different set of tires.

The significance of the roll and sideslip angles has led to many publications documenting various methods of their estimation. This includes a large body of research involving the use of GPS to estimate sideslip kinematicly, as opposed to using vehicle models. An important distinction among these is the number of GPS antennas used, as many current vehicles are already instrumented with a single antenna system as opposed to double or even triple antenna systems. Several researchers have proposed using single antenna GPS systems to improve sideslip estimation methods. A simple approach is presented in [2] and contrasted with results from a model based estimator. A more cascaded approach is presented in [8], where a sideslip estimate is first obtained with a yaw rate gyroscope. This estimate is used to aid sideslip estimation with a lateral accelerometer. This approach is expanded even further in [6]. Even more authors approach kinematic sideslip estimation using dual antenna GPS systems. Ryu uses a dual antenna GPS receiver with INS to estimate vehicle velocities, sideslip, roll, and road grade [40]. In [9], Bevly and Ryu present Kalman filtering methods for vehicle state estimation using both single and dual antenna GPS/INS systems. A planar model is no longer assumed, as the inertial measurements are compensated for the roll effects [9]. In [7] and [15], the authors use a dual antenna system to obtain estimates of the lateral states, including sideslip. These estimates are then used to determine tire parameters.

The kinematic approach to state estimation has also been a popular way to approach roll angle estimation. Tseng [44] presents a novel method of estimating both roll and pitch based on the inertial mechanization equation without using GPS at all. Yet far more authors elect to take advantage of GPS information. The previously referenced works [9], [6], and [40] focus on roll estimation as well sideslip. In [6] the lumped effects of roll and the accelerometer

bias are estimated using a single antenna GPS and then separated using a low pass filter. The approach in [9] first uses a single antenna system. The estimate of the roll is then extrapolated from the estimate of the lateral accelerometer bias. Since the accelerometer is not compensated for roll, the roll effects dominate the bias estimate and therefore this provides a good method for approximating roll. The authors then compare this method with results from integration with a dual antenna GPS system. Dual antenna systems can measure roll directly, and as a result GPS/INS integrations which use such systems can produce more accurate roll estimates. The dual antenna approach is also adopted in [40], which was one of the first papers to demonstrate this potential. The roll information can also be used with a vehicle model to separate suspension roll from the road bank, as done in [38]. Those authors go on to show in [39] that roll and pitch have significant effects on sideslip estimation. It follows then that it is important to include pitch information into the estimation algorithm. In urban areas the design limit for road grade can be up to 9% at speeds of 60 mph, and up to 12% for lower speeds around 30 mph. Rural roads can go up to 10% at design speeds of 40 mph and 8% for 60 mph design speeds [33]. Therefore knowledge of the pitch will be necessary to fully exploit these GPS/INS estimation schemes on steeper roadways.

Jansson estimates the road grade by combining GPS information with barometer and torque measurements into a Kalman filter [29]. Sahlholm and Johansson take a similar model based approach using drive line sensors and GPS; however they additionally present a method for recursively improving the grade estimate with new passes over the same road [41]. Lingman and Schmidtbauer also use a longitudinal vehicle model and Kalman filtering techniques to estimate both vehicle mass and road grade. This is done without any GPS information [30]. All of these are done under the context of longitudinal vehicle control, as opposed to lateral dynamic control applications. Bae and Ryu describe two methods for road grade estimation using GPS that are much more suitable for lateral estimation and control purposes [3]. These methods involve measuring total pitch directly with a dual antenna

GPS receiver or taking the arctangent of the ratio of the vehicle up and forward velocities obtained from a single antenna [3]. This method is expanded in [4], where the up and forward velocities are estimated using a simple Kalman filter. Since pitch rate sensors are uncommon on most commercial vehicles, the road grade estimate affords the opportunity of replacing the pitch rate gyro with the increasingly common single antenna GPS unit.

All of the above authors use "modular" sensor fusion approaches as opposed to those that are "unified" to estimate the roll, pitch, and sideslip angles. Speaking very specifically, the term "unified" is meant to signify a filter which incorporates *all* position, velocity, and attitude information of the three dimensional, six degree-of-freedom (DOF) model into *one single filter*. The term "modular" is meant to signify two things. First, it means that all of the states are not coupled together into one single filter, rather there are separate or cascaded filters for various sets of states. These separate filters may or may not be coupled together indirectly, sharing information about certain variables, but the distinction here is that this sharing is done *outside* of the filter. An example of a modular approach is [39], where the authors use one filter for the heading state and gyro bias and another for the forward and lateral velocities and accelerometer biases. Second, the term modular conveys that although there might be one single filter in the overall estimator, this filter does not include states for *all* six degrees of freedom. An example of this would be [15], where the filter is based on a planar model (the authors compensate for roll effects *outside* of the filter).

The following works consider approaches which are based on a unified integration scheme. The authors of [10] present a navigator based on a single antenna GPS integrated with a low cost INS in a loosely coupled integration. They evaluate its performance for positioning and present some experimental results. The authors of [16] use vehicle constraints, such as assuming no lateral or vertical velocity, to improve a standard loosely coupled GPS/INS implementation. These assumptions are used for performance improvement by many authors. In [18] the authors apply the constraints to a tightly coupled GPS/INS architecture. The authors of [20] also apply to velocity constraints to the loosely coupled

filter and evaluate its performance. The difference between these works and this thesis is three-fold. First, all of these works use a 6 DOF IMU/INS, whereas this work only uses a 5 DOF INS. Second, all of these works assume that the lateral velocity is zero *at all times*, whereas that assumption is made in this thesis only on the basis of certain conditions. Finally, all of these works evaluate the performance primarily with regard to position, velocity, or attitude as opposed to sideslip. Indeed they make the assumption of zero sideslip at all times, whereas in this thesis the sideslip estimation performance is the chief objective. In [26] the authors perform an observability analysis on the loosely coupled GPS/INS filter and find that under certain maneuvers, all states can be observed. It is also possible to use magnetometers and magnetic sensors embedded into the roadway to aid the filter. Yang and Farrell demonstrate this in [45] by creating a vehicle state estimation system having three layers of redundancy which uses magnetometers, GPS, and INS to determine the vehicle states. The accuracies and observabilities of the different estimates are discussed regarding the availability of each of the sensors. They show that adding the magnetometers increases the observability and eliminates the acceleration requirements.

The performance of the loosely coupled filter for position and velocity determination is well studied. However, the loosely coupled filter uses a pitch rate gyroscope which is not available on commercial vehicles. Therefore the filter's performance must be evaluated in light of this sensor reduction. Furthermore, the sideslip estimation performance of the loosely coupled filter is not well documented, because the loosely coupled filter is usually employed for general navigation purposes. The potential for sideslip estimation using this filter is briefly discussed in [19], but the performance is not analyzed. Some example plots of sideslip estimates are indeed shown, although these were not generated using a loosely coupled filter, rather they are estimates from the method presented in [15]. There is even less, if any, documented studies on the sideslip estimation performance of the loosely coupled filter when the pitch rate gyroscope is removed. Evaluating this performance is a primary goal of this thesis. The yaw information which is added to the filter based on certain conditions

(as described later) is also a new element. This constraint is practically the same as the one added to the estimators in [8], [2], but there is no documentation regarding how this additional yaw information affects the performance of the loosely coupled filter. Such analysis is an important part of understanding the sideslip estimation performance, and it is included in this thesis. Finally, there is a fair amount of literature documenting the observability characteristics of the system described by the loosely coupled filter, specifically noting that the observability of the system depends on the dynamics. Yet these results cannot be taken for granted regarding the filter used in this thesis, because of the aforementioned changes to the system. This thesis presents a summary of the observability results of the loosely coupled system, followed by an observability analysis of the presented system. Therefore this thesis is distinguished from the previous works by the lack of the pitch rate gyroscope, the sideslip estimation performance analysis, the conditionally included yaw constraint, and the observability analysis of this new system.

1.2 Contributions

The goal of this work is to achieve good sideslip, attitude, and velocity estimation in addition to inertial sensor bias identification by combining GPS with sensors which are already present in current RSC system sensor clusters. Furthermore, the operating window of the system needs to be expanded by making the system robust to all road geometries. The sensors used in this work are the GPS (single-antenna, 1Hz), accelerometers in x, y, z, a roll rate gyroscope, and a yaw rate gyroscope. There is no pitch rate gyroscope, because these are not present on commercial vehicles. Additionally, this work shows the capability of using GPS to identify changes in tire pressure without using pressure sensors. A method of using the yaw rate signal, improved with the bias estimates from the state estimator, and the steer angle sensor to estimate steering misalignment is also presented. All of these goals seek to use GPS to increase the intelligent safety capabilities of commercial vehicles without any additional sensor costs. Following are the contributions contained in this thesis:

- Development of an algorithm (initial results presented by the author in [37]) to provide sideslip and attitude estimates using ESC sensors and single antenna GPS.
- Analysis of sideslip and roll estimation performance of the algorithm.
- Analysis of performance improvements provided by conditionally added course measurements.
- Observability analysis of the algorithm.
- Performance analysis of modified "modular" method (initial results presented by the author in [36]).
- Development of a method to detect tire pressure changes using GPS and wheel speed sensors.
- Development of a method to detect steering misalignment using bias estimates from the state estimator, a yaw rate sensor, and the steer angle sensor.

The first contribution of this work is the development of a fully integrated state estimation algorithm, for the specific purpose of estimating sideslip, roll, and inertial biases, which uses only GPS and sensors present in RSC systems. The algorithm will be referred to throughout the thesis as the (au)tomotive (nav)igation (AUNAV) estimator. The AUNAV estimator is originally based on the well known "loosely coupled" GPS/INS integration scheme [19], [23]. The differences between the AUNAV estimator and the basic loosely coupled filter are that the AUNAV estimator does not use pitch rate gyroscope information and that it incorporates course information from the GPS, when it detects that the vehicle is not turning, in order to improve the sideslip estimation performance. The "navigation" term in the acronym points back to the original loosely coupled filter's purpose, and to the fact that the AUNAV estimator still possesses the same navigation functionality even though navigation is not its primary purpose. The AUNAV algorithm was first presented in [37],

showing initial sideslip and roll estimation results. Experimental data from two tests on low friction surfaces was used to compare the sideslip and roll estimates from the AUNAV filter with those produced by the commercially available Oxford RT3000 GPS/INS unit. It was shown that for both maneuvers the AUNAV filter produced roll angle estimates within one degree of the RT estimate. The sideslip estimate was within one degree of the RT estimate for the first run and within 2 degrees on the second run, which was much more challenging from an observability standpoint.

In this thesis the estimation performance of the AUNAV filter for sideslip, roll, and inertial bias estimation is analyzed with experimental data. The observability of the AUNAV filter is also analyzed, and the convergence expected from this analysis is studied by simulation and experiment. A comparison of the estimation performance of the modified "modular" estimator (MME) with the AUNAV filter is also given. The modular filter was developed by Bevly in [6] and represents a different approach to GPS/INS sensor fusion. In [36], the author of this thesis modified the filter by removing the pitch gyroscope and substituting the road grade estimate for the Euler pitch angle. The sideslip estimation performance of the MME when in the presence of larger road grades was analyzed. In this thesis, the MME estimation performance is analyzed for the same conditions as the AUNAV system for comparison.

Another contribution is a method to indirectly detect changes in tire pressure using only the sensors stated above in addition to wheel speed measurements, which are ubiquitous on commercial vehicles today. The approach is based on the idea that the "effective" or "rolling" radius of the tire will vary as a function of tire pressure. First, a simple method of estimating the rolling radius using GPS and wheel speed sensors is discussed. This is a recursive version of the batch least squares used in [13]. The effects of tuning and slip on the estimate are investigated, and simulations are presented to validate the estimate. Further validation is shown with experimental data. Next, experimental data is shown which illustrates how the radius estimate varies according to tire pressure. The repeatability and potential problems

of this method for inferring tire pressure are discussed. A simple method of detecting tire pressure loss well within the TREAD Act requirements using only the estimate of the rolling radius is put forward. The TREAD Act is a law requiring new vehicles to have Tire Pressure Monitoring Systems (TPMS) as standard features, and it contains some required performance specifications [1]. Future work to statistically improve the pressure loss detection algorithm is also discussed.

The final contribution is a method of detecting steering misalignments using the yaw rate sensor and the steering wheel angle sensor. The yaw rate signal is improved using the yaw rate bias estimate produced by the AUNAV estimator. The premise of the detection logic is that if the vehicle is driving straight, then the steer angle ought to be very close to zero. Steering effects from road crown can cause an offset, but these are ignored in this work as these effects remain the same throughout all experiments. It is shown that adjustments to the front right tire toe angle, which cause a misalignment, are detectable using the corrected yaw rate signal and the steering wheel angle signal.

Chapter 2

GPS/INS Integration Algorithms

2.1 Background

2.1.1 Kalman Filtering

A linear Kalman filter is simply a classical estimator in state variable format that incorporates the statistical knowledge of the system and sensors into the calculation of an "optimal" estimator gain L. In doing so the Kalman filter also calculates an estimate of the variances of the state estimation errors at each time step. This covariance matrix P can provide reliable confidence bounds on the estimates under certain conditions. Kalman filtering consists of two steps, referred to here as the measurement update and the time update. For GPS/INS integration applications the measurement update usually runs at a lower frequency than the time update. A new estimator gain L is calculated every time a new measurement arrives. This is contrary to basic pole-placement estimation, where the gain is constant. The innovations (the differences between the measurements and the state estimates) are taken and multiplied by the estimator gain, and this new quantity is added to the state estimate, just like a basic estimator. The covariance is updated at this time interval as well. The following equations describe the measurement update step [14]:

$$L = P_k \cdot C^T \cdot \left(C \cdot P_k \cdot C^T + R\right)^{-1} \tag{2.1}$$

$$P_k = (I - LC)\, P_k \tag{2.2}$$

$$\hat{X}_k = \hat{X}_k + L\left(Y - C\hat{X}_k\right) \tag{2.3}$$

In Equations (2.1-2.3), C represents the measurement matrix, $\hat{X}_k$ represents the current state estimate, and Y represents the measurement vector. It should be noted that neither the gain nor the covariance matrix depends on the measurement innovations or any inputs to the system. This follows from the strict list of assumptions that must be true to satisfy both the optimality of the Kalman filter and the accuracy of the covariance estimates. The time update consists of taking a model of the system, in either continuous or discrete representation, and propagating it forward in time just like a traditional estimator. Continuous equations are used during the time update in this thesis, therefore the overall filter is a continuous-discrete Kalman filter. At each time step the system input is measured, the rate of change of the state estimate is calculated, and the result is integrated. The rate of change of the variances must also be calculated and integrated. Trapezoidal integration is found to be sufficient for this work. The following are the equations for the time update [14]:

$$\dot{\hat{X}}_k = A\hat{X}_k + B_u u_k \tag{2.4}$$

$$\hat{X}_{k+1} = \hat{X}_k + \frac{1}{2}\Delta t\left(\dot{\hat{X}}_k + \dot{\hat{X}}_{k-1}\right) \tag{2.5}$$

$$\dot{P}_k = AP_k + P_kA^T + B_wQB_w^T \tag{2.6}$$

$$P_{k+1} = P_k + \frac{1}{2}\Delta t\left(\dot{P}_k + \dot{P}_{k-1}\right) \tag{2.7}$$

The matrices B_u and B_w are the system input and noise input matrices, respectively. It can be seen here that other factors affect the quality of the estimates, the optimality of the Kalman filter, and the accuracy of the error variance estimate P. First, the filter

assumes a perfect model in A, B_u, and B_w. Absolute perfection in this regard is highly unlikely. However, given that the models used for this work are kinematic sensor models, modeling uncertainty is not a major problem. Second, it is assumed that the statistics of the noise are perfectly known as well. This is not that bad of an assumption either, since the process noise is actually sensor noise which can be approximated from sampled data. Although such an approximation is not perfect, it is possible to achieve sufficient accuracy by analyzing sampled data. Third, it is assumed that the measurement errors are uncorrelated. The measurements in this work are the position and velocity solutions from a stand alone GPS receiver. These are likely the outputs of Kalman filters themselves, and would therefore have time correlated error, violating this assumption. Yet it has been shown that the impact from this is minimal, and may be overcome by simply increasing the tuning values for the measurements [23]. Therefore despite the stringent assumptions required to satisfy the filter's optimality and error variance estimation, the Kalman filter does perform very well for this application. Furthermore the structure of the Kalman filter offers many advantages for GPS/INS integration. GPS receivers offer highly accurate, unbiased estimates of a vehicle's velocity, but these are output at slower update rates on most receivers. They can also suffer from loss of signal in certain environments such as "urban canyons" or heavily wooded regions. INS systems boast much higher update rates but suffer from biases and errors that grow over time. Combining the two truly offers the best of both: an accurate, unbiased estimate of the vehicle states at a high update rate which can handle a loss of GPS satellite coverage for short periods of time.

2.1.2 Sensor Rotations and Coordinate Frames

The accelerometers and the gyroscopes of the IMU provide measurements which are resolved in the coordinate frame of the IMU (the sensor frame), and not the reference coordinate frame (the navigation frame). The two methods of GPS/INS integration presented in this work each use a different navigation frame. The MME filter uses the "local tangent"

frame. This X axis of this system always points in the forward direction of the vehicle (see Figure 1.1), but it is not aligned in the pitch direction with the X axis of the body frame. The Z axis of the vehicle points to the center of the reference ellipsoid. The origin of this frame is at projection of the center of gravity onto the local tangent plane. This frame can be thought of as traveling with the vehicle, but it does not pitch or roll. It's XY plane is always aligned with the local tangent plane, and it yaws with the vehicle. The AUNAV filter uses the North-East-Down frame. The X axis of this frame points due North, and the Z axis points to the center of the reference ellipsoid. The only difference between the two coordinate frames is a yaw angle rotation, specifically the yaw angle of the vehicle. In both navigation frames the Y axis is chosen to complete the right-handed coordinate systems. The body frame is in general not aligned with the navigation frame, potentially due to hills, banked roads, or suspension deflections caused by dynamic maneuvers or vehicle loading. It is important, then, to resolve both the IMU accelerometer and gyroscope signals into the common navigation frame through a series of steps known as IMU mechanization.

The following equations show how to resolve the IMU into the navigation frame. The accelerometers are rotated using the standard series of three body fixed Euler rotations. This not only brings them into the common navigation frame but also removes the gravitational effects present in the X and Y accelerometers. For rotations into the local tangent plane, the yaw angle ψ in (2.10) is 0. For rotations into the NED frame, ψ represents the vehicle heading. Equation (2.8) shows how the rotation matrix R_b^n is used to rotate the measured acceleration signals $\underline{a}^b$ from the body frame b into the navigation frame n.

$$\underline{a}^n = R_b^n \underline{a}^b \tag{2.8}$$

Technically, the inertial sensors only provide measurements in the body frame if the IMU is perfectly aligned with the vehicle body. This is rarely the case, however misalignments are usually very small and are not a focus of this thesis. The rotation matrix R_b^n is formed by a series of three body-fixed rotations which are yaw (ψ), pitch (θ), and roll (ϕ). Equation

(2.9) defines the rotation matrix in terms of the three individual rotations, while equations (2.10 - 2.12) define the matrices for each of the three rotations.

$$R_b^n = \left(R_n^b\right)^{-1} = \left(R_\phi R_\theta R_\psi\right)^{-1} \tag{2.9}$$

$$R_\psi = \begin{bmatrix} \cos(\psi) & \sin(\psi) & 0 \\ -\sin(\psi) & \cos(\psi) & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{2.10}$$

$$R_\theta = \begin{bmatrix} \cos(\theta) & 0 & -\sin(\theta) \\ 0 & 1 & 0 \\ \sin(\theta) & 0 & \cos(\theta) \end{bmatrix} \tag{2.11}$$

$$R_\phi = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\phi) & \sin(\phi) \\ 0 & -\sin(\phi) & \cos(\phi) \end{bmatrix} \tag{2.12}$$

The angular signals are resolved into navigation frame using the mechanization equations, expressed in the mechanization matrix $F_{\underline{\psi}}$. Equation (2.13) shows how the measured gyroscope signals $\underline{\omega}^b$ are resolved from the body frame b into the navigation frame n.

$$\underline{\omega}^n = F_{\underline{\psi}} \underline{\omega}^b \tag{2.13}$$

The mechanization matrix is a function of the level angles (roll and pitch). It is defined in equation (2.14).

$$F_{\underline{\psi}} = \frac{1}{\cos(\theta)} \begin{bmatrix} 1 & \sin(\phi)\sin(\theta) & \cos(\phi)\sin(\theta) \\ 0 & \cos(\phi)\cos(\theta) & -\sin(\phi)\sin(\theta) \\ 0 & \sin(\phi) & \cos(\phi) \end{bmatrix} \tag{2.14}$$

2.1.3 Inertial Sensor Models

Inertial sensors are advantageous in that they do not suffer from any sort of loss of signal, as long as they don't fail, and that they output signals at a very high update rate, which is necessary for control. However they do have several disadvantages. Both accelerometers and rate gyroscopes have static or "turn on" biases in addition to a moving bias which is known as "drift". The accelerometer (and gyroscope) drifts are modeled as first order Markov processes [6], being driven by white noise μ and having a time constant τ. They additionally suffer from noise which can be modeled as additive Gaussian white noise. The sensors also potentially have a scale factor which scales the true value being measured [6]. It has been shown that this scale factor can be estimated, and the model assumed in this work assumes that no scale factor error is present or that it has already been estimated. Equations (2.15-2.18) show the inertial sensor models:

$$\underline{a}^b_{meas} = \underline{a}^b_{true} + R^b_n \underline{g} + \underline{\delta f} + \underline{\epsilon}_{accel} \tag{2.15}$$

$$\underline{\omega}^b_{meas} = \underline{\omega}^b_{true} + \underline{\delta\omega} + \underline{\epsilon}_{gyro} \tag{2.16}$$

$$\delta f_i, \delta\omega_i = b_{\text{“}turn-on''} + b_{walk} \tag{2.17}$$

$$\dot{b}_{walk} \approx \frac{1}{\tau}\left(\mu - b_{walk}\right) \tag{2.18}$$

Note that Equations (2.17-2.18) apply for the biases of both the accelerometers and gyroscopes. The vectors $\underline{a}$ and $\underline{\omega}$ represent the acceleration and rotation rate vectors. The vectors $\underline{\delta f}$ and $\underline{\delta\omega}$ are the accelerometer and gyroscope biases. The "b" terms represent specific components of the biases, and the $\underline{\epsilon}$ vectors represent the noise.

2.1.4 GPS Sensor Models

When only inertial sensors are available, their biases are usually compensated through certain assumptions such as a flat road. If these do not hold, for example driving on a hill,

other information such as GPS is needed to compensate for the biases. Even after all of the biases are accounted for, the white noise remains and will still corrupt the solutions. Integrating this noise will potentially result in unbounded error growth. An additional sensor such as GPS is therefore necessary to bound the error growth. The GPS sensor model highlights the well documented advantages of GPS [6]. In the GPS/INS integration strategy used in this thesis, the GPS receiver outputs position solutions as latitude, longitude, and altitude while the velocity solutions are output in the North-East-Down (NED) navigation frame. The velocity signals are unbiased and contain no scale factor error. The errors present on the GPS positions and velocities can be modeled as uncorrelated Gaussian random noise, as seen in Equation (2.19).

$$\begin{bmatrix} \lambda \\ \varphi \\ h \end{bmatrix}_{meas} = \begin{bmatrix} \lambda \\ \varphi \\ h \end{bmatrix} + \underline{\eta}_{lla}$$
$$\underline{V}^n_{meas} = \underline{V}^n + \underline{\eta}_v \tag{2.19}$$

In this equation, λ represents the latitude, φ represents the longitude, and h represents the altitude. $\underline{V}^n_{meas}$ is the measured velocity, $\underline{V}^n$ is the true velocity, and $\underline{\eta}_i$ is the noise vector for each solution. There do exist more detailed models of GPS errors, but this simplified model is representative of the error behavior at the position and velocity level and is sufficient for this work.

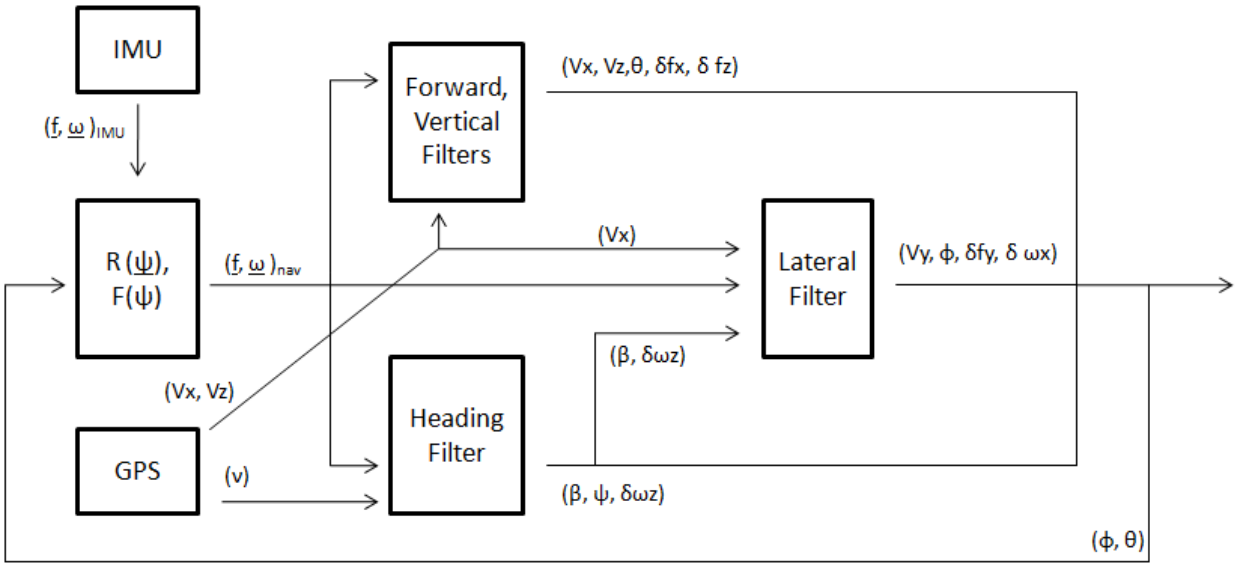

Figure 2.1: Diagram of the Modified Modular Estimator.

2.2 GPS/INS: The Modified Modular Estimator

2.2.1 Estimation Strategy Overview

The modified modular estimator is an extension of the estimator presented in [6]. The estimation strategy is largely directed by the sensors and corresponding measurements available. This thesis assumes the following sensors to be available: a five degree of freedom (DOF) IMU consisting of accelerometers mounted in the body x, y, and z axes along with roll rate and yaw rate gyroscopes, and a single antenna GPS receiver. While 5 DOF IMU clusters are not ubiquitous in the automotive world, they can be found in certain production vehicles with a roll stability control system or with a rollover curtain system. The following paragraph gives a broad, big picture outline of the estimation strategy.

Figure 2.1 is a block diagram showing an overview of the process. First the IMU signals are resolved into the navigation frame using the current attitude angle estimates. This removes the biases resulting from gravity which are present in the x and y accelerometer measurements. Second, Kalman filters are used to estimate the vertical and longitudinal velocities as described in [4]. These velocities are then used to calculate an accurate estimate

of the road grade, which is used as a substitute for the true pitch angle. Next, a Kalman filter which estimates the vehicle heading and the yaw gyroscope bias is used, together with the GPS course measurement, to obtain an initial estimate of the sideslip. Finally, the initial sideslip estimate is combined with the GPS velocity measurement to produce a derived "measurement" of the lateral velocity for a lateral state Kalman filter. This Kalman filter estimates the vehicle's lateral velocity, lateral accelerometer bias, roll angle, and roll rate gyroscope bias. The final sideslip estimate is calculated in a straightforward manner from the lateral velocity estimate [8], [9]. It is available at the frequency of the IMU, while the initial sideslip estimate is only available at the frequency of the GPS.

2.2.2 Road Grade Estimation

The slope of the road in the forward direction is generally referred to as the road grade. It can be presented in two forms: either as the actual slope of road (% grade) or as the angle that the road makes with the horizon; where conversion between the two is a matter of simple trigonometry. Determination of this angle can be accomplished using the ratio of the vertical and horizontal speeds, assuming that the vehicle is moving (so as not to have a zero in the denominator). GPS receivers output a vertical speed and a speed-over-ground velocity vector, where the speed over ground is the vehicle velocity vector in the local navigation frame. The magnitude of this vector can be taken to be the longitudinal speed of the vehicle in the navigation frame, assuming little to no sideslip. Therefore the arctangent of the two speeds can be taken to find the grade angle. This method has been proven to produce high quality, unbiased estimates of the road grade, yet it should be noted that any bounce motions that the vehicle experiences will affect the grade estimate. However, this has been shown not to significantly diminish the estimator's performance [4].

Coupling the GPS system and the IMU together in a Kalman filter structure offers many advantages. The following equations, which are simply scalar versions of the vector Equation (2.15) , show how this is possible:

$$a_z \approx \dot{V}_z + g_z + \delta f_z + \epsilon_{accel} \tag{2.20}$$

$$a_x \approx \dot{V}_x + g_z \sin\theta + \delta f_x + \epsilon_{accel} \tag{2.21}$$

where a_z, a_x represent the accelerometer measurements in the z and x directions, $\dot{V}_z$ and $\dot{V}_x$ are the true accelerations, g_z is gravity, θ is the total Euler pitch, $\delta f_{z,x}$ is the inherent sensor bias, and ϵ is the sensor noise. The Kalman filter now takes the following form, where the accelerometer measurement is the input and the GPS velocity is the measurement:

$$\begin{bmatrix} \dot{\hat{V}}_z \\ \delta\dot{\hat{f}}_z \end{bmatrix} = \begin{bmatrix} 0 & -1 \\ 0 & -\frac{1}{T_m} \end{bmatrix} \begin{bmatrix} \hat{V}_z \\ \delta\hat{f}_z \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \end{bmatrix} (a_z - g_z) + \begin{bmatrix} 1 & 0 \\ 0 & \frac{1}{T_m} \end{bmatrix} \epsilon \tag{2.22}$$

$$y = \begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} \hat{V}_z \\ \delta\hat{f}_z \end{bmatrix} + \eta_z \tag{2.23}$$

$$\begin{bmatrix} \dot{\hat{V}}_x \\ \delta\dot{\hat{f}}_x \end{bmatrix} = \begin{bmatrix} 0 & -1 \\ 0 & -\frac{1}{T_m} \end{bmatrix} \begin{bmatrix} \hat{V}_x \\ \delta\hat{f}_x \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \end{bmatrix} (a_x - g\sin\theta) \begin{bmatrix} 1 & 0 \\ 0 & \frac{1}{T_m} \end{bmatrix} \epsilon \tag{2.24}$$

$$y = \begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} \hat{V}_x \\ \delta\hat{f}_x \end{bmatrix} + \eta_x \tag{2.25}$$

where η and ϵ are the measurement noise and process noises, respectively, $\hat{V}_{z,x}$ represents the velocity estimates, and $\delta\hat{f}_{z,x}$ represent the bias estimates. T_m represents the bias time constants. Both of these estimator models are observable. The road grade estimate is the arctangent of the two speed estimates.

$$\hat{\theta} = -\arctan\frac{\hat{V}_z}{\hat{V}_x} \tag{2.26}$$

Since the input to the Kalman filter is an accelerometer, the process noise is taken to be noise from the accelerometer where the noise driving the bias Markov process is included as well. Both are taken to be Gaussian white noise, as shown below.

$$\epsilon = \begin{bmatrix} \epsilon_{accel} \\ \mu \end{bmatrix} \tag{2.27}$$

$$\epsilon_{accel} \sim N(0, \sigma^2_{accel}) \tag{2.28}$$

$$\mu \sim N(0, \sigma^2_m) \tag{2.29}$$

$$E\{\epsilon\epsilon^T\} = Q = \begin{bmatrix} \sigma^2_{accel} & 0 \\ 0 & \sigma^2_m \end{bmatrix} \tag{2.30}$$

The measurement noise is the noise on the GPS velocity signal, which is also assumed to be Gaussian white noise. Due to satellite orientations the noise is generally higher on the vertical speed measurement than on the speed over ground measurement. The measurement noise is taken to be

$$\eta_x \sim N(0, \sigma^2_{GPSx}) \tag{2.31}$$

$$\eta_z \sim N(0, \sigma^2_{GPSz}) \tag{2.32}$$

$$E\{\eta_x \eta_x^T\} = R_x = \sigma^2_{GPSx} \tag{2.33}$$

$$E\{\eta_z \eta_z^T\} = R_z = \sigma^2_{GPSz} \tag{2.34}$$

This filter structure combines the advantages of the high accuracy, unbiased GPS velocity measurements with the high update rate signals from the IMU, thereby offering high quality velocity estimates (and therefore a high quality road grade estimate) at a frequency suitable for control signals. The GPS measurement enables correction for any accelerometer biases, while the accelerometer allows for continued velocity tracking during GPS outages over short time intervals [4].

2.2.3 Heading Estimation

Vehicle sideslip is defined in two ways. It is the difference between the vehicle course (the direction that the vehicle is traveling) and the vehicle heading (the direction that the vehicle is pointed). It is also the arctangent of the ratio of lateral speed to longitudinal speed (all references to course and velocities "of the vehicle" are with respect to the IMU location, which is near the center of gravity.) The vehicle course and longitudinal speed are readily available with a single antenna GPS receiver, but neither heading nor lateral velocity are measurable with only one antenna. This renders sideslip technically unobservable, but it does not make it impossible to produce useful sideslip estimates under certain conditions. A yaw rate gyro could be integrated to obtain a heading estimate which could be subtracted from the course measurement to obtain sideslip. However doing so first requires overcoming several problems resulting from integration. The first obstacle is any bias that may be present in the sensor. This can be estimated and removed using a Kalman filter under the condition that the vehicle is driving straight. When this is true, the course angle and the heading angle are the same. A notable exception occurs on a banked road, causing a small steady sideslip, although this would be a rare scenario for straight sections of roadway. Common road crowns will yield on a very slight steady sideslip angle. Since GPS measures course and the course is equal to heading when the vehicle is driving straight, the estimation can be "switched on" during this scenario. The Kalman filter could then estimate any biases in the gyroscope and remove them. When a turning maneuver is initiated the estimation is "switched off" and the yaw gyro is integrated to determine heading. The difference in this heading estimate and the course measurement becomes the sideslip estimate. This process is seen the equations below, beginning with the gyroscope model, which is the scalar version of (2.16).

$$\omega_r \approx \dot{\psi} + \delta\omega_r + \epsilon_{gyro} \tag{2.35}$$

The previous equation, along with (2.17), now yields the Kalman filter.

$$\begin{bmatrix} \dot{\hat{\psi}} \\ \delta\dot{\hat{\omega}}_r \end{bmatrix} = \begin{bmatrix} 0 & -1 \\ 0 & -\frac{1}{T_m} \end{bmatrix} \begin{bmatrix} \hat{\psi} \\ \delta\hat{\omega}_r \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \end{bmatrix} (\omega_r) + \begin{bmatrix} 1 & 0 \\ 0 & \frac{1}{T_m} \end{bmatrix} \epsilon_r \tag{2.36}$$

When driving straight:

$$y = \begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} \hat{\psi} \\ \delta\hat{\omega}_r \end{bmatrix} + \eta_\nu = \nu^{GPS} \tag{2.37}$$

Otherwise, the estimation is switched off:

$$y = \begin{bmatrix} 0 & 0 \end{bmatrix} \begin{bmatrix} \hat{\psi} \\ \delta\hat{\omega}_r \end{bmatrix} + \eta_\nu = \nu^{GPS} \tag{2.38}$$

When the vehicle is driving straight, the C matrix is set to [1 0] and measurement updates are performed. The C matrix is set to [0 0] when turning. In both cases the sideslip estimate is the difference between the course measurement and the heading estimate:

$$\hat{\beta}_0 = \nu^{GPS} - \hat{\psi} \tag{2.39}$$

As in the case of the road grade estimation filters, the input to the system is an inertial sensor measurement. The process noise is the noise on the gyroscope signal and the noise driving the random walk. Both are assumed to be zero mean Gaussian white noise:

$$\epsilon_r = \begin{bmatrix} \epsilon_{gyro} \\ \mu \end{bmatrix} \tag{2.40}$$

$$\epsilon_{gyro} \sim N(0, \sigma^2_{gyro}) \tag{2.41}$$

$$\mu \sim N(0, \sigma^2_m) \tag{2.42}$$

$$E\{\epsilon_r \epsilon_r^T\} = Q = \begin{bmatrix} \sigma_{accel}^2 & 0 \\ 0 & \sigma_m^2 \end{bmatrix} \tag{2.43}$$

The measurement noise is the noise on the GPS course measurement, and is assumed to be zero mean Gaussian white noise as well. However, the accuracy of this measurement increases with speed, giving a variance that is a function of speed.

$$\eta_\nu \sim N(0, \sigma_{GPS\nu}^2) \tag{2.44}$$

$$E\{\eta_\nu \eta_\nu^T\} = R_\nu = \sigma_{GPS\nu}^2 \tag{2.45}$$

$$\sigma_{GPS\nu}^2 = \frac{\sigma_\nu^2}{V} \tag{2.46}$$

The second problem lies in determining whether or not the vehicle is going straight. It is imperative to the estimator performance that the biases are estimated accurately prior to periods of integrating the gyroscope signal. This makes it necessary to include logic statements to determine whether or not the vehicle is driving straight, which is done using the the yaw gyro signal. The basic logic is that if the absolute value of the yaw rate is less than some threshold then the vehicle is deemed to be driving straight. Complications arise, however, from the noise on the signal. The resulting situation becomes a trade-off in the sensitivity to turning motion verses false alarms caused by the noise. If the thresholds are set too close to the noise floor the estimator will "believe" the vehicle is turning during many time instances in which it is not. This means that the heading estimation is turned off, and no gyro bias correction is done. In this work the GPS measurement only comes in at 1Hz, so losing that measurement can be very costly. The primary limitation of this sideslip estimator is its dependency on periods of straight driving to zero out the gyro bias, and every missed course measurement results in a lost second of straight driving. At moderate speeds on winding roads, for example, one second of straight driving can be rare. Therefore the tuning of these thresholds plays an important role in the overall performance of the estimator.

The other obvious problem with this method comes from integrating a noisy sensor signal. Some improvement could be found by low pass filtering the signal, but all of the noise cannot be removed. Therefore this estimation scheme is limited as well in terms of the amount of time that the gyro can be integrated before the error grows too large.

2.2.4 Lateral Velocity

As previously noted, sideslip can be defined as the arctangent of the ratio of the lateral speed of the vehicle to its longitudinal speed. The longitudinal speed can be estimated accurately by combining wheel speed sensors, longitudinal accelerometers, and GPS, but the lateral speed is not measurable with only a single antenna GPS system. Using the sideslip estimate from the heading Kalman filter described above, a lateral velocity "measurement" can be generated and therefore a lateral state estimator can be introduced in the form of a Kalman filter. The following equations form the foundation of this estimator.

$$a_y \approx \dot{V}_y + V_x\dot{\psi} + g_z \sin\phi + \delta f_y + \epsilon_{accel} \tag{2.47}$$

$$\omega_p \approx \dot{\phi} + \delta\omega_p + \epsilon_{gyro} \tag{2.48}$$

As was the case before, this sensor model assumes no scale factor error. The biases for both the lateral accelerometer and the roll rate gyro are taken to be Markov random process driven by Gaussian white noise as described previously. There are two sources of bias in the lateral accelerometer measurement, one being the roll component of gravity and the other being the sensor's random walk. Both of these influences have the exact same effect on the velocity error, therefore the filter will not be able to distinguish the two. If a direct measurement of roll were available, from a double antenna GPS system for example, the two would be independently observable. Yet since roll is not measurable with a single antenna, the two states must be lumped together into one. The resulting simplification of the equation follows from using the small angle approximation and lumping the two terms together.

$$a_y \approx \dot{V}_y + V_x\dot{\psi} + g_z\left(\phi + \delta f_y\right) + \epsilon_{accel} \tag{2.49}$$

This results in an estimator in the form:

$$\begin{aligned}
\begin{bmatrix} \dot{\hat{V}}_y \\ \left(\dot{\hat{\phi}} + \delta\dot{\hat{f}}_y\right) \\ \delta\dot{\hat{\omega}}_p \end{bmatrix} &= \begin{bmatrix} 0 & -g_z & 0 \\ 0 & 0 & -1 \\ 0 & 0 & -1/T_m \end{bmatrix} \begin{bmatrix} \hat{V}_y \\ \left(\hat{\phi} + \delta\hat{f}_y\right) \\ \delta\hat{\omega}_p \end{bmatrix} \\
&+ \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} a_y - \hat{V}_x\left(\omega_r - \delta\hat{\omega}_r\right) \\ \omega_p \end{bmatrix} \\
&+ \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1/T_m \end{bmatrix} \begin{bmatrix} \epsilon_{accel} \\ \epsilon_{gyro} \\ \mu \end{bmatrix}
\end{aligned} \tag{2.50}$$

$$y = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} \hat{V}_y \\ \left(\hat{\phi} + \delta\hat{f}_y\right) \\ \hat{b}_p \end{bmatrix} + \eta_y = V^{GPS}\sin\left(\hat{\beta}_0\right) \tag{2.51}$$

$$\hat{\beta} = \arctan\left(\frac{\hat{V}_y}{\hat{V}_x}\right) \tag{2.52}$$

The process noise for the system arises from the IMU measurements. Since the measurement is constructed using the sideslip estimate from the heading filter, the measurement noise is approximated as the noise on the GPS course measurement.

$$E\left\{\omega\omega^T\right\} = Q = \begin{bmatrix} \sigma^2_{accel} + V^2\sigma^2_{gyro} & 0 & 0 \\ 0 & \sigma^2_{gyro} & 0 \\ 0 & 0 & \sigma^2_m \end{bmatrix} \tag{2.53}$$

$$E\{\eta_\nu \eta_\nu^T\} = R_\nu = \sigma^2_{GPS\nu} \tag{2.54}$$

It was stated above that the Kalman filter cannot separate the roll and the lateral accelerometer bias. The reason is because the two states look exactly the same in terms of velocity errors. However, they do not share the same frequency characteristics. That is, they do not change in the same way. This difference can be exploited using complimentary low and high pass filters as shown in [6]. The following equations show the Laplace representation of the complimentary filters.

$$\begin{aligned} \hat{\phi} &= \frac{T_m s}{T_m s + 1}\left(\hat{\phi} + \delta\hat{f}_y\right) \\ \delta\hat{f}_y &= \frac{1}{T_m s + 1}\left(\hat{\phi} + \delta\hat{f}_y\right) \end{aligned} \tag{2.55}$$

Since the roll angle changes much faster than the bias, this filtering approach yields an accurate estimate of both states independently.

2.3 GPS/INS: The Automotive Navigation (AUNAV) Estimator

The AUNAV estimator is most easily understood as a modified version of the generic loosely coupled GPS/INS blending strategy. Therefore this section is divided into two subsections, the first of which details the standard loosely coupled algorithm. After this, the modifications which are made to complete the AUNAV filter are discussed.

2.3.1 The Loosely Coupled Algorithm

The "loosely coupled" GPS/INS method of integration is a well documented technique for blending the GPS and INS navigation solutions [19]. Figure 2.2 gives a high level view of the integration. The INS and GPS position and velocity solutions are compared, and the difference is input into the extended Kalman filter (EKF) as a measurement. This is

because the states of the filter are the differences between the true states and INS estimates, as opposed to actual position and velocity states. The filter outputs estimated corrections to the INS solution, which are added to the INS solution. The corrected INS solution serves as the final estimate. The EKF also estimates the inertial errors of the INS, which are fed back into the INS to continuously improve the estimation process. The system functions basically as two independent navigators (the GPS and INS) and one EKF. When using lower grade inertial navigation systems, the loosely coupled approach treats the GPS solutions as truth [19]. Therefore the states of the EKF practically become the differences between the INS and GPS solutions. It should be noted that the inertial bias estimates are not error states, but rather they represent estimates of the actual biases themselves.

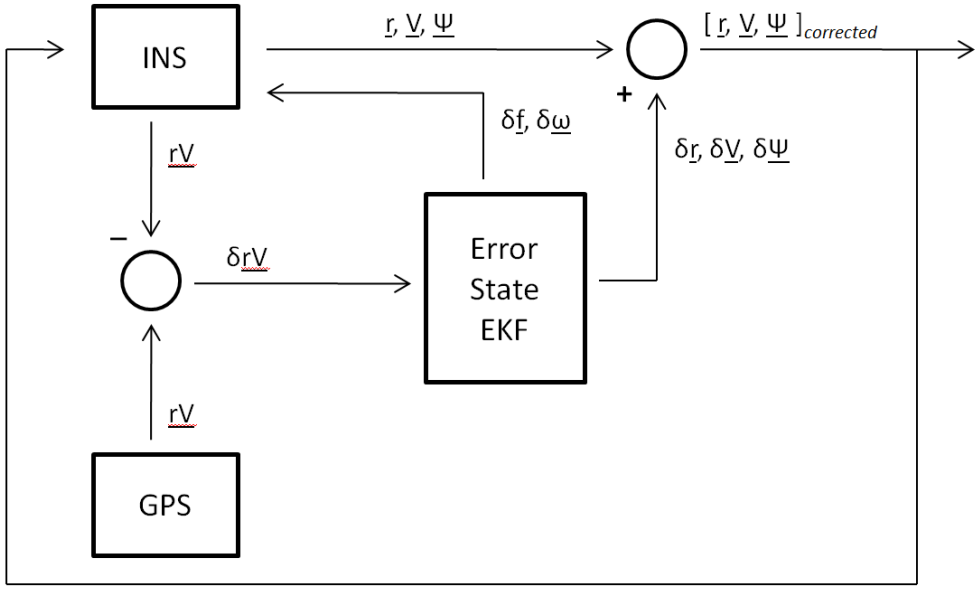

Figure 2.2: Diagram of the Standard Loosely Coupled Integration Strategy.

The details are of the loosely coupled approach are described as follows. The states of the Kalman filter, $\delta\hat{\underline{X}} = \left[\begin{array}{ccccc} \delta\hat{\underline{r}}, & \delta\hat{\underline{V}}, & \delta\hat{\underline{\psi}}, & \delta\hat{\underline{f}}, & \delta\hat{\underline{\omega}} \end{array} \right]'$, are the estimates of the errors in the INS position-velocity-attitude (PVA) solution and the biases of the inertial sensors. The first step of the overall system is to calculate the INS solution. The inertial sensors are rotated into the navigation frame and then propagated forward in time via trapezoidal

integration to obtain the PVA estimates as in Equations (2.56-2.62), a process also known as IMU mechanization.

$$\dot{\underline{\hat{\psi}}}_k = F_{\underline{\psi}}(\underline{\omega}^{INS} - \delta\underline{\hat{\omega}}) \tag{2.56}$$

$$\dot{\underline{\hat{V}}}_k = R_b^n(\underline{a}^{INS} - \delta\underline{\hat{f}}) - \underline{g} \tag{2.57}$$

$$\dot{\underline{\hat{r}}}_k = T\underline{\hat{V}}_k \tag{2.58}$$

$$\underline{\hat{\psi}}_{k+1} = \underline{\hat{\psi}}_k + \frac{1}{2}\Delta t\left(\dot{\underline{\hat{\psi}}}_k + \dot{\underline{\hat{\psi}}}_{k-1}\right) \tag{2.59}$$

$$\underline{\hat{V}}_{k+1} = \underline{\hat{V}}_k + \frac{1}{2}\Delta t\left(\dot{\underline{\hat{V}}}_k + \dot{\underline{\hat{V}}}_{k-1}\right) \tag{2.60}$$

$$\underline{\hat{r}}_{k+1} = \underline{\hat{r}}_k + \frac{1}{2}\Delta t\left(\dot{\underline{\hat{r}}}_k + \dot{\underline{\hat{r}}}_{k-1}\right) \tag{2.61}$$

$$T = \begin{bmatrix} \frac{1}{R_n+h} & 0 & 0 \\ 0 & \frac{1}{(R_e+h)\cos(\lambda)} & 0 \\ 0 & 0 & -1 \end{bmatrix} \tag{2.62}$$

In Equations (2.56-2.57), R_b^n represents the rotation matrix from the body frame to the NED frame, and $F_{\underline{\psi}}$ represents the mechanization matrix to align the gyroscopes with the NED frame. These are outlined in Section 2.1.2 in more detail. The matrix T transforms the velocities from the NED frame into rates of latitude, longitude, and altitude. R_e and R_n are parameters of the reference ellipsoid. During this step the EKF also propagates the uncertainty of the error estimates forward using standard Kalman filter covariance update equations linearized about the current state estimates. This is shown in Equation (2.63),

where J is the Jacobian (defined later), and B_w is simply a 15×15 matrix with the vector $([0_{3x3}, R_b^n, R_b^n, R_b^n, R_b^n])$ on the diagonal and 0's everywhere else.

$$\dot{P}_k = JP_k + P_k J^T + B_w Q B_w^T \tag{2.63}$$

Equation (2.63) is then integrated according to Equation (2.7) to complete the covariance propagation. Equations (2.56-2.63) are done at the update rate of the INS. The error estimates are not propagated during this time, because they are reset to zero each time they are added to the INS solutions. When GPS measurements arrive, the difference between the GPS and INS position and velocity solutions are taken in Equation (2.64).

$$Y = \begin{bmatrix} \underline{r} \\ \underline{V} \end{bmatrix}_{GPS} - \begin{bmatrix} \underline{\hat{r}} \\ \underline{\hat{V}} \end{bmatrix} \tag{2.64}$$

This difference serves as measurement for the error state Kalman filter. At this point the Kalman filter calculates the gain L, error residuals, and the covariance matrix as per the standard Kalman filter measurement equations outlined in Section 2.1.1; followed by updating the error states as in Equation (2.65).

$$\delta\underline{\hat{X}}_k = \delta\underline{\hat{X}}_k + L\left(Y - C\delta\underline{\hat{X}}_k\right) \tag{2.65}$$

The term C in Equation (2.65) represents the measurement model. It comes from the fact that the GPS outputs position and velocity measurements. The C matrix is defined below.

$$C = \begin{bmatrix} I_{6x6} & 0_{6x9} \end{bmatrix}_{6x15} \tag{2.66}$$

The estimates of the INS velocity and attitude errors are then added to the current INS solution, as per Equations (2.67-2.69), and are subsequently reset to zero.

$$\underline{\hat{r}} = \underline{\hat{r}} - \delta\underline{\hat{r}} \tag{2.67}$$

$$\underline{\hat{V}} = \underline{\hat{V}} - \delta\underline{\hat{V}} \tag{2.68}$$

$$(R_b^n)_k = \left(I + \left(\delta\underline{\hat{\psi}}\times\right)\right)(R_b^n)_k \tag{2.69}$$

The notation $(\delta\underline{\hat{\psi}}\times)$ indicates the skew symmetric form of the vector $\delta\underline{\hat{\psi}}$. The Euler angles are computed from R_b^n as in (2.70-2.72) [19].

$$\phi_k = \arctan\left(\frac{R_b^n\,(3,2)_k}{R_b^n\,(3,3)_k}\right) \tag{2.70}$$

$$\theta_k = -\arcsin\left(R_b^n\,(3,1)_k\right) \tag{2.71}$$

$$\psi_k = \arctan\left(\frac{R_b^n\,(2,1)_k}{R_b^n\,(1,1)_k}\right) \tag{2.72}$$

At each time step, whether during the INS estimate propagation or during a GPS update, the sideslip estimate is calculated in Equation (2.73) [19].

$$\hat{\beta} = \arctan\left(\frac{\hat{V}_{east}}{\hat{V}_{north}}\right) - \hat{\psi}_{yaw} \tag{2.73}$$

Finally, the estimated inertial sensor errors $\delta\underline{\hat{f}}$ and $\delta\underline{\hat{\omega}}$ are fed back into the INS (the subtracted terms in Equations (2.56-2.57). This is not necessary for high grade inertial units, but it is absolutely imperative in this work due to the relatively large errors in the automotive grade sensors. This feedback is known as a closed loop implementation as opposed to open loop.

The extended Kalman filter is based on the inertial sensor error models, Equations (2.15-2.18), and on the error propagation equations of the INS position, velocity, and attitude solutions which are Equations (2.74-2.76), which are simplified versions of those found in [19]. These are the equations used to form the Jacobian matrix which is required to update the covariances. δ signifies INS estimation error.

$$\delta\dot{\underline{r}} \approx T\delta\underline{V} \tag{2.74}$$

$$\delta\dot{\underline{V}} \approx \left(\left(R_b^n \cdot \underline{a}^{INS}\right)\times\right) \cdot \delta\underline{\psi} + R_b^n \cdot \delta\underline{f} \tag{2.75}$$

$$\delta\dot{\underline{\psi}} \approx -R_b^n \cdot \delta\underline{\omega} \tag{2.76}$$

$$J = \begin{bmatrix} 0_3 & T & 0_3 & 0_3 & 0_3 \\ 0_3 & 0_3 & \left(\left(R_b^n \cdot \underline{a}^{INS}\right)\times\right) & R_b^n & 0_3 \\ 0_3 & 0_3 & 0_3 & 0_3 & -R_b^n \\ 0_3 & 0_3 & 0_3 & -\frac{1}{\tau}\cdot I_3 & 0_3 \\ 0_3 & 0_3 & 0_3 & 0_3 & -\frac{1}{\tau}\cdot I_3 \end{bmatrix} \tag{2.77}$$

2.3.2 Modifications to the Loosely Coupled Algorithm

There are two primary changes that need to be made to the loosely coupled filter in order to accurately estimate sideslip and roll using only sensors available on current vehicles. First, the pitch rate gyroscope must be removed, as these are not currently available on commercial vehicles. Therefore the impact of this sensor reduction on the estimation performance of the loosely coupled estimator must be studied. It is hypothesized that the AUNAV filter will be able to accurately estimate the sideslip and roll despite this sensor reduction. It is also expected that even without the pitch rate gyroscope the AUNAV estimator will be able to estimate the low frequency component of pitch. This is important since the road grade is the low frequency component and it has much higher amplitudes than the higher frequency suspension pitch changes. Both of these expectations are confirmed in Chapter 4. Additionally, the observability of the new system requires evaluation. This is done in Chapter 3. It is also important to note that because the pitch rate gyroscope bias is one of the states of the EKF in the loosely coupled filter, the number of states of the AUNAV

estimator is reduced by 1. The AUNAV estimator has 14 states, therefore all of the EKF matrices must be resized accordingly.

The second change that is made to the loosely coupled estimator is that a measurement of the vehicle's course angle from the GPS is added when the vehicle is driving straight. The loosely coupled estimator suffers from observability issues when there is little excitation present [19, 28]. Many authors seek to overcome the problem by adding velocity constraints [16, 18, 20]. They make the assumptions that the vehicle's lateral and vertical velocities in the vehicle frame are both zero at all times. These constraints are then added in the form of a virtual measurement update. Practically speaking, the latter assumption simply means that the vehicle is constrained to the road. The former assumption, however, is tantamount to assuming that there is no sideslip. This is acceptable for navigation purposes, but it will obviously not do given the explicit goal of estimating sideslip. So in order to add constraints that do not violate the estimation of sideslip, a virtual measurement of the course angle is created when the vehicle is driving straight. Recall the definition of vehicle sideslip given by Equation (1.1). If the vehicle is driving straight, and assuming that the side slope (bank) of the road is small, then the sideslip will be practically zero. If this is true, then from Equation (1.1) it is clear that the course ν and yaw ψ angles are equivalent. The course angle can be calculated from the GPS north and east velocities as shown in Equation (2.78).

$$\nu^{GPS} = \arctan\left(\frac{V_{east}}{V_{north}}\right)^{GPS} \tag{2.78}$$

When the vehicle is driving straight, the course angle measurement is used as a measurement of the yaw angle. The new measurement vector is shown below.

$$Y = \begin{bmatrix} \underline{r} \\ \underline{V} \\ \nu \end{bmatrix}_{GPS} - \begin{bmatrix} \underline{\hat{r}} \\ \underline{\hat{V}} \\ \hat{\psi}_{yaw} \end{bmatrix} \tag{2.79}$$

It is important that the filter *not* use the course measurement when the vehicle is turning, because the sideslip will cause the necessary assumptions to be violated, resulting in the course angle being *unequal* to the yaw angle. Using the course measurement under these circumstances will result in the filter falsely attributing the inevitable error to other sources, thereby corrupting the other state estimates. Therefore a means of switching is required for the filter to toggle between using the course measurement and not using it. This is done simply by setting the term in C corresponding to the course measurement to 1 or 0, depending on whether or not the measurement is being used. When the vehicle is turning, the course measurement is not used and the C matrix is defined as shown below.

$$C = \begin{bmatrix} I_{6x6} & 0_{6x2} & 0_{6x1} & 0_{6x5} \\ 0_{1x6} & 0_{1x2} & 0 & 0_{1x5} \end{bmatrix}_{7x14} \tag{2.80}$$

When the vehicle is driving straight, the course measurement is used and the C matrix is defined as shown below.

$$C = \begin{bmatrix} I_{6x6} & 0_{6x2} & 0_{6x1} & 0_{6x5} \\ 0_{1x6} & 0_{1x2} & 1 & 0_{1x5} \end{bmatrix}_{7x14} \tag{2.81}$$

Note that the column dimension of the C matrix is now 14 instead of 15, because of the reduction of the pitch rate gyroscope bias state. Adding the course measurement keeps the attitude errors bounded during periods of straight driving, as there is no lateral excitation with which to relate the north and east velocities to the accelerations.

It is not trivial to determine whether or not the vehicle is turning. At first glance, it may seem simple enough to use the steering angle. A simple law could perhaps be used that says if the steer angle is above a certain threshold, then the vehicle is turning. The problem with this is that if the wheels are misaligned, or if there is a large enough road crown (a slight road bank that improves water drainage), then the steer angle will have some constant offset from zero. The lateral accelerometer could also potentially be used, but it is potentially

subject to gravitational effects resulting from inaccurate roll angle compensation. In this thesis, the yaw rate gyroscope is used to determine whether or not the vehicle is driving straight. A flag is set if the yaw rate signal has been within 3 $^{deg}/_{s}$, consecutively, for a certain period of time. If this flag is "true", then the vehicle is assumed to be driving straight. If the absolute value of the yaw rate signal rises above the threshold even once, then the flag is reset to false. The detection logic requires that the absolute value yaw rate signal be below the threshold consecutively for a certain period of time in order to toggle the flag to true (driving straight), as opposed to simply setting the flag to true upon the first yaw rate signal below the threshold. This is done to avoid the problem of zero crossings of the yaw rate during a turning maneuver. The yaw rate will be zero for a short period of time, for example, during a sinusoidal steering maneuver. The vehicle is still turning in this case, therefore it would be an error to assume straight driving conditions. The time requirement on the detection logic helps mitigate this effect. There is no corresponding time requirement to toggle the flag back to false (turning). This is in order to be conservative, so as not to introduce error into the states by assuming that there is no sideslip when in fact there is. The threshold of 3 $^{deg}/_{s}$ was chosen as a conservative value, because experimental tests showed that sideslip was generally extremely close to zero for yaw rates below this.

There are several problems with the current method of turning detection. First, the effects of the yaw rate bias need to be studied. For now, the threshold is simply set large enough to accommodate the bias. The bias on the gyroscope used is quite small (on the order of hundredths of a degree per second), so this is not a major problem. For lower quality gyroscopes, however, it could be. There are potential problems as well with using the yaw rate gyroscope bias estimate to try to mitigate the problem. For example, it could result in an unstable feedback type of situation, where a large initial error in the bias estimate causes yaw rate signal to always fail the test, thereby impeding the estimation of the bias. The other problem is the sensor noise. If the noise on the gyro is too high it becomes difficult to estimate slowly growing sideslip, because raising the detection threshold will mean that low rates of

turning will fall into the zero sideslip assumption. Setting the threshold correctly is also important. It makes sense to set it to some integer multiple of the noise standard deviation. The problem is that drivers rarely drive perfectly straight, and setting the threshold in this way would only trigger the yaw course measurement under perfectly straight conditions. In effect, the threshold becomes a trade off between the sensitivity of the estimator and how often the course measurements are applied. Furthermore, the time requirement is a tunable parameter. Setting the value too low means that zero crossings which occur slowly will result in the estimator assuming zero sideslip in the middle of the turn. Setting the value too large limits the instances where the valuable course information is used. Therefore both the signal threshold and the time threshold are important parameters of the overall AUNAV system which must be tuned carefully.

2.3.3 Conclusion

In this chapter some fundamentals of GPS systems, inertial measurement units, and Kalman filtering were discussed. Following this, the algorithm for the modified modular estimator was presented. In contrast to the AUNAV estimator, the MME algorithm consists of several distinct Kalman filters, each estimating states along different axes. Therefore there is qualitatively less coupling between the states of all of the different filters in the MME when compared with the coupling between the states of the AUNAV estimator. The MME is a modified version of the estimator presented in [6]. The modification consists of the removal of the pitch rate gyroscope. The pitch angle estimate in the MME comes from the road grade estimate, which is a function of the vertical and horizontal velocity estimates.

Finally, the development of the algorithm for the AUNAV estimator was presented. The AUNAV estimator is developed by making two important modifications to a loosely coupled GPS/INS filter. The classic loosely coupled algorithm is first discussed, followed by a discussion of the modifications. The first modification is the removal of the pitch rate gyroscope (and its corresponding bias state in the EKF). This is necessary because pitch

rate gyroscopes are not currently available on commercial vehicles. The second modification is the GPS course measurement which is used when the vehicle is driving straight. This is necessary to improve the observability when the vehicle is driving straight, because there is no lateral acceleration with which to relate heading errors to velocity errors. Turn detection is done using the yaw rate gyroscope.

Chapter 3

Observability of the AUNAV Estimator

3.1 Definitions

The notion of observability for a particular system is obviously of great importance for estimation algorithms. The observability of a system indicates whether or not it is possible to estimate the states of the system based on the given sensor configuration. As an overly simplistic example, consider the problem of determining the direction that a vehicle is facing (also known as vehicle heading). Certainly this state could not be determined if the only sensor available is a thermometer in the cabin of the vehicle. This is obvious, because the cabin temperature information is utterly unrelated to the vehicle's heading. Yet what if GPS north and east velocity measurements were available? This information is much more related to heading, but is it enough, and under what conditions? These questions of relating the information from the sensors to knowledge of the desired states are at the heart of observability. For linear, time invariant, deterministic systems, observability is simply a function of the pair of the system dynamics matrix and the measurement matrix. Changing the measurement matrix, which implies changing the sensor configuration, or changing the system dynamics matrix, which defines how the states are related to one another dynamically, will affect the observability properties. This is intuitive. For linear time-varying or non-linear systems things become more complicated. Let's begin the discussion by focusing first on linear, deterministic systems which are not time invariant. The problems imposed by non-linearity and by stochastic influences will be discussed afterward. A formal definition of observability is given by [43].

Definition 1. *A system is said to be observable if the initial state $x(t_0)$ can be determined from the output $y(t)$ over the finite time interval $[t_0, t_f]$.*

Definition 1 assumes that the input $u(t)$ is known. This definition serves as a general definition of what is meant by "observability" for all types of systems, yet it is far from the only definition in the literature. In fact the researcher cannot be too careful when considering the terminology of observability, as more specific definitions and conditions also abound. Silverman and Meadows offer three more discriminating definitions of different types of observability for linear time variant systems [42]. They are *complete*, *total*, and *uniform* observability. It is worth noting, as an aside, an example of inconsistent terminology found in the literature. In [35] the authors, while citing Silverman and speaking of the exact same three observability definitions, refer to them as *complete*, *differential*, and *instantaneous* observability. In this work, the author adopts the terminology of Silverman. Here, complete observability carries the same definition of observability as Definition 1. Total and uniform observability are defined as follows [42].

Definition 2. *A system is said to be totally observable on an interval $[t_0, t_f]$ if it is completely observable on all subintervals of $[t_0, t_f]$.*

Definition 3. *A system is said to be uniformly observable on an interval $[t_0, t_f]$ if the matrix $Q_0(t)$ is full rank for all t on $[t_0, t_f]$.*

Definition 3 is the strongest of the three conditions. The matrix $Q_0(t)$ is defined as follows. If a linear, continuous system is given by (3.1), then the matrix $Q_0(t)$ is defined as in (3.2).

$$\begin{aligned} \dot{x}(t) &= A(t)\,x(t) + B(t)\,u(t) \\ y(t) &= C(t)\,x(t) \end{aligned} \tag{3.1}$$

$$
\begin{aligned}
Q_0(t) &= \begin{bmatrix} S_0(t) & S_1(t) & \cdots & S_{n-1}(t) \end{bmatrix} \\
S_{k+1}(t) &= A'(t) S_k(t) + \dot{S}_k(t) \\
S_0(t) &= C'(t)
\end{aligned} \tag{3.2}
$$

The general observability condition (Definition 1) is a necessary one for estimation purposes. For systems which are linear, deterministic, and time invariant, it is also sufficient, but this is not true of the loosely coupled filter. It is important to remember moving forward that proofs of convergence are not in view, rather the focus is on verifications that the system meets the minimum condition for estimation. It will be seen that many times, the loosely coupled filter does not. It will now be considered how these definitions have been applied in the literature to the problem at hand, namely, the loosely coupled GPS/INS filter.

3.2 Applications in Literature

In [21] Goshen-Meskin and Bar-Itzhack develop a theoretical method of analyzing the observability of linear time varying systems by considering them as piece-wise continuous systems. They first show that often times it is valid to consider a time varying system as a sequence of consecutive time invariant systems if the time varying system meets certain conditions. This allows the authors to view the observability of the overall system as a simpler function of the observability of each time segment. This idea is then applied in [22] to the in flight alignment (IFA) problem for INS units in aircraft. Various maneuvers are reduced to distinct and consecutive time segments of constant systems, such as a constant acceleration or a constant radius turn, and the overall observability is studied. In this way it is possible to consider whether certain maneuvers will result in an observable system. Said another way, it makes it possible to theoretically decide what is an effective sequence of motions for IFA. They conclude that, for the 12 state loosely coupled GPS/INS, any first segment has an

observability matrix of rank 9. Adding another distinct segment increases the rank to 11, and any third increases it to full rank. The order of the segments does not matter, and repeating segments has no effect. All segments in that work consist of distinct, constant accelerations in either the north, east, or down directions. This analysis is very similar in approach to most other works in the literature. That is, researchers are not evaluating the observability of the "system", but rather asking whether or not the system is observable along certain trajectories. The fact is, this system and those like it are neither linear nor time invariant, so the analysis is restricted to local observability analyses of individual trajectories. Yet this a very useful endeavor as an analysis or design tool, allowing the designer to draw conclusive and defined operating boundaries outside of which the filter can be guaranteed *to fail*. Since the system is actually non-linear, the results for the local trajectory represent a *necessary* condition, not a *sufficient* condition. This analysis of the observability of the nonlinear system along a certain trajectory will be referred to hereafter as the observability "of the trajectory" for simplicity.

Another important work is that of Rhee *et al.* [35]. Here the authors look at the observability of several trajectories using the definitions and conditions defined in [42]. First they look at the case of constant linear acceleration, treating it as a time invariant system as in [22]. Using well known observability tests such as the Hautus test and the standard LTI observability matrix, they find that the total system is unobservable and that the attitude angles are the unobservable modes. This brings up a very important discussion. In this case, Rhee *et al.* found that the observability matrix was rank deficient by 3. This means that there are 12 observable *modes*, which is not the same thing as 12 observable *states*. Modes may be states or linear functions of states. In this case, the 3 unobservable modes correspond directly to states (the 3 attitude angles). However, 6 of the 12 observable modes do not correspond to individual states, rather they are a function of other states (specifically the attitudes and inertial biases). While there are 12 observable modes, there are only 7 observable states. These are position, velocity, and the vertical accelerometer bias. The remaining 5 observable

modes are functions of the other 8 unobservable states. This means that the *combination* of the remaining attitudes and inertial biases is observable, but that the filter is unable in this case to separate them individually. This is consistent with [22], and is basically equivalent to applying their test to just the one acceleration segment alone. In short, the filter cannot distinguish between the attitudes and sensor biases (except for the vertical accelerometer bias); however the combination of the attitude errors and the sensor biases is observable. Yet it is extremely important for the filter to be able to distinguish sensor biases from attitude, therefore knowing the circumstances under which the attitude and bias estimates cannot be independently observed is crucial. Rhee also considers the uniform observability of the case of non-constant linear acceleration or constant rotation. This is done by applying the definition of uniform observability (Definition 3) found in [42]. In doing so they find that maneuvers of either type increase the number of uniformly observable modes by two. For the case of non-constant axial acceleration, two attitude angles are made observable. The attitude angle around the jerk vector is the state which remains unobservable. For the constant turn case, the authors find that the yaw angle remains unobservable. In each case, the attitude angles which are made observable by the maneuver might also cause the biases to become independently observable as well. Recall that for the case of no excitation (driving straight) there are functions of the attitudes and biases which are observable, even though the attitudes and biases aren't observable independently. Making the attitude angles observable effectively decreases the number of unknowns in the equations, making the biases observable also. This makes sense intuitively. Consider the lateral accelerometer model, Equation (2.47). If the only biasing effects are from the roll angle ϕ and the sensor bias δf_y, and if the roll angle is observable or known, then it follows that the bias likewise is observable.

Uniform observability (Definition 3) is the strongest condition of the three, and the authors do not address the general observability condition (Definition 1) for the time variant case. This makes it difficult to quantitatively compare these results with others in the

literature. Hong *et al.* [26] also examines the single antenna loosely coupled GPS/INS filter. They find that the attitude and bias states are unobservable if the system can be represented as being time-invariant (i.e. when only undergoing constant axial acceleration). Furthermore they conclude that linear acceleration changes enhance the overall observability of these states. These findings are qualitatively consistent with those previously mentioned, and they also align with the intuitive presentation of observability given in [19].

3.3 Observability Simulations

The above findings in the literature paint a clear picture of the observability of the loosely coupled filter concerning most of the relevant trajectories and operating conditions. However these findings are all for the standard algorithm, and the observability of the modified algorithm in this work still needs to be considered. The modifications to the filter include removing a state (the pitch rate gyro bias), an input (the pitch rate gyro), and adding a measurement (the yaw/course measurement). Adding the measurement of yaw will certainly improve the observability, and the degree to which it does so will be shown in this section.

When considering the observability of the modified algorithm, the general definition of observability given by Definition 1 is used. As a means of evaluating a system's observability, Stengel [43] gives an equation for the observability matrix of a linear time-varying system.

$$O_{LTV}\left(t_f, t_0\right) = \int_{t_0}^{t_f} \Phi^T\left(\tau, t_0\right) H^T H \Phi\left(\tau, t_0\right) d\tau \tag{3.3}$$

The system is observable if $O_{LTV}\left(t_f, t_0\right)$ is non-singular, where $O_{LTV}\left(t_f, t_0\right)$ represents the observability matrix for the system on the interval $[t_0, t_f]$, $\Phi\left(\tau, t_0\right)$ represents the state transition matrix from t_0 to t_f, and H represents the measurement matrix. The state transition matrix is obtained by multiplying in series the state transition matrices of each discrete time step.

$$\Phi(k,0) = J_d(k)J_d(k-1)\cdots J_d(0) \tag{3.4}$$

The state transition matrix of each time step is obtained by discretizing the Jacobian of the system dynamics. This is done according to (3.5) using the matrix exponential function provided by Matlab, *expm()*

$$J_d = e^{(J*T_s)} \tag{3.5}$$

where T_s represents the sample time. The AUNAV estimator is a nonlinear system, therefore it bears restating that the observability results are local to the trajectory about which the system is linearized. It is also important to note that for the observability analysis the system is linearized not about state estimates, but about the *true trajectory*. This is because the problem is to determine whether or not a certain trajectory (maneuver) theoretically results in an observable time varying system. The modified algorithm is investigated by analyzing the rank of the matrix $O_{LTV}(t_f, t_0)$. If it is full rank, then it is also non-singular and the trajectory of interest is fully observable on the time interval $[t_0, t_f]$. If the trajectory is observable, then the filter will converge if it is tuned appropriately. The analysis from the literature provides further insight into the case where $O_{LTV}(t_f, t_0)$ is rank deficient. If we first analyze the modified filter without including the yaw constraint (i.e. the standard loosely coupled filter without having the pitch rate gyro or the pitch rate bias state), we should expect to see results in accord with those discussed in Section 3.2. Several simulations were performed in Carsim to validate this expectation. A note here regarding the following results is necessary. The following plots show the rank of $O_{LTV}(t_f, t_0)$ over the course of the simulations. In general, the rank test of $O_{LTV}(t_f, t_0)$ shows whether or not the system is observable *on the interval* $[t_0, t_f]$. This means that if a maneuver is performed which increases the rank to full, during some time interval, then the rank will thereafter remain full and the system will be declared observable until the end of that interval. What is really

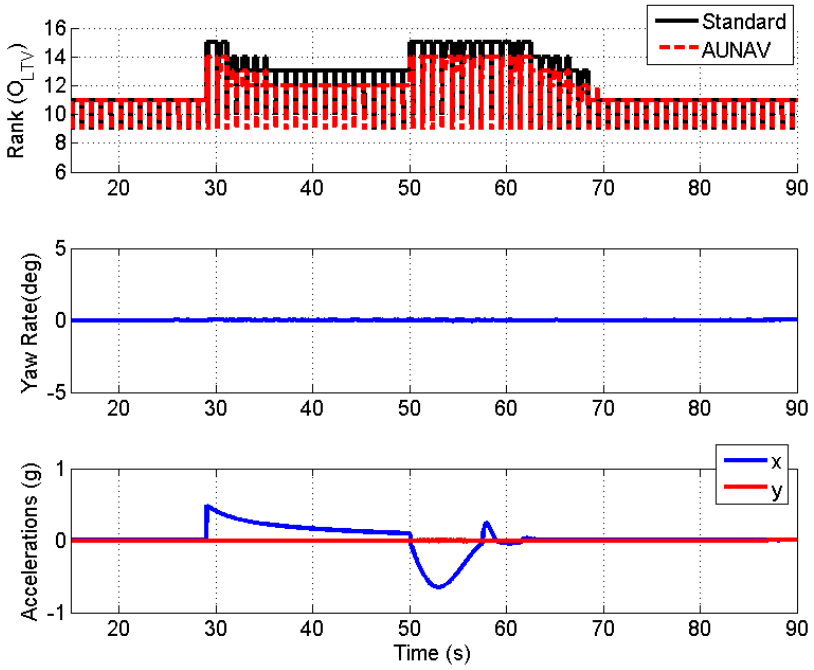

Figure 3.1: Observability of the Loosely Coupled and AUNAV Filters During Longitudinal Dynamics.

of interest here is to show the impact of various maneuvers regarding observability. Because of this, small intervals of one second were chosen to evaluate O_{LTV}.

Two simulations were done in Carsim in order to test the observability of each trajectory. All values from the simulations are true values, no noise or other errors were added, because the analysis is done concerning the linearization about the *true* trajectory. The first test involves the vehicle driving in a straight line with a period of forward acceleration followed by deceleration. Figure 3.1 shows the rank of O_{LTV} during this test, with subplots showing the yaw rate and accelerations. Full rank is 15 for the standard loosely coupled filter and 14 for the modified filter. Figure 3.2 shows the results from a simulation in which the vehicle drives straight, enters a steady state turn, and resumes straight driving. The rank of O_{LTV} during these tests confirms the findings in the literature. Specifically, the system is not fully observable under constant acceleration, but requires changes in acceleration to reach full observability. It can be seen in Figure 3.1 that as the vehicle begins to accelerate the system

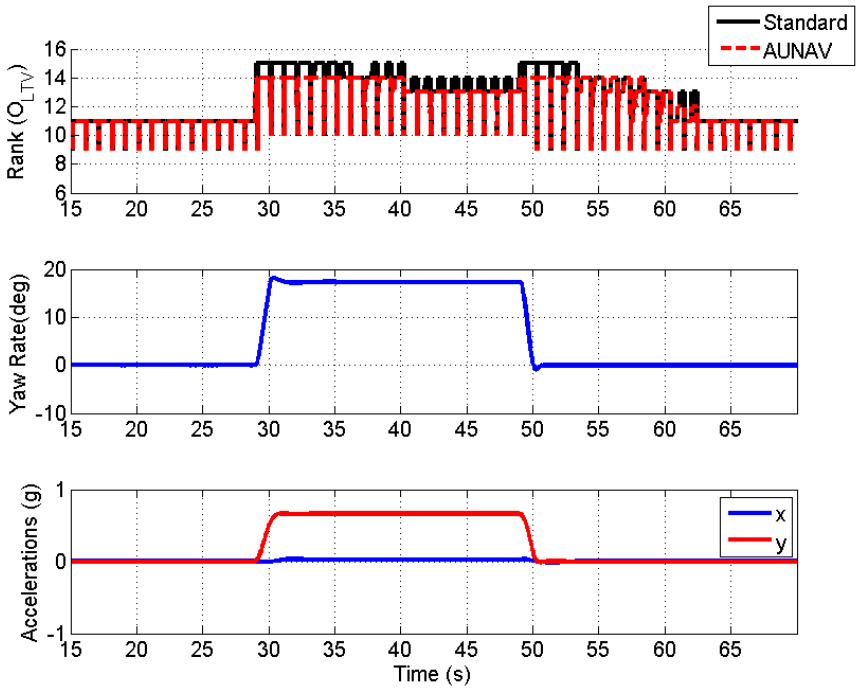

Figure 3.2: Observability of the Loosely Coupled and AUNAV Filters During Lateral Dynamics.

becomes observable, but as the acceleration becomes steady the rank drops back to 11. The same is seen when the vehicle decelerates. Figure 3.2 confirms this also, showing the same behavior in regard to lateral acceleration. Under constant acceleration, the observability matrix is rank deficient by three. Constant acceleration in another direction increases the rank by two, confirming the results in [21]. The filter only reaches full rank when a change in acceleration occurs, confirming the results in [35]. These results show that the modified system, without the extra yaw constraint, behaves just like the standard system in regards to observability. Therefore the conclusions of Section 3.2 apply to the modified filter when the yaw constraint is not applied. This gives some insight into the anticipated performance of the filter during constant acceleration operation (straight steady driving). Rhee found that in this situation the attitude and bias states are unobservable [35]. This means that the estimates will likely be biased, yet because the *combined effect* of the leveling angles and

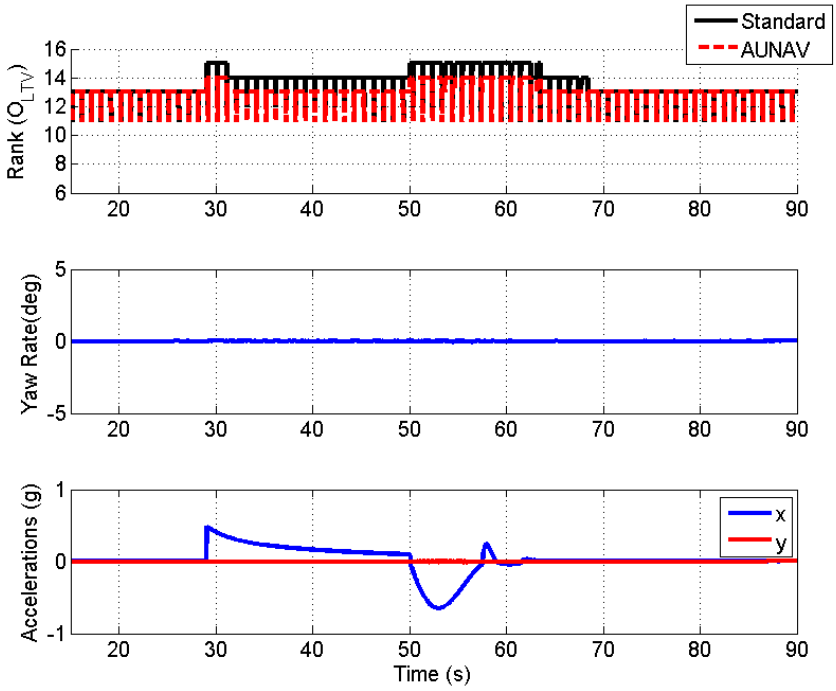

Figure 3.3: Observability of the LC Filter Compared with the AUNAV Filter During Longitudinal Dynamics when the Course Measurement is Conditionally Added to the AUNAV Filter.

the accelerometer biases is observable, the errors will be bounded. This cannot be said of the yaw angle, which is known to exhibit drift during this time.

Figures 3.3 and 3.4 show the results from the same two simulations, only this time the yaw constraint is imposed when the straight driving condition is met. There is no difference in the rank of O_{LTV} under dynamics, but the rank is closer to being full during straight driving. In this case the observability matrix is only deficient by one, as opposed to being deficient by three. The improvement of two more observable states is easily explained. During this time, the filter has direct yaw information available (because sideslip is assumed to be zero), which makes the yaw angle observable. Furthermore, since the integral of the yaw rate gyroscope is measurable, the bias of this gyro becomes observable. In this way the modification overcomes the problem of drifting yaw estimates.

While the previous simulation results show when the filter does and does not meet the necessary conditions for estimation, the convergence of the filter still requires investigation.

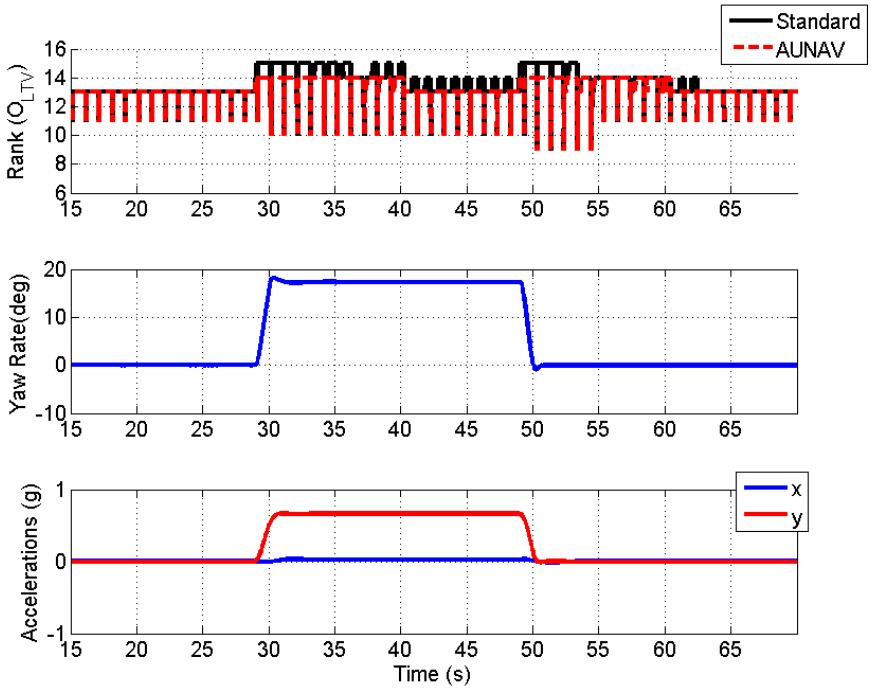

Figure 3.4: Observability of the LC Filter Compared with the AUNAV Filter During Lateral Dynamics when the Course Measurement is Conditionally Added to the AUNAV Filter.

If the system were in fact linear and deterministic, observability would imply convergence. However large initial errors or unmodeled disturbances could cause convergence to an incorrect local minimum. What's more, it has been shown in [5] that the stochastic elements by themselves can cause the filter to diverge if the filter is improperly tuned. The authors there describe the distinction that results between standard observability and "stochastic" observability. The system can meet the necessary conditions for observability and yet diverge due to poor tuning of P, Q, R, due to too large values of P_0, or due to large errors in the initial estimates. So filter convergence of the nonlinear estimator can only be expected if the trajectory is observable, the tuning is appropriate, and the initial error is not too large.

The following simulations demonstrate the behavior of the filter under such conditions. These simulations consist of the combination of the two prior simulations. That is, the vehicle enters a steady state turn, disengages from the turn, and accelerates and decelerates while driving straight. The simulation is run first without the yaw constraint, followed by

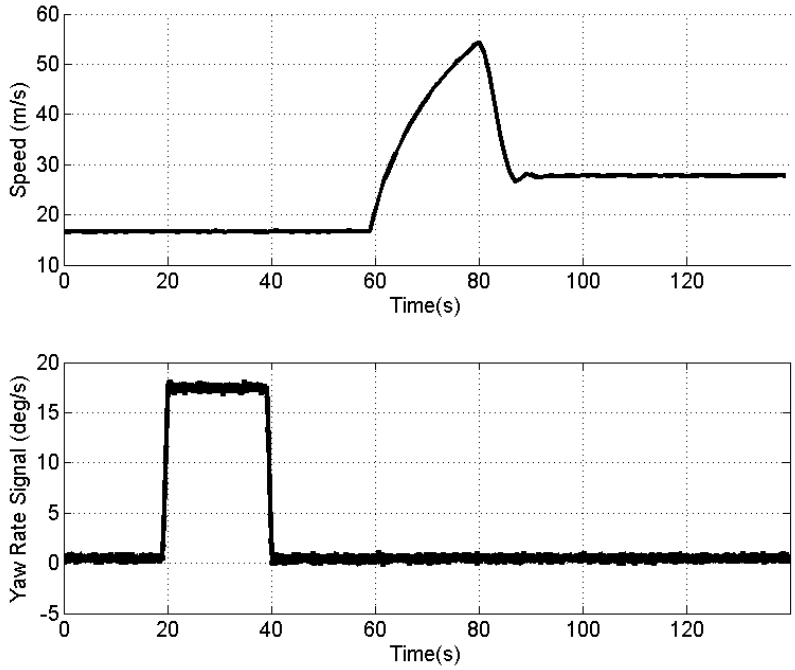

Figure 3.5: Speed and Yaw Rate Profile of Convergence Test Simulation.

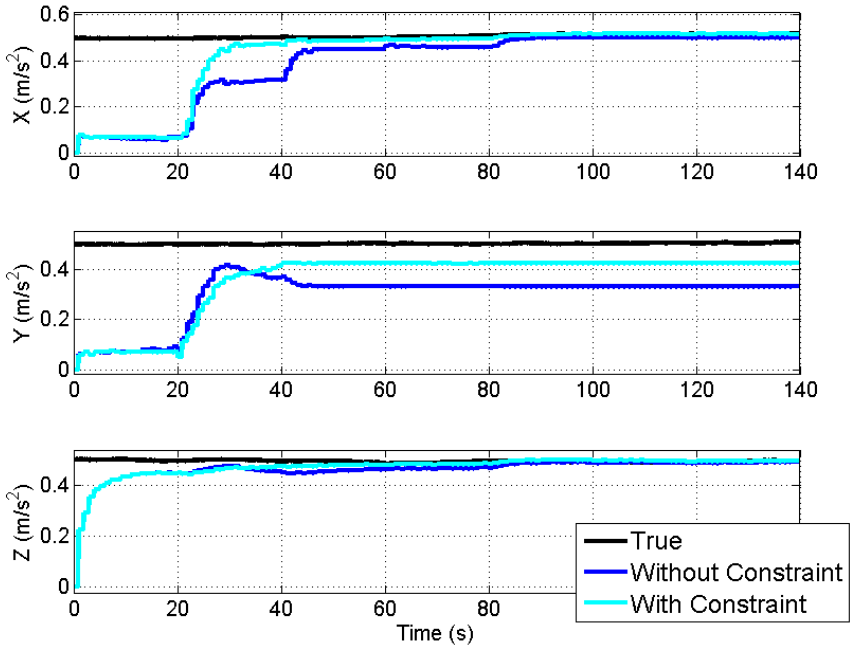

Figure 3.6: Convergence of the AUNAV Accelerometer Bias Estimates During Simulation.

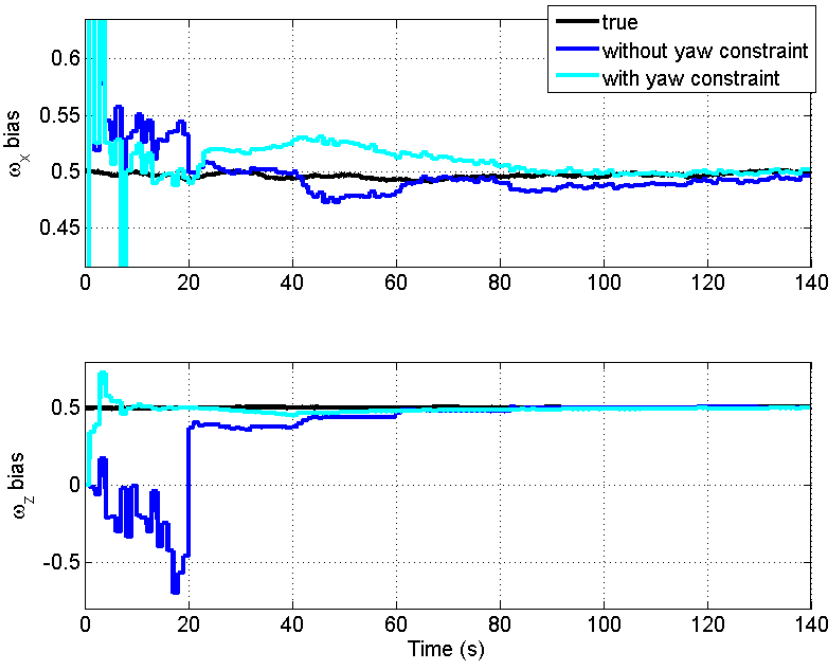

Figure 3.7: Convergence of the AUNAV Gyroscope Bias Estimates During Simulation.

a run with the constraint. Figure 3.5 shows the speed and yaw rate for the simulations. Figures 3.6-3.11 show the estimates of the bias states and the leveling angles, where the blue signal represents the estimates without the yaw constraint and the cyan represents the estimates with the constraint.

Let's first consider the case without the yaw constraint (blue). It can be seen in Figure 3.6 that the x and y accelerometer bias estimates do in fact converge toward the true value during periods where the observability test is full rank (during acceleration changes). The y accelerometer bias shows in particular that bias convergence is strongest when the acceleration change is along the axis collinear with that bias. It is also observed that the z axis accelerometer bias converges regardless of the dynamics. Both of these results are in accord with the results in [22], [26] and [35]. Figure 3.7 shows the behavior of the gyroscope biases. The yaw rate gyro bias (ω_z bias) also behaves as expected, converging toward the true value during changes in acceleration. By contrast, the roll rate gyroscope bias (ω_x bias) converges rapidly, regardless of the dynamics. This result is in disagreement with the results in [35].

There the authors find that when there is no excitation (straight driving at a constant speed), an unobservable mode is given by

$$x_{3x} = \delta\omega_p - \Omega_z\delta\theta + \Omega_y\delta\psi \tag{3.6}$$

where x_{3x} is the particular unobservable mode in question (using the author's notation), $\delta\omega_p$ is the roll rate gyroscope bias, and $\delta\theta$ and $\delta\psi$ are the pitch and yaw errors respectively (using the notation of this thesis). The terms Ω_z and Ω_y relate to the Earth's rotation rate, see [35] for details. It is observed that if the second two terms in Equation (3.6) are small, then

$$x_{3x} \approx \delta\omega_p \tag{3.7}$$

Since x_{3x} is an observable mode even with no excitation, then the roll rate gyroscope bias is observable under these conditions. While in reality it is the sum of the bias and the Earth rotation terms which is observable, the rotational terms are small enough compared with the bias that they can be disregarded. Carsim does not include any simulated effects from the Earth's rotation, which explains the convergence in simulation. In order to verify that the bias estimate converges using the true sensors, experimental data was analyzed. In this experiment, the vehicle is driven on a mostly straight road at close to a constant speed. Details of the sensors and vehicle are given in Chapter 4. Figure 3.8 shows the roll rate gyroscope bias estimates. Analysis of static data for this run reveals a bias of 0.139 $^{deg}/_{sec}$. This was found by taking the mean of the roll rate signal over 25 seconds of static data (collected at 100Hz). Is is observed that even during periods of very little excitation, the bias estimate converges to very near 0.14 $^{deg}/_{sec}$, confirming the simulation results and the hypothesis that the Earth rotational terms are small enough to neglect. Figure 3.9 is included to show the forward speed and yaw rate, showing that apart from the period $500s < t < 700s$ there is very little excitation.

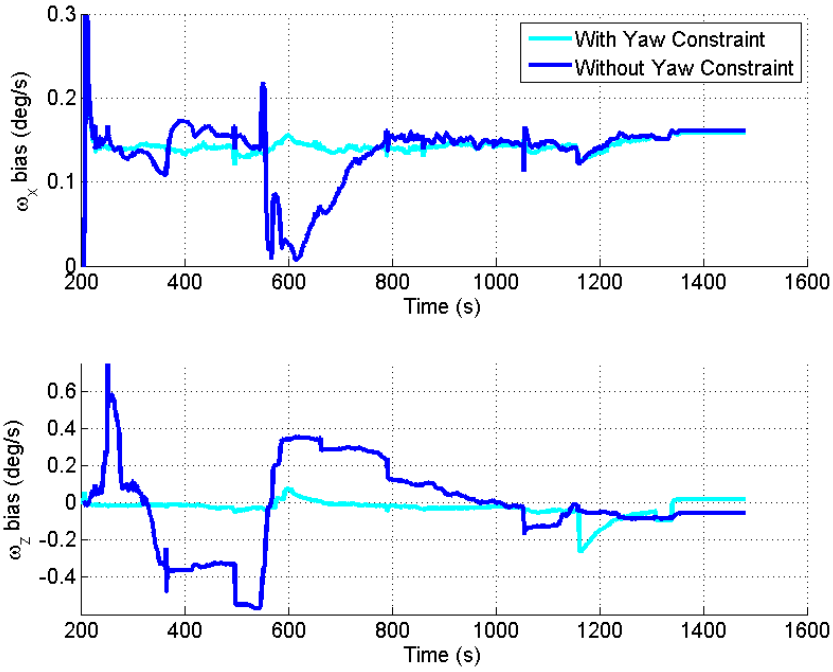

Figure 3.8: Convergence of the AUNAV Gyroscope Bias Estimates During Experimental Testing with Limited Dynamics.

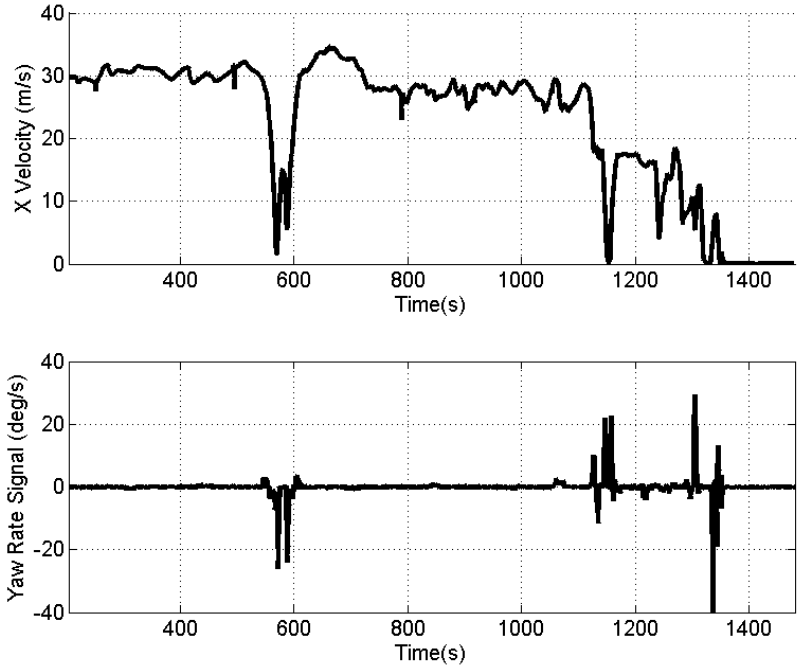

Figure 3.9: Velocity and Yaw Rate from Experimental Testing.

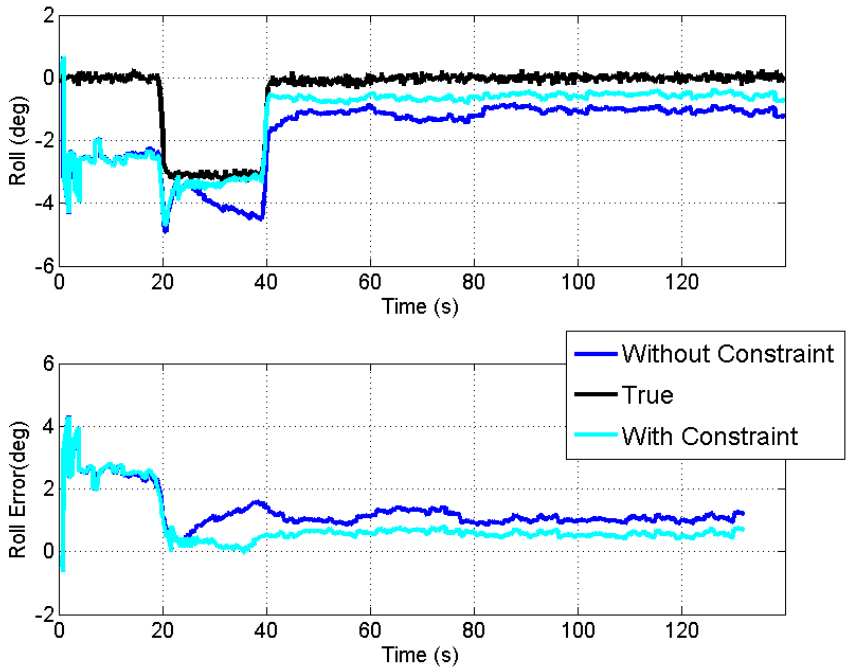

Figure 3.10: AUNAV Roll Angle Estimate Convergence During Simulated Test.

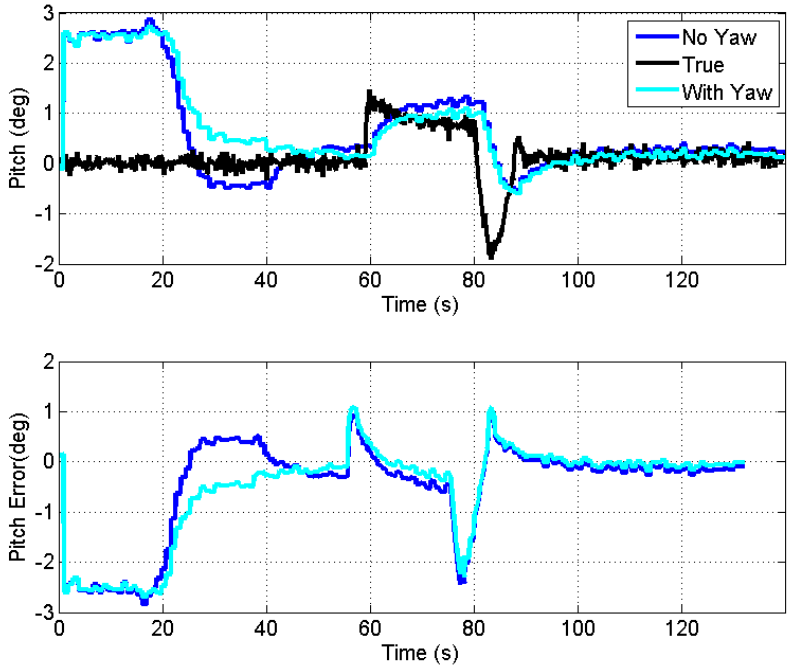

Figure 3.11: AUNAV Pitch Angle Estimate Convergence During Simulated Test.

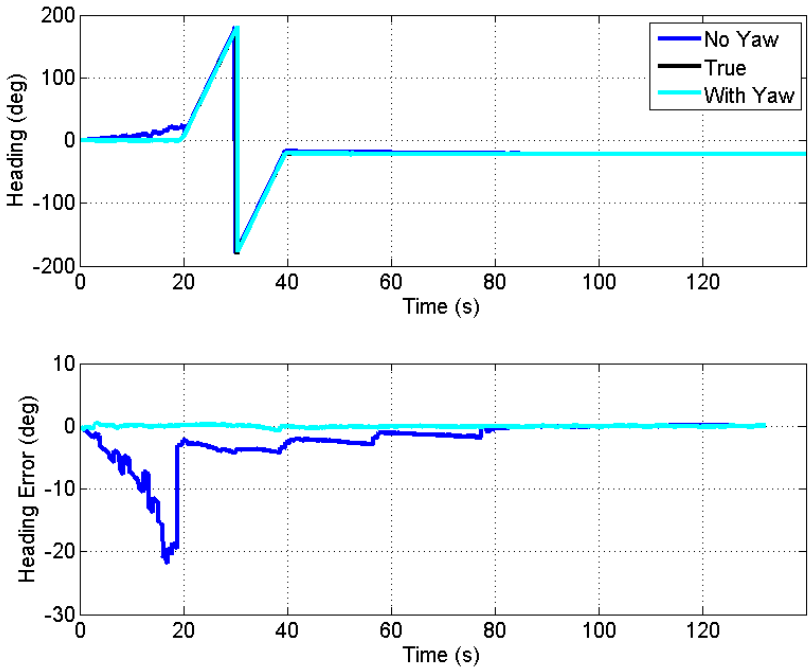

Figure 3.12: AUNAV Yaw Angle Estimate Convergence During Simulated Test.

Since the combined effect of the leveling angles and the x and y accelerometer biases is observable, it is expected that the leveling angle estimates would converge as the bias estimates converge. In particular, it is expected to see this kind of coupling between the roll angle estimate and the y accelerometer bias, and between pitch angle and the x accelerometer bias. Figures 3.10 and 3.11 show that this is the case. The roll angle, like the y accelerometer bias, converges most rapidly during changes in lateral acceleration, while the pitch angle converges most rapidly when its bias counterpart does as well. Let's now consider the simulations which include the yaw constraint (cyan). It can be seen in Figures 3.6 and 3.7 that the x and y biases show better convergence this time, as do the roll and yaw gyro biases. In fact convergence for the yaw gyro bias is no longer a function of the dynamics, as expected. A slight improvement is seen in the roll angle, while no improvement can be detected in the pitch angle. Figure 3.12 offers a comparison of the yaw angle estimation performance with and without the yaw constraint. It can be seen that the yaw angle does in fact drift when the constraint is not imposed and the yaw gyro bias remains unobservable.

This drift diminishes as dynamics occur to enable an accurate estimation of the bias. Despite this, the performance with the yaw constraint is far superior, showing no drift whatsoever during the run. However, a very slow drift would occur during a sustained steady state turn.

3.4 Conclusion

Several key states are unobservable for the loosely coupled filter when the vehicle is undergoing limited dynamics. While this represents the chief limitation of the filter, it is far from a death blow to its potential. The unobservable states in question are the attitude states and the bias states. The bias states by nature change very slowly over time, relative to changes in the level angles, making it possible to broadly define two general operating scenarios of the filter: an "initialization" stage and a "tracking stage". The initialization stage represents the period before the bias estimates converge to near the true values. For the purpose of illustration, the biases could be thought of as static during this phase, and once they are acceptably identified the attitude and sideslip estimates become trustworthy. This marks the transition into the tracking phase. During this phase, since the bias estimates are close enough to the true values, the filter outputs good attitude and sideslip estimates. Furthermore, the slowly changing bias can now be viewed basically as a slow drift with a small initialization error. Normal driving conditions *likely* provide enough dynamics to track these small changes, thereby keeping bias estimate errors from growing dangerously large. This is a good avenue for future research.

It is also important to reconsider a result detailed in [35] regarding the difference between the observability of a function of several states which are not independently observable. The authors showed that in the scenario of constant acceleration the combination of the attitude and the inertial sensor biases remains observable despite the fact that these states are not *independently* observable. This means that if the filter has already entered the "tracking phase", meaning it has already identified the biases, then it follows that the attitude states will also be known, even though they are technically unobservable. Said another way, once

the biases have been identified, there is no danger filter from the vehicle entering into the unobservable scenario of simple straight driving for a period of time. The amount of time for which the attitude will remain known is a function of how quickly the sensor biases drift relative to how much dynamics the vehicle experiences. It is also a function of tuning. Poor tuning could cause the bias estimate to change rapidly, even though the true bias does not. From this it is possible to infer a problematic scenario. If the vehicle is traveling cross country through a region like Texas, there will likely be very prolonged periods of driving straight without lane changes, braking, forward accelerations (or anything even remotely interesting in general) . Such a scenario would likely send the filter back to the initialization phase. However it may be possible to consider additional constraints, such as a flat road assumption, to mitigate this effect and perhaps even prevent it. It may also be possible to use the drift rate characteristics of the IMU to quantify how long it is safe to remain in the unobservable state once the biases have been identified.

Chapter 4

Experimental Validation of the AUNAV Estimator

In order to validate the algorithm, tests were performed in a 2007 Infiniti G35 at Auburn University's National Center for Asphalt Technology, or NCAT. NCAT offers a couple of valuable resources to this project. The first is a skid pad area for high dynamic testing. This region is a large, approximately flat, section of asphalt where maneuvers such as double lane changes or J-turns are safe to conduct. The other, primary resource of NCAT is the 1.7 mile oval loop. The oval has two straight segments and two segments of banked (8 degree) turns. The track provides the ability to collect data on longer experiments which is especially important for bias estimation validation.

The G35 is outfitted with several sensors including a Novatel Propak V3 antenna and receiver, a Septentrio PolaRx2 three-antenna GPS attitude determination system, a Crossbow IMU 440, and a CAN reader which provides access to the vehicle's on board wheel speed and steer angle signals. Although the Crossbow is a six degree of freedom IMU, the pitch rate gyro is not used in this work. The Crossbow is of higher quality than automotive inertial sensors, and this difference is discussed later in this chapter. The Novatel receiver is capable of receiving differential corrections, however corrections were not used in order that the experiment might more closely resemble the sensor suite of a market vehicle. The Novatel and Crossbow are the sensors used for the estimator, while the Septentrio provides reference "truth" values for sideslip, roll, pitch, and yaw. The Novatel provides position and velocity measurements, while the Crossbow provides accelerations and rotation rates in the frame of the IMU. Sideslip is obtained from the Septentrio via its velocity and heading solutions. Equation (4.1) is used to calculate the reference sideslip

$$\beta = \arctan(\frac{V_y^{Sep}}{V_x^{Sep}}) \tag{4.1}$$

The main Septentrio antenna is located near the rear axle of the vehicle (along the X axis). Therefore V_y^{Sep} must be translated to the Novatel antenna location via Equation (4.2).

$$V_y^{Sep} = -V_n^{Sep}\sin(\psi) + V_e^{Sep}\cos(\psi) + \Delta_x\omega_r \tag{4.2}$$

In this equation Δ_x represents the distance in the body frame from the Novatel antenna to the Septentrio main antenna. There is a slight yaw angle alignment difference between Septentrio and the body frame. This is found by comparing the course measurement from the Novatel, which is an unbiased measure of the velocity vector, and the yaw measurement from the Septentrio. If the Septentrio is aligned perfectly, then the yaw measurement should match the course measurement when the vehicle is driving straight. It was found that there is a difference of one half of a degree between the two, and the Septentrio heading measurement was adjusted accordingly. Since the sideslip is a function of heading, this adjustment also pertains to the sideslip. All plots of the Septentrio sideslip are of the adjusted values.

The conclusions from Chapter 3 motivate two separate experiment scenarios: one for the "initialization" phase and one for the "tracking" phase. The tracking phase is mainly concerned with state estimation under ESC type conditions, which by nature last over short time durations, while the initialization phase is concerned with long term bias estimation.

4.1 Initialization

Figures 4.1 and 4.2 show the trajectory, speed, and yaw rate of the initialization experiment. The purpose of this test is to validate the results of the observability theory and show that under sufficient dynamics the filter can estimate both the attitude and the bias states accurately. Two scenarios were created during this test in order to show this. One scenario has minimal dynamics and the other has exciting dynamics. This was done by varying how

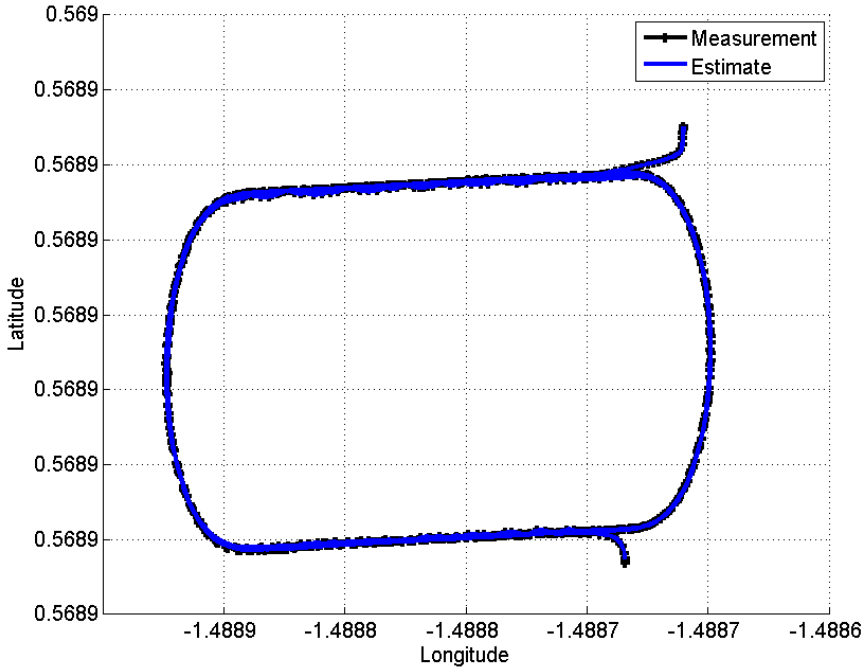

Figure 4.1: Vehicle Trajectory During Initialization Experiment.

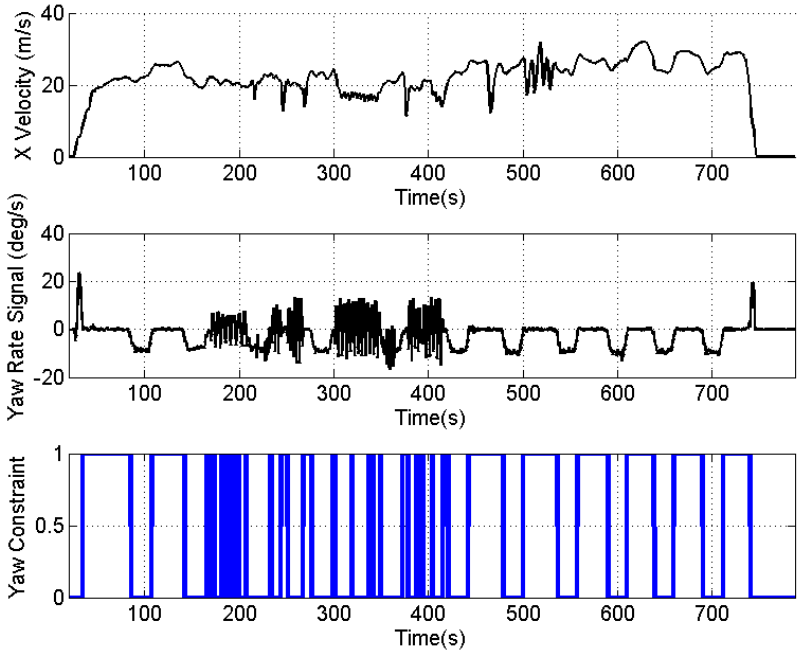

Figure 4.2: Profile of Dynamic Conditions During Initialization Experiment.

the vehicle was driven during the straightaways. The vehicle was driven normally around the turns every time. During several laps at the beginning and the end of the experiment the vehicle was driven straight and at a nearly constant speed, as much as possible, on the straightaways. These laps serves as a comparison for the behavior of the filter under dynamics. The dynamic laps occur during the time $t \approx 200 - 400s$. Here a sinusoidal steering input was applied on the straight sections. Intermittent braking and acceleration were also applied, but these were relatively fewer and less exciting than the steering inputs. The yaw rate signal in Figure 4.2 shows the lateral excitation, while the spikes in the velocity signal show the braking and acceleration inputs.

It is necessary to show that the AUNAV estimator can accurately estimate the accelerometer and gyroscope biases. However, obtaining a "true" value for the accelerometer biases is difficult. An accurate value for gyroscope biases can be obtained by analyzing static data from either before or after an experiment. This is because, as shown in Equation (2.16), the sensor bias is the only biasing element corrupting the sensor. It can therefore be isolated by taking the mean of the static data. The difficulty for the accelerometers lies in the fact that the gravitational effects appear exactly like biases. Simply taking the mean of static data would include the effects of any roll and pitch. The Septentrio can be used to estimate the accelerometer biases, because the attitude measurements which it provides cause the AUNAV estimator to be fully observable. This is precisely how the "reference" values of the accelerometer biases were obtained. The attitude measurements from the Septentrio were included in the AUNAV estimator for the purpose of obtaining the reference accelerometer biases. The addition of the attitude measurements was only for this purpose; the AUNAV estimator does not use these measurements in general. One more detail should be noted regarding the reference accelerometer biases. The Septentrio roll and pitch measurements have errors in them also, and these errors will corrupt any estimates of the accelerometer biases. Therefore an artificial bias of 0.6 m/s^2 was added to all three accelerometer signals in order that the bias effects would be relatively much larger than the Septentrio errors.

The purpose of the initialization test was to emulate the long term bias estimation performance of the AUNAV estimator. The sinusoidal steering was done in order to accelerate the convergence so as not to need an extraordinarily long data set. The idea was to provide as much dynamics within the data set and to show the biases converging to the true values. To this end the tuning was also modified. If the filter could be tuned more aggressively during the favorable dynamic conditions, the convergence could potentially be accelerated further. This was done by scaling the values of the Q matrix which correspond to the accelerometer bias states according to the yaw rate, as in Equation (4.3).

$$Q_{ii} = \frac{Q_{ii}\omega_z}{1^{deg}/_s} \tag{4.3}$$

The notation Q_{ii} refers to the Q matrix values which correspond to the accelerometer bias driving noise. This equation simply scales the Q value according to the current yaw rate signal, normalized to $1^{deg}/_s$. The adaptive scaling did in fact prove to accelerate the convergence time. Figure 4.3 shows the lateral accelerometer bias estimation converging to the reference value during the period of lateral excitation ($t \approx 200 - 400s$). Recall that the reference accelerometer bias values were obtained in the manner described previously in this section using the Septentrio. This result is in accord with the observability results discussed in Section 3.2. It shows that under dynamic maneuvers corresponding to the direction of the accelerometer of interest, the bias estimate for that accelerometer will be observable. It is also important to consider the affect this has on the attitude estimates. As was previously stated, it is expected that if the bias estimates are correct, then the attitude estimates will likewise be correct because the combination of the two are observable together. All of the states are on some level coupled in the filter, but the strongest coupling between the lateral accelerometer and attitude is with the roll angle. This is because errors in the roll angle estimate cause the exact same effect in the velocity and position estimates as lateral accelerometer bias.

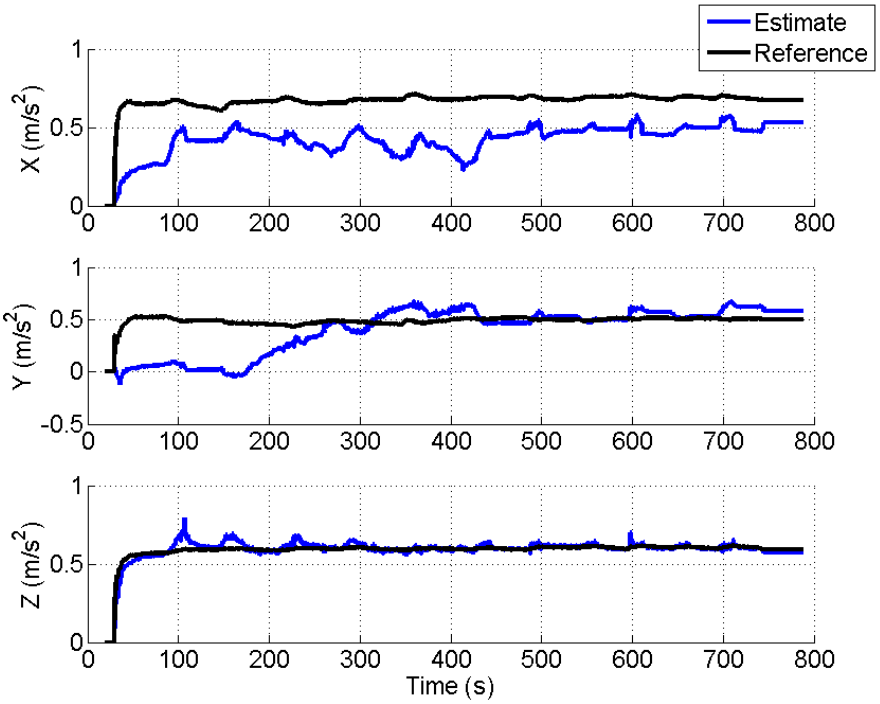

Figure 4.3: AUNAV Accelerometer Bias Estimation and Convergence During Initialization Experiment.

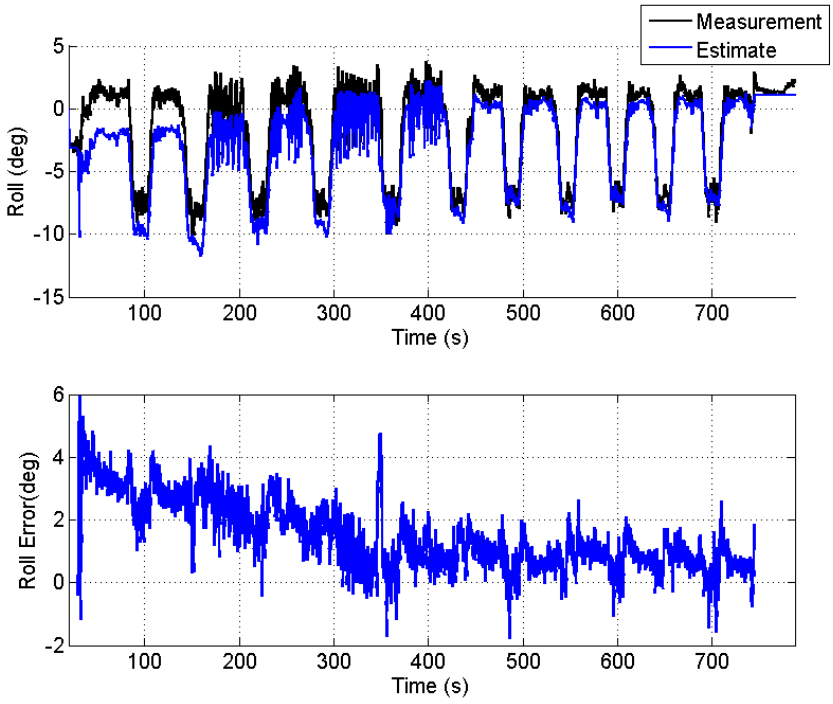

Figure 4.4: AUNAV Roll Angle Estimation and Convergence During Initialization Experiment.

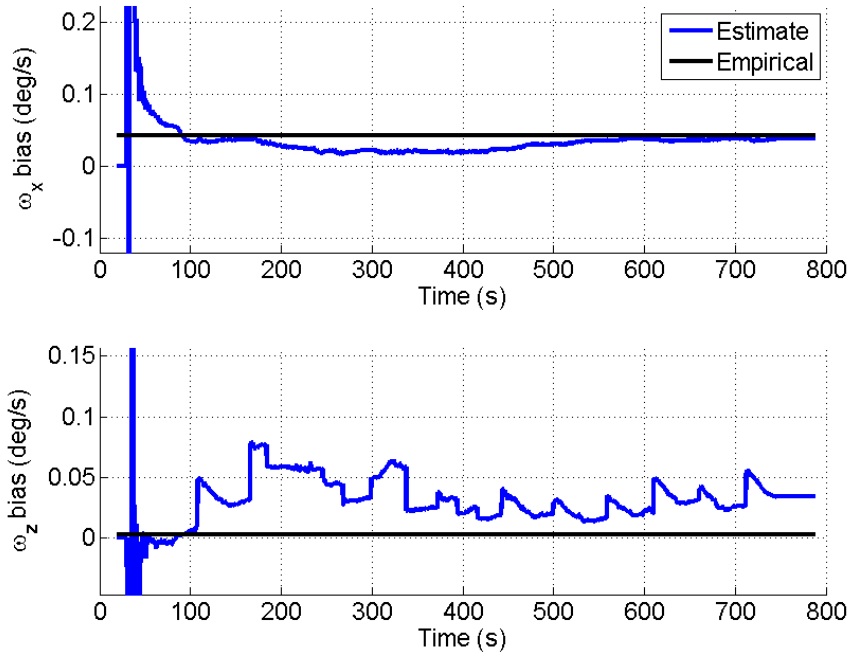

Figure 4.5: AUNAV Gyroscope Bias Estimation and Convergence During Initialization Experiment.

Figure 4.4 shows the roll angle estimate, which is initially very biased. This demonstrates the filter's inability to separate it's effects from those of the accelerometer bias during periods of low dynamics, rendering it an unobservable state. However, recall that the net effect of the two states remains observable, thus the estimates do not drift off. As soon as the vehicle experiences lateral excitation the roll angle begins to converge to the true value. It continues to do so until the excitation subsides. This is to be expected according to the theory, because the lateral accelerometer bias estimate is also converging towards its true value.

The gyroscope biases are also an important part of the initialization process. Figure 4.5 shows the estimates compared with the empirically obtained values. The roll rate gyroscope bias estimate converges to the empirical value. This confirms the results from Section 3.3. The yaw rate gyroscope bias estimate does not display the same convergence as the roll rate bias estimate. Nor does it exhibit the same convergence here as it does in simulation in Section 3.3. This is because in the real world experiment the bias estimate is also capturing

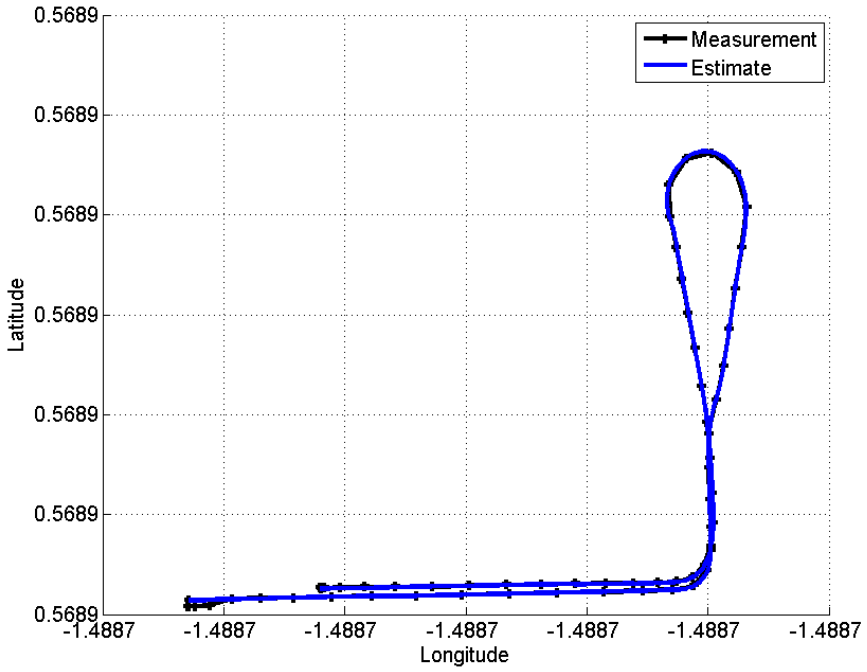

Figure 4.6: Vehicle Trajectory for Dynamic Test on the NCAT Skid Pad.

any unmodeled scale factor or misalignment errors. These are not present in simulation. Nevertheless, the error in the bias estimate remains very small, and the overall filter operation is not greatly affected by it. Future work will examine means of capturing scale factor and misalignment errors properly, so that they do not appear in the bias estimate.

4.2 Tracking

Figures 4.6 and 4.7 show the trajectory and some general dynamics of the experiment for the tracking phase. The third subplot in Figure 4.7 simply indicates when the yaw constraint condition is triggered; it is included because the constraint condition plays a large role in the overall operation. The results from this run are seen in Figures 4.8-4.11, starting with the velocity estimates. The velocity estimates follow the true values well, which is not surprising, given that each has a direct measurement for correction. These estimates are also smooth, as opposed to being jagged and having a sawtooth like appearance, which is a

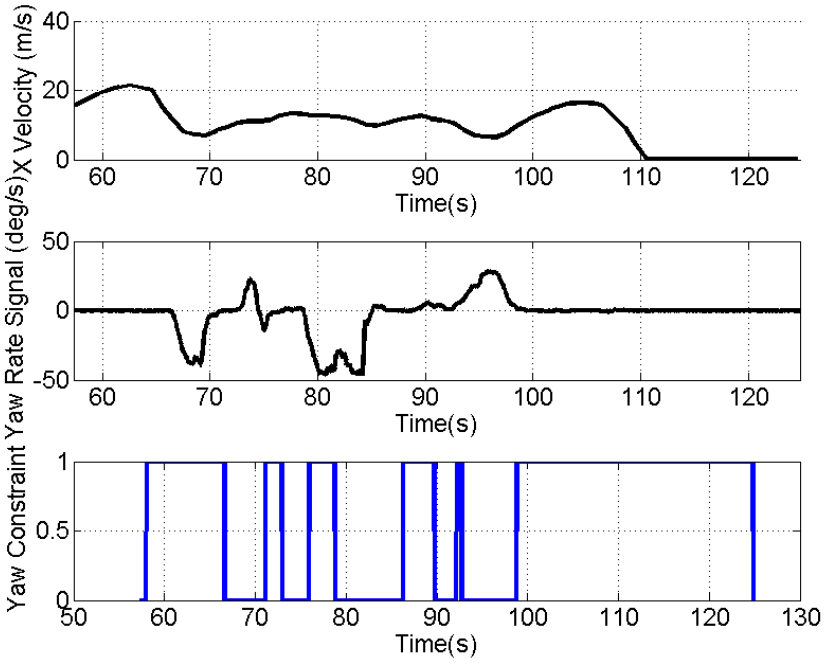

Figure 4.7: Profile of Dynamic Conditions During Dynamic Test on the NCAT Skid Pad.

good initial indication that the rest of the estimates are also tracking well. The presence of jagged features in the velocity estimates would indicate that either the attitude states, the accelerometer bias states, or both, were not being estimated properly. A histogram of the velocity residuals is shown in Figure 4.9. The velocity residuals are mostly white Gaussian, another sign indicating proper Kalman filter operation. Although the East velocity residuals do not appear perfectly Gaussian, they do appear close to Gaussian and they are of very small magnitudes. The red line represents the normal probability density function given by the mean and standard deviations of the residuals.

Figure 4.10 displays the sideslip estimate in blue with the Septentrio reference in black. It is observed from the plot that overall the AUNAV estimator accurately estimates the sideslip angle. Yet it is also observed that a large error in the estimate appears at $t \approx 84.5s$. This error occurs during a GPS update, and the filter subsequently corrects for it. It takes the AUNAV estimator 2 seconds to bring the error to within 1 degree and close to 5 seconds to bring the error near zero. The error occurs at a time when the vehicle is experiencing

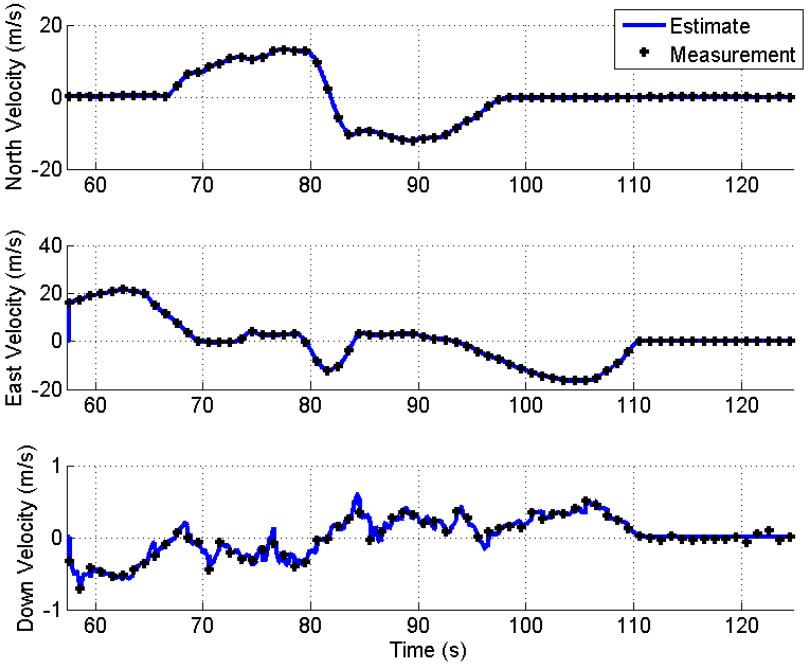

Figure 4.8: AUNAV Velocity Estimates for Dynamic Test.

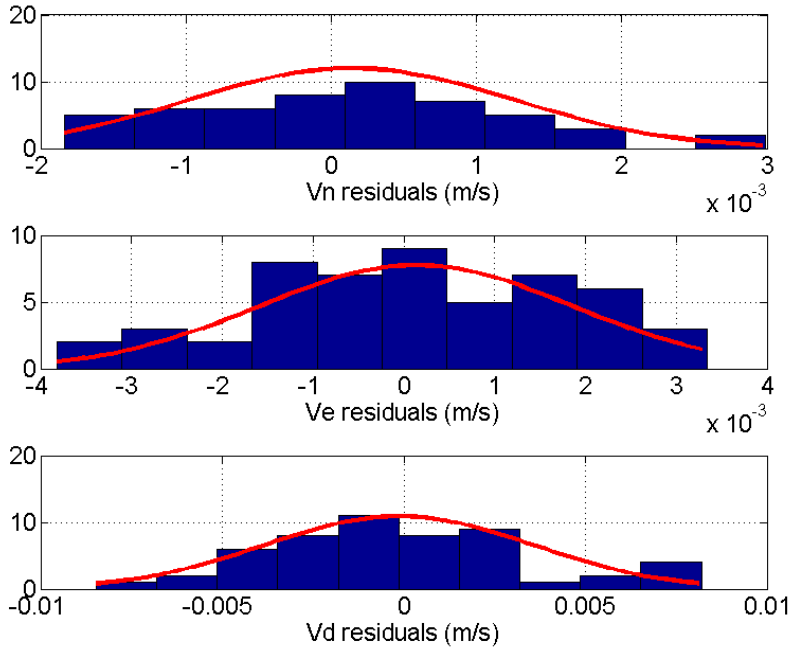

Figure 4.9: Velocity Estimation Residuals for Dynamic Test.

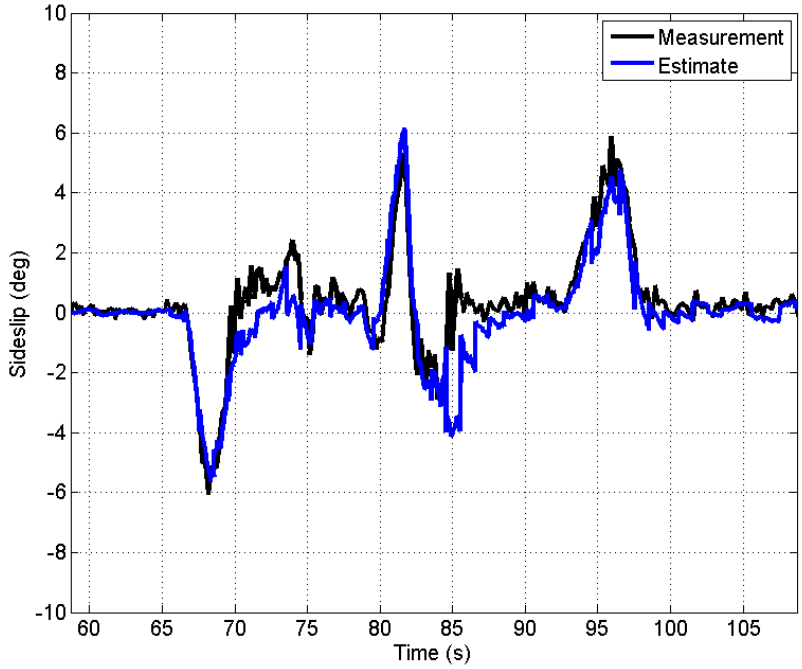

Figure 4.10: AUNAV Sideslip Estimation During Dynamic Test.

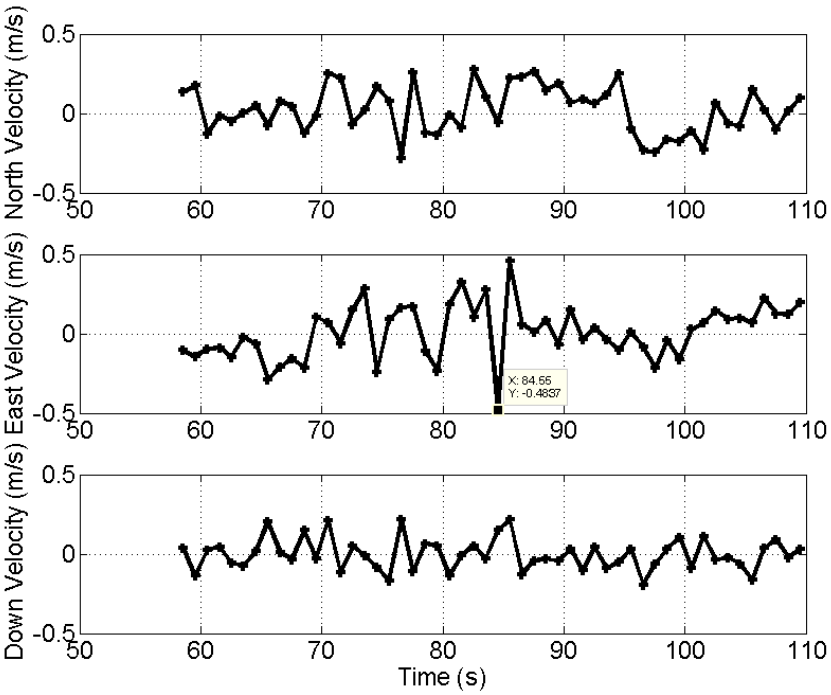

Figure 4.11: Velocity Innovations During Dynamic Test.

very heavy roll dynamics, suggesting that the lever arm could be a factor. The effects of the lever arm in terms of m/s follow the relationship given by Equation (4.4).

$$\delta \underline{V} = \underline{\omega} \times \delta \underline{l} \tag{4.4}$$

Here, $\delta \underline{V}$ represents the difference between the velocities at the Novatel antenna and the IMU location, $\underline{\omega}$ is the rotation rate vector, and $\delta \underline{l}$ is the vector from the IMU to the antenna in the body frame. The yaw dynamics will not cause any noticeable effect through the lever arm, since the Novatel is mounted almost directly over the IMU when considered in the XY plane. However the pitch and roll dynamics interact with the Z displacement, which is in this case approximately 1 meter. For this data run the roll rate at $t \approx 84.5s$ is -0.31rad/s. This means that in terms of general speed, not specifically North and East velocities, the error introduced is ≈0.3m/s. The standard deviation of the noise on either the North or East velocities is only 0.05m/s, meaning that the lever arm is having a relatively large effect here. A plot of the velocity innovations is given in Figure 4.11, showing that at $t \approx 84.5s$ the innovation on the East velocity is -0.49m/s, which is larger than the other innovations for this state. This indicates that the lever arm is having an adverse effect and is the likely culprit of the sideslip estimate jump at $t \approx 84.5s$. It should be noted that at $t \approx 84.5s$ the heading angle is 170 degrees, which means that this roll motion almost exclusively influences the East velocity measurement.

The roll angle estimate is shown in Figure 4.12. The estimate closely follows the reference value, with errors mostly staying within 1 degree in either direction. As was mentioned at the beginning of this section, there is some drift in the reference attitude solution which needs to be considered when comparing the AUNAV estimate with the Septentrio reference. It is possible that the estimate is either slightly closer to or slightly further from the true value. What is important here is to show comparable performance to the reference system. Future work will lengthen the baseline of the Septentrio, increasing the accuracy of the reference roll angle.

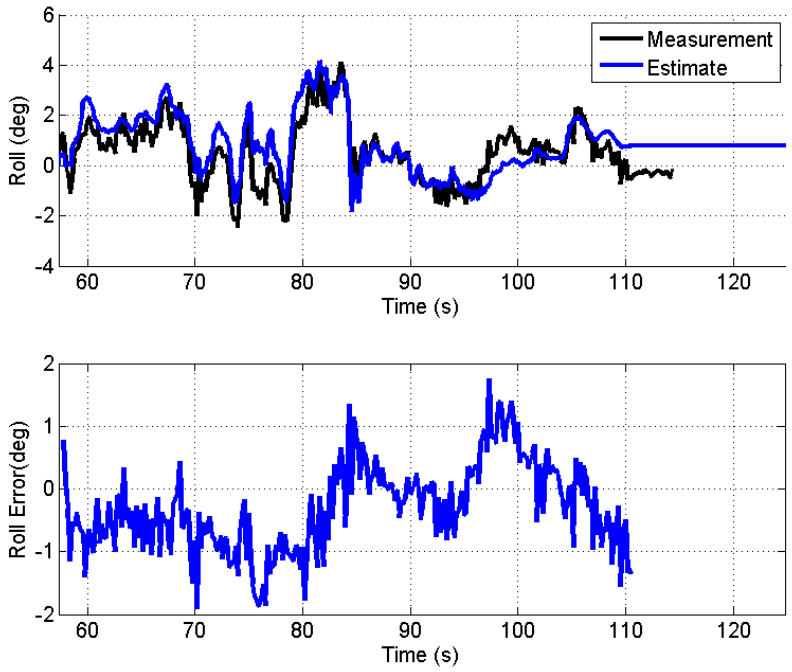

Figure 4.12: AUNAV Roll Estimate for Dynamic Test.

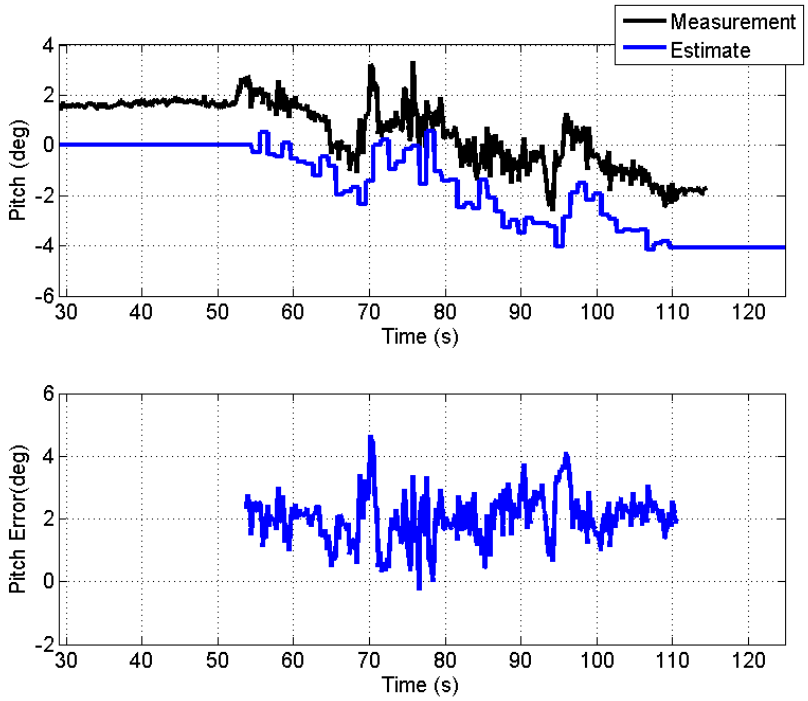

Figure 4.13: AUNAV Pitch Estimate for Dynamic Test.

The pitch estimate is unique among the attitude states, as there is no rate gyro to aid the estimation process. The estimate therefore is only updated at the period of the GPS, giving it a much lower bandwidth. This characteristic of the performance is clearly seen in Figure 4.13. Despite the low bandwidth the filter tracks the general dynamics acceptably well, albeit with a slight lag. The most obvious characteristic of the plot, however, is the bias in the estimate of around two degrees. It is hypothesized that this is due to a mounting angle difference between the Septentrio and the Crossbow, meaning that the estimate isn't biased at all. This hypothesis could be tested by using the Septentrio attitude solution as a measurement to the filter and observing the bias states. Adding these measurements makes the filter totally observable at all times. The pitch measurement would cause the filter to interpret the relative pitch difference as a bias in the forward accelerometer by essentially forcing the pitch estimate to the the Septentrio measurement. The resulting forward accelerometer bias could then be interpreted as the relative pitch and solved for via the simple relationship given in (4.5).

$$\delta f^x_{grav} \approx g \sin(\theta) \tag{4.5}$$

This does assume that the relative pitch angle effects are much greater than the inherent bias, which is a good assumption for the Crossbow. The Septentrio was used to identify the bias resulting from the relative pitch mounting angle, and the result was a bias of approximately $-0.375 m/s^2$. This was done using the nearly 50 seconds of static data from the same data run. A plot of the pitch angle and the bias estimate is shown in Figure 4.14. Computing a relative angle from this via Equation (4.5) yields an angle of -2 degrees, which matches the "bias" in the pitch. Figure 4.15 shows two subplots, the first of which shows the pitch estimate, the Septentrio measurement, and the road grade estimate. Recall that this estimate, obtained as described in Chapter 2, is unbiased. Subplot 1 of Figure 4.15 shows that the Septentrio is very close the road grade estimate. The second subplot shows the

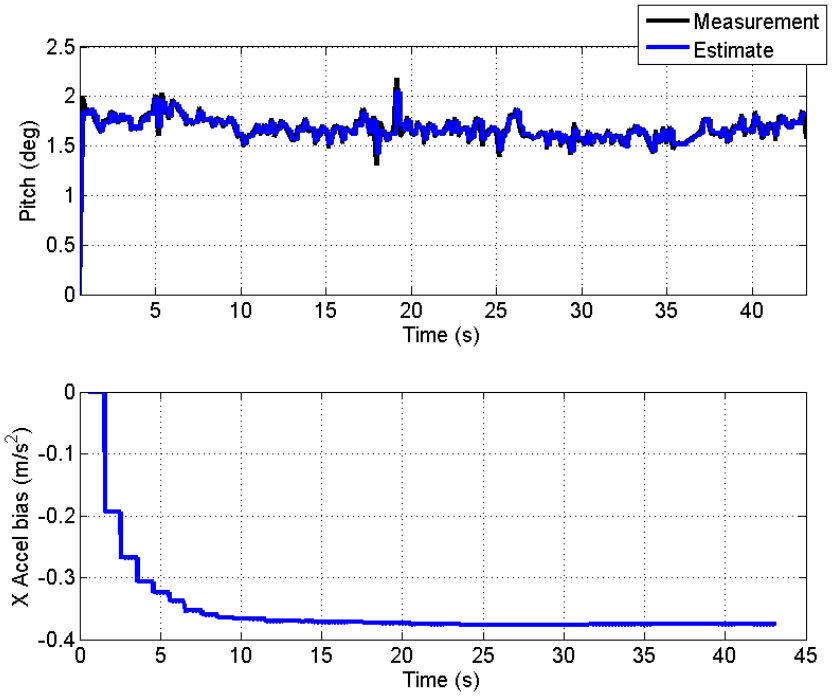

Figure 4.14: Pitch Estimate when Using Septentrio Pitch Information as an Extra Measurement.

differences between the Septentrio measurement and each estimate (pitch and road grade). The road grade difference is practically 0 degrees in steady state, while the pitch estimate is very near 2 degrees. A line indicating the calculated relative pitch, obtained using Equation (4.5), is shown in green. It is observed that the difference between the pitch estimate and the Septentrio measurement matches the estimated relative pitch. It is concluded that the Septentrio mounted pitch angle is level with the road. This does not mean that the suspension pitch is necessarily zero, only the sum of the suspension pitch and the Septentrio mounting angle is zero. Furthermore it is concluded that the Septentrio and Crossbow have a 2 degree mounting angle difference, and that this mounting angle difference constitutes the majority of the "error" between the AUNAV pitch angle estimate and the Septentrio measurement.

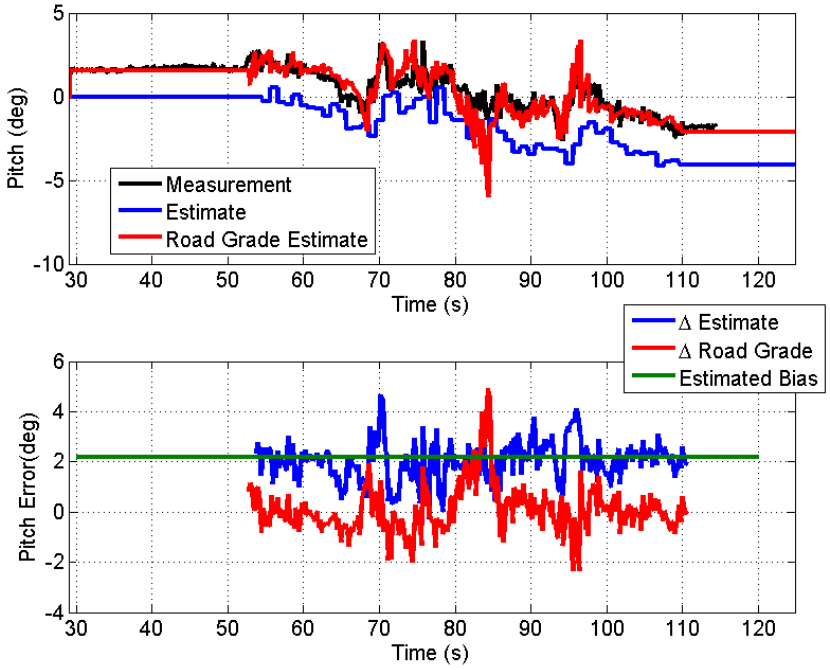

Figure 4.15: Comparison of the AUNAV Pitch Estimate with the Septentrio Measurement and the Road Grade Estimate.

4.3 Slowly Growing Sideslip

One of the major limitations for ESC systems is detecting low rates of sideslip growth. That is, it is difficult to estimate the sideslip angle when it builds up slowly over time. The rate of sideslip growth presents a threshold of operation below which the ESC system cannot detect. The cause of this can be understood by considering the noise of the sensors. Many estimation schemes rely heavily on the lateral accelerometer for information. Recall the definition of sideslip given by Equation (1.2), and that it is predominately related to the lateral velocity. Therefore from this definition the sideslip growth rate will be predominately related to the lateral velocity rate, i.e. lateral acceleration. Recall also that the lateral accelerometer is subject not only to its own noise and moving bias, but also from biased effects from incorrect roll angle compensation. These three error sources combine to effectively raise the threshold below which true lateral accelerations cannot be detected. Said another way, the error sources decrease the signal to noise ratio.

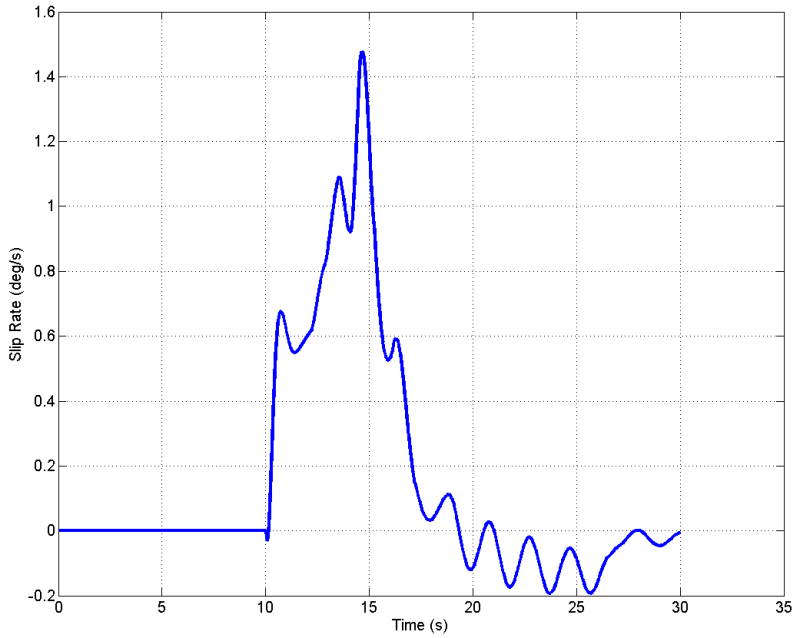

Figure 4.16: Rate of Sideslip Growth During Simulated Test.

The AUNAV estimator avoids this problem by approaching sideslip using the definition given by Equation (1.1), which is a function of the North and East velocity states and the heading state. The velocity estimates are unbiased, and only contribute small amounts of noise to the overall error of Equation (1.1). The main error source is the yaw angle, yet this is precisely where the advantages of the AUNAV estimator and GPS are leveraged. With most systems, the yaw rate bias would present a major problem here and cause rapid drift in the yaw angle. This drift would be indistinguishable from slowly growing sideslip. However recall that yaw information from GPS is applied when the vehicle is driving straight, resulting in observability for the yaw gyro bias and allowing the system to effectively remove the bias. Therefore practically the only source of drift in the yaw angle is noise, which causes a much lower rate of drift than bias does, making it possible to estimate lower rates of sideslip growth.

It is desirable to demonstrate the effectiveness of the AUNAV estimator to estimate sideslip during slow rates of growth. Simulations were performed in Carsim in order to

initially validate the estimation performance during slow sideslip buildup. Figure 4.16 shows the slip rate for the simulation of the slow sideslip buildup. In this simulation, a 4WD vehicle was briefly driven straight followed by a linearly increasing steering input. The slip rate peaks at only $1.5^{deg}/_{s}$. The estimation performance is shown in Figure 4.17. The estimator is in fact able to estimate the slow sideslip buildup for a relatively long period of time in terms of ESC events. Notice that the estimate does begin to drift as soon as the slip rate slows down. The peak slip rate is extremely low in the first place, and once it decreases further the filter cannot estimate it. Also notice that the error drift is still very slow. This is because the drift is entirely due to drift in the yaw estimate, which is itself only due to drift from noise integration. In this run the yaw constraint was true until $t = 11.1s$. The last GPS update that occurred while the constraint was true occurred at $t = 11.0s$, which is *after* the turning maneuver begins. In this case, the assumption of zero sideslip is technically false, but applying the course measurement doesn't introduce any noteworthy error. However this does not guarantee that improperly applying the course measurement in other cases would be equally harmless.

A closer examination of the sideslip shown in Figure 4.10 (the experimental data) reveals that the slip rate during the third turn is relatively small. The average slip rate for the third turn is $1.8^{deg}/_{s}$. For reference, the average slip rates for turns 1 and 2 are $-4^{deg}/_{s}$ and $4.5^{deg}/_{s}$, respectively. This is calculated by taking the slope of the black line (measured slip value) at the points highlighted in Figure 4.18. There is some variation in this slope near the top of the peak, but some of this is noise. Since the sideslip growth rate on this turn is very close to that of the simulation, this portion of the run is further examined to validate the estimation performance. Figure 4.18 shows that the estimator is able to estimate the sideslip during this time, demonstrating the sideslip estimation performance during slow sideslip buildup. Also note that the estimate is accurately maintained throughout the maneuver. This result is in agreement with the simulation, because in simulation the estimate did not exhibit much drift until after short periods of time. In Figure 4.18 there does not appear to be consistent

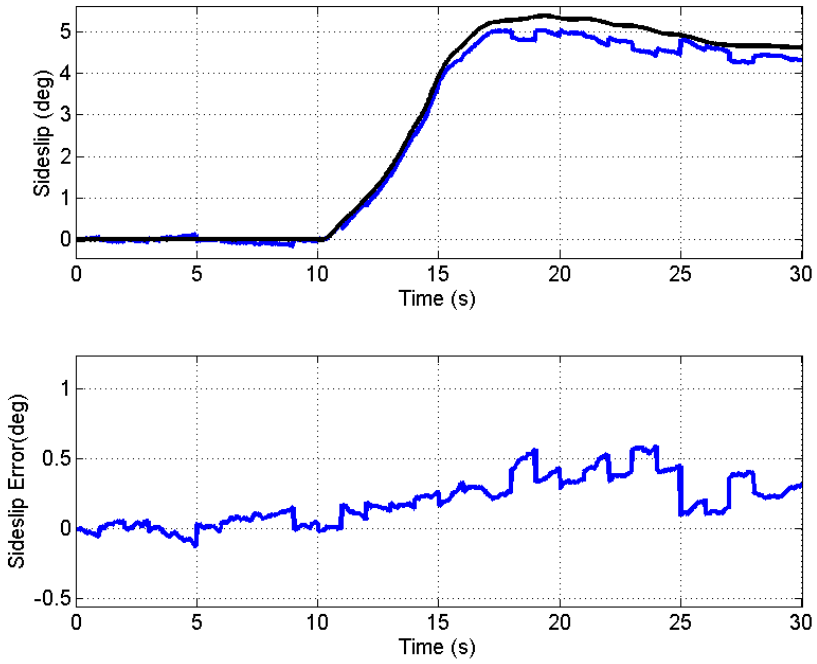

Figure 4.17: AUNAV Sideslip Estimate During Simulated Test.

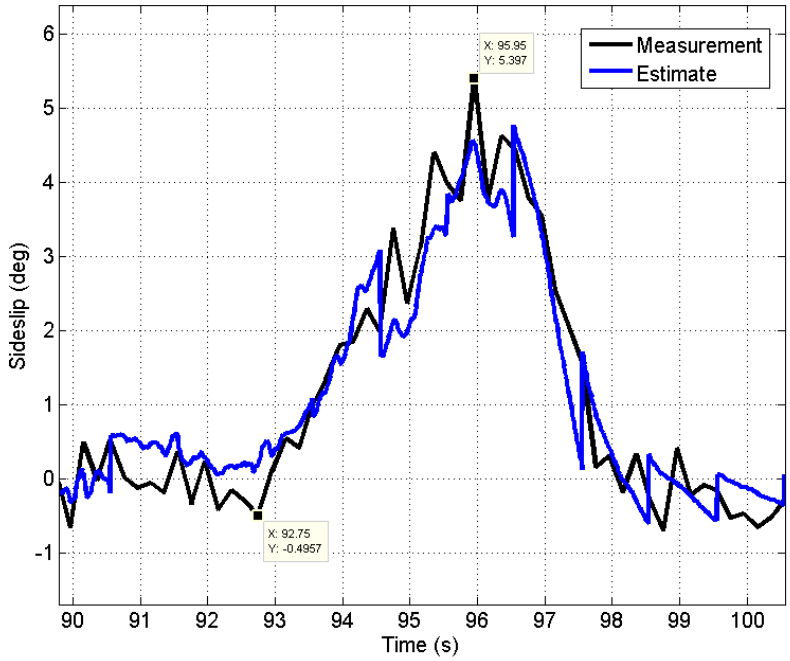

Figure 4.18: AUNAV Sideslip Estimate During Experimental Test with Slowly Growing Sideslip.

drift in the estimate, although certainly there is some noise. In conclusion, the AUNAV estimator is able to accurately estimate sideslip growth rates as low as $\approx 1.5^{deg}/_{s}$.

4.4 Conclusion

Experimental tests were performed to validate the estimation performance for the AUNAV estimator. The tests were divided into two phases. The initialization experiment was performed on the NCAT loop. Laps were driven in which sinusoidal steering inputs were provided in order to demonstrate the convergence of the bias and attitude states in the presence of dynamics, even despite large initialization errors on the bias states. A reference filter which incorporates the Septentrio attitude information was used for comparison of the bias states. It was shown that given enough dynamics over time the filter will accurately identify the biases. It was also shown that as the bias estimates converge, the corresponding attitude estimates converge also. In effect, the AUNAV estimator accurately separates the sensor biases from the attitudes. This result agrees with the conclusions from Chapter 3. The dynamic estimation phase, where it was assumed that the biases had already been sufficiently identified, was discussed second. Dynamic maneuvers producing sizable sideslip and roll were performed in the tracking experiment, and it was shown that the AUNAV estimator accurately estimates the vehicle states. An instance of lever arm effects negatively affecting the estimation performance was discussed. Finally it was shown that the AUNAV estimator can estimate slowly growing sideslip. In conclusion, the AUNAV estimator can provide the state estimation performance necessary for ESC and RSC systems.

Chapter 5

Experimental Comparison with the MME Estimator

The modular filter was also validated with experimental data, using the data from the same two experiments already discussed in Chapter 4. This assures that the performance of the two filters can be properly compared. It is worth noting that the two previous experiments were designed specifically to address the nature of operation of the AUNAV estimator. In particular, the first experiment was designed to validate the initialization phase. The second experiment assumed that the "initialization" was successfully completed and so the focus was on estimating states during dynamic events pertinent to ESC systems. It may be asked if it is appropriate to use these same experiments to validate the modular filter, as it may behave differently. In actuality, the modular filter suffers many of the same weaknesses in terms of initial errors and bias estimation. Recalling the issue from the AUNAV estimator, the need for the initialization time arises from the fact that if the biases are not properly determined they will exert a negative impact on the estimation quality of the other states. In fact this effect is even more pronounced for the modular case. It should be no surprise to find this similarity in the behavior of the two filters, as they are based on the same fundamental relationships. The only real differences between the two are the coordinate frames into which the states are resolved, the inclusion of position states in the AUNAV estimator, and in the relative level of coupling between all of the states. That is, the AUNAV estimator includes more of the coupling relationships, whereas by definition the modular filter does not.

5.1 Initialization

In this section consideration is given to the bias estimation. For the AUNAV estimator the discussion is grounded in the larger idea of observability, and this is also true here. However, as the modular approach entails many separate filters as opposed to one large filter, we cannot speak of the overall observability of the whole system. It can be said from the start that the forward and vertical velocity filters are always fully observable, and that when driving straight the heading filter is observable.

In terms of the biases, this means that there will always be reliable estimates of the vertical accelerometer bias and that the yaw gyro bias will be reliable so long as the vehicle does not undergo extremely long periods of uninterrupted turning. There is a caveat to the observability of the forward accelerometer bias. Strictly speaking, this filter is always observable. However, the filter does not include any pitch information inherently, rather the accelerometer is rotated into the appropriate frame and gravity is removed in processes outside of the filter. So any observability analysis of the forward velocity and forward accelerometer bias is inherently decoupled from any notion of whether or not the pitch is observable. In short, the net effect of the forward accelerometer bias and any pitch angle estimation error is always observable, and this combined effect is what the forward velocity filter estimates. Therefore if the pitch estimate is accurate, the filter's estimate of the bias will accurately represent the bias inherent in the sensor alone.

The lateral state filter is different from the other three in that at no time is a direct measurement of the lateral velocity or roll available. The lateral velocity "measurement" for this filter comes from the combination of the GPS velocity and the sideslip estimate from the heading filter, as given in Equation (2.52). From this standpoint the accuracy of the sideslip estimate from the lateral filter will never be more accurate than that of the heading filter (although it is available at the INS frequency as opposed to the GPS frequency). If the assumption is made that the sideslip estimate from the heading filter is

accurate, meaning that the measurement for the lateral filter is accurate, then some valuable conclusions regarding the observability of the other two states in the filter can be made.

The pair of the system matrix and the measurement matrix, each shown in Equations (2.50) and (2.51), is linear, time invariant, and always observable for the lateral state filter. This means that the roll rate gyroscope bias is always observable, and that the lumped effect of the lateral accelerometer bias and the roll angle is always observable. This result is in accord with intuition. Recall that for the AUNAV estimator the combination of the leveling angles and the biases is observable when the INS is experiencing constant acceleration. It is reasonable to expect a similar result here. Likewise, recall that for the AUNAV estimator the leveling angles and the biases could be separated (independently observed) under enough dynamics. The same is true here, only in this case additional low pass and high pass filtering is applied to the output of the lateral filter externally. Therefore the expectation is that if the heading filter estimate of sideslip is accurate, the lateral filter will produce accurate estimates of the sideslip, the roll rate gyroscope bias, and the net accelerometer bias. If the INS experiences dynamics, the net accelerometer bias can be separated into roll and true bias.

A simulation was performed in Carsim to test this expectation. The vehicle was simulated driving straight at 19.5 $^{m}/_{s}$ on a flat road, followed by a prolonged period of sinusoidal steering inputs, followed again by straight driving. This simulation is much like the actual test performed to validate the initialization of the AUNAV estimator. Figure 5.1 shows the yaw rate. Figure 5.2 shows the estimate of the lumped state ($\phi + b_y$), along with the true roll and the true bias in the bottom plot. It can be seen that the lumped state has an approximately constant offset and a dynamic component. The offset corresponds to the effects of the sensor bias, which is relatively constant over the run. The variation of the bias shown in the figure is less than $0.02 m/s^2$. The dynamic component is due to the changing roll angle. The filtering of the lumped state is shown in Figures 5.3 and 5.4. First, in 5.3, the roll angle is correctly filtered out during the transient maneuvers. Then in 5.4 the lateral

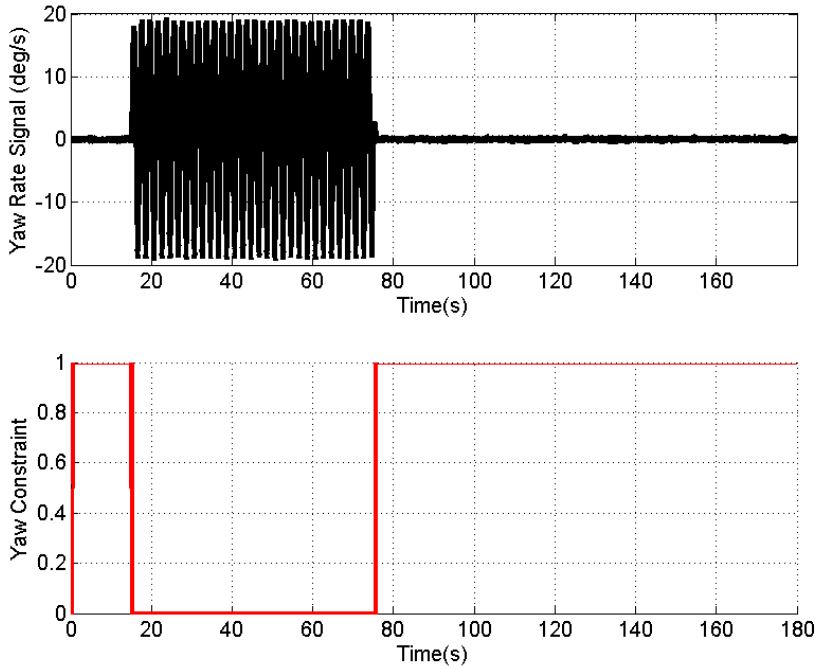

Figure 5.1: Yaw Rate and Yaw Constraint Signals of Initialization Simulated Test.

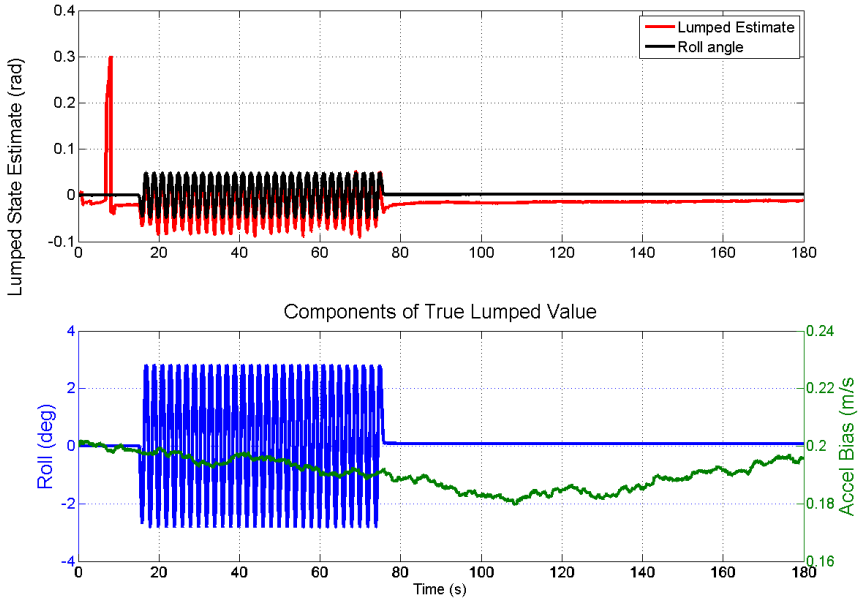

Figure 5.2: MME Estimate of the Lumped State Compared with True Simulation Values.

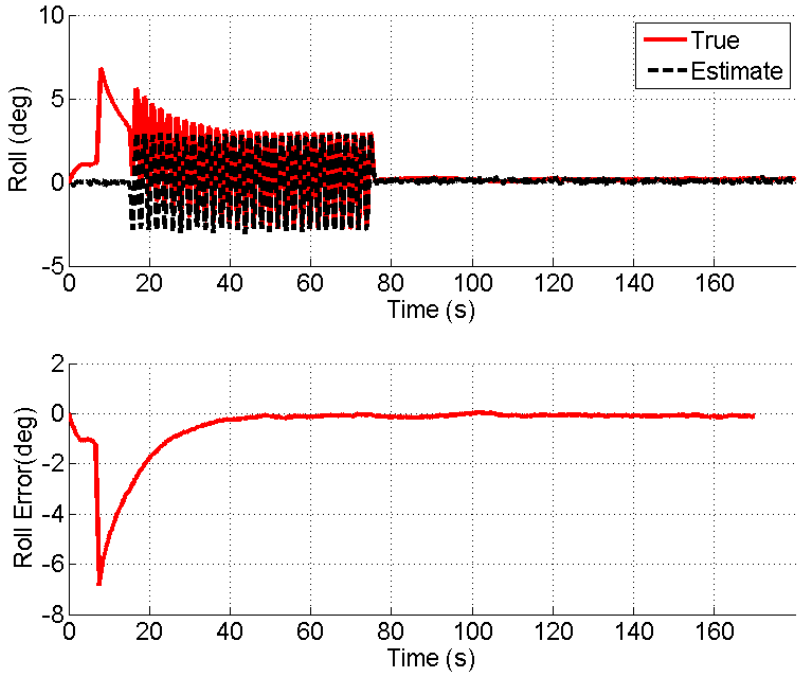

Figure 5.3: Convergence of the MME Roll Angle Estimate During Simulated Initialization Test.

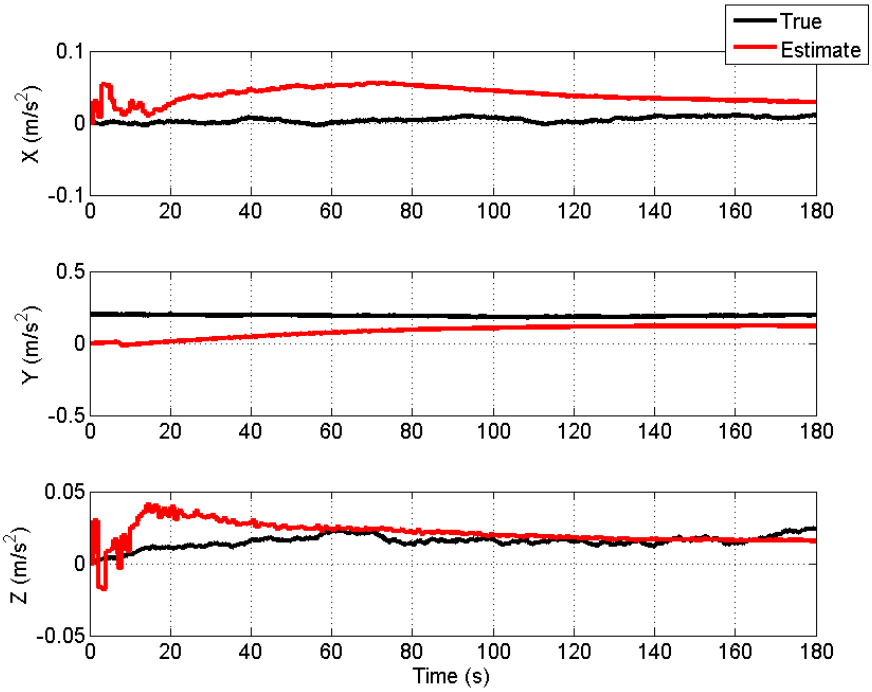

Figure 5.4: Convergence of the MME Accelerometer Bias Estimates During Simulated Initialization Test.

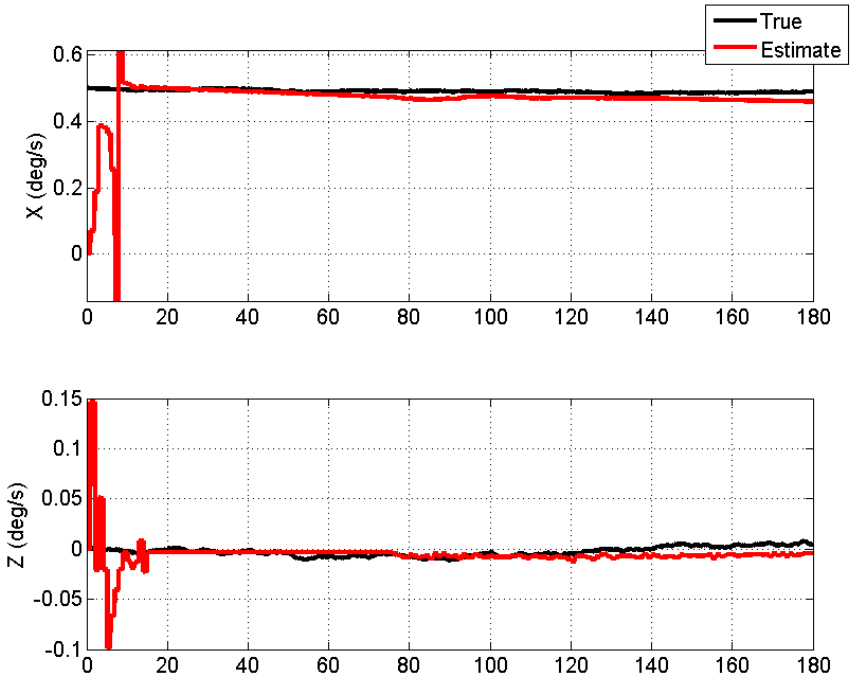

Figure 5.5: Convergence of the MME Gyroscope Bias Estimates During Simulated Initialization Test.

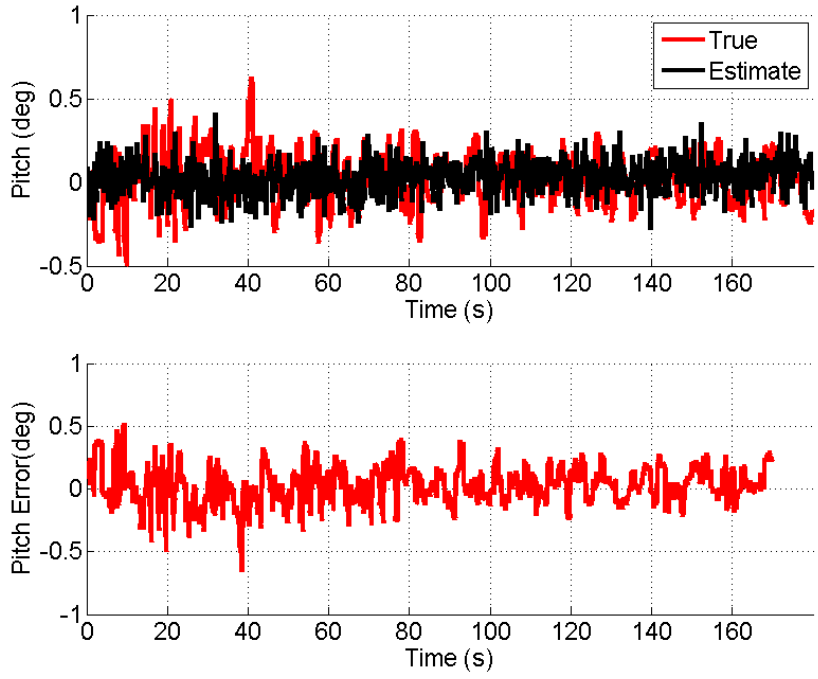

Figure 5.6: Convergence of the Pitch Angle Estimate During Simulated Initialization Test.

accelerometer bias is seen converging to the true bias. This process takes much longer, as the filter expects the bias to change very slowly. Figures 5.5 and 5.6 show that the gyroscope biases and the pitch angle are tracked correctly.

Experimental validation was done by applying the filter to the same validation data used for the AUNAV estimator. Figures 5.7 and 5.8 show the yaw rate and velocity estimates for this run, while Figures (5.9 - 5.11) show the estimation results of the leveling angles and the bias states. First, the forward and vertical velocity estimates show good performance in Figure 5.8. It is expected that this would be the case, because both filters are observable all the time. In order to compare these results with the AUNAV estimator, the AUNAV north and east velocity estimates are converted into forward velocity by simply taking the vector magnitude of the north and east velocities. There is no noticeable difference between the AUNAV estimator and the MME for the velocity estimates. The roll estimate is shown in Figure 5.9. Recall that the roll estimate is obtained by high pass filtering the lumped state of the lateral filter, as described by Equation (2.55). The estimate requires time to converge under dynamics because of this external filtering which is required to separate it out from the lumped state. The MME estimate of roll behaves very much like the AUNAV roll estimate in this case. However, it is observed that the AUNAV estimate is more accurate than the MME estimate once the high dynamics have settled out.

The lateral accelerometer bias also requires time to converge. It is shown in Figure 5.10, where the reference values for the accelerometer biases are the same those described in Section 4.1. Here the MME lateral accelerometer bias estimate is shown converging to the reference bias estimate over time. These two figures demonstrate that the additional filtering of the lumped state from the lateral state filter can be successfully used to isolate the roll and lateral accelerometer bias under dynamic conditions. The MME lateral bias estimate converges more slowly than the AUNAV estimate. This is because the filtering for the MME estimate is strictly a low pass filter, while the AUNAV estimator takes advantage of all of the relationships built into the system dynamics matrix. The vertical accelerometer

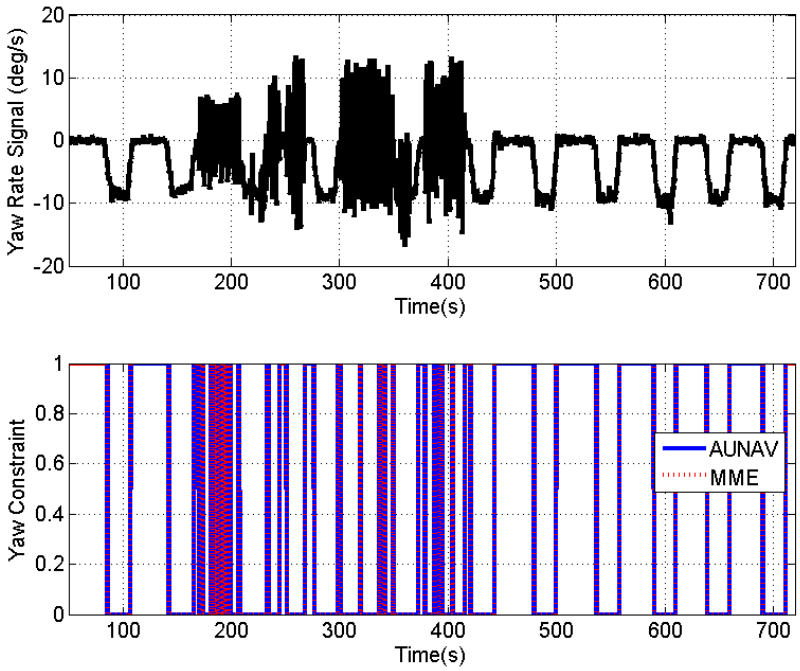

Figure 5.7: Yaw Rate and Yaw Constraint Signals During Initialization Experiment.

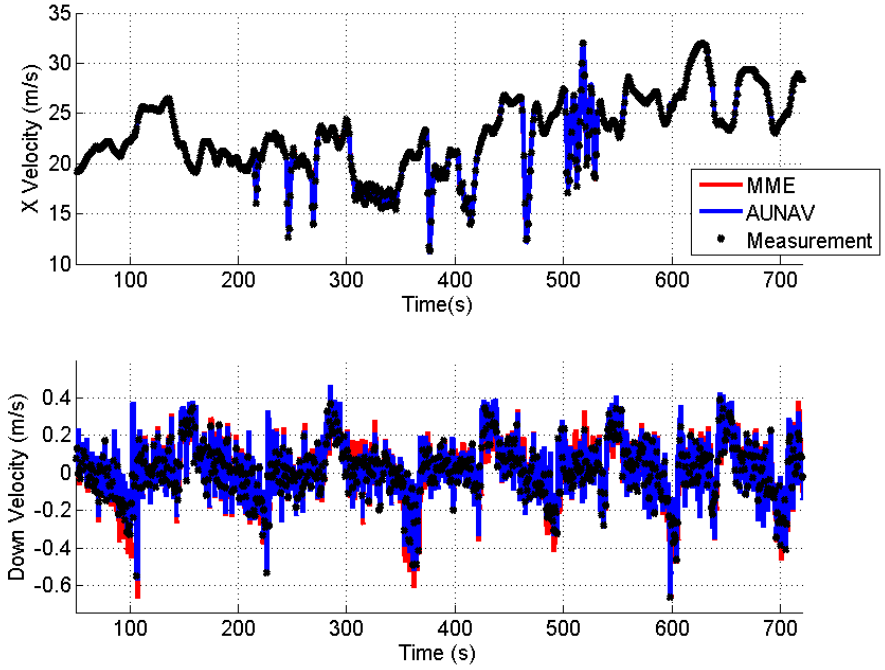

Figure 5.8: Comparison of Velocity Estimates During Initialization Experiment.

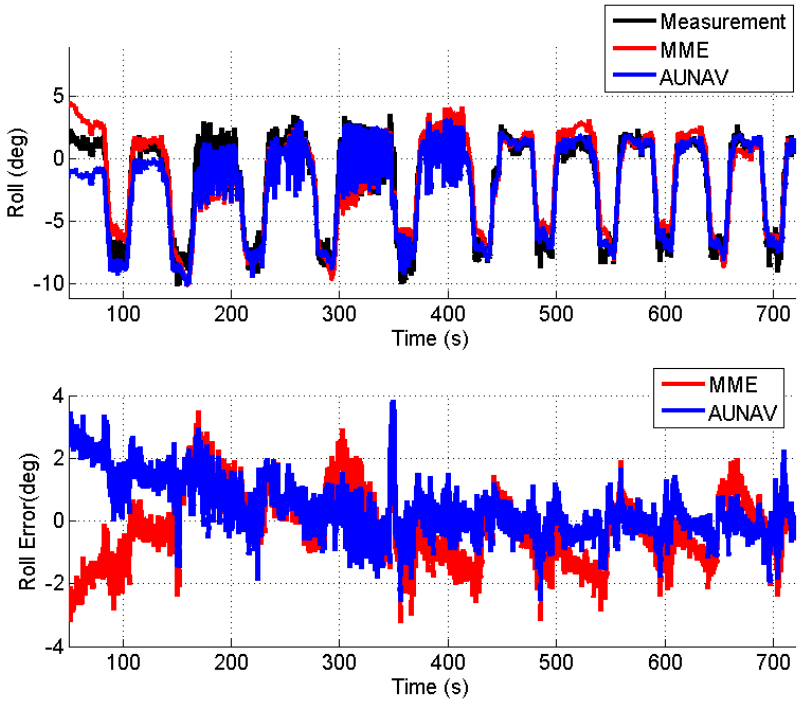

Figure 5.9: Comparison of Roll Angle Estimation Convergence During Initialization Experiment.

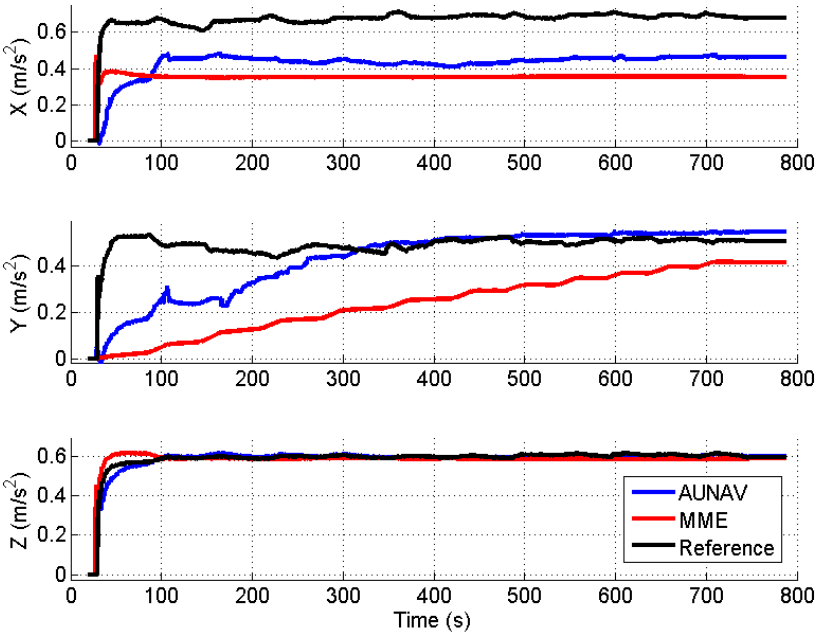

Figure 5.10: Comparison of Accelerometer Bias Estimation Convergence During Initialization Experiment.

bias in Figure 5.10 shows excellent performance, behaving just like the estimate from the AUNAV estimator. The forward accelerometer bias estimate has a clear offset from the reference. Recall that it was shown in Section 4.2 that there is a 2 degree pitch angle mounting difference between Septentrio and Crossbow. When this 2 degrees of pitch is converted into a gravitational bias according to Equation (4.5) it equates to $0.34 {}^{m}/_{s^2}$, which is almost exactly the difference between the reference forward accelerometer bias and the MME bias estimate. Therefore it is concluded that the relative pitch angle accounts for practically all of the difference between the reference forward accelerometer bias and the MME estimated bias. The difference between the AUNAV estimate of forward accelerometer bias and the MME estimate is that the MME estimator uses direct road grade information, from the road grade estimate, and assumes that there is no suspension pitch. The AUNAV estimator estimates the true pitch angle and does not make this assumption. The gravitational effects from the suspension pitch account for the difference between the AUNAV and MME forward accelerometer bias estimates. The pitch angle estimates are shown in Figure 5.11. The MME estimate is closer to the Septentrio measurement than the AUNAV estimate, having a mean difference of almost zero degrees. This is expected, given the findings shown in Section 4.2, where it was shown that the Septentrio is almost perfectly aligned with the road grade. The slight difference between the Septentrio measurement and the MME estimate is due to the variability of the Septentrio pitch measurement. The Septentrio pitch angle solution has a stated error standard deviation of 0.6 degrees. This does not mean that the error is zero mean. Rather, the mean of the error has a standard deviation of 0.6 degrees. This means that 95% of the time the pitch angle will be within plus or minus 1.2 degrees (2σ) of the true value. This variability not only affects the analysis of the pitch estimates but also the analysis of the forward accelerometer bias, because the Septentrio pitch information is used to obtain the reference bias.

The gyroscope bias estimates are shown in Figs. 5.12. The yaw rate gyroscope bias estimate displays similar characteristics to the yaw rate bias estimate from the AUNAV

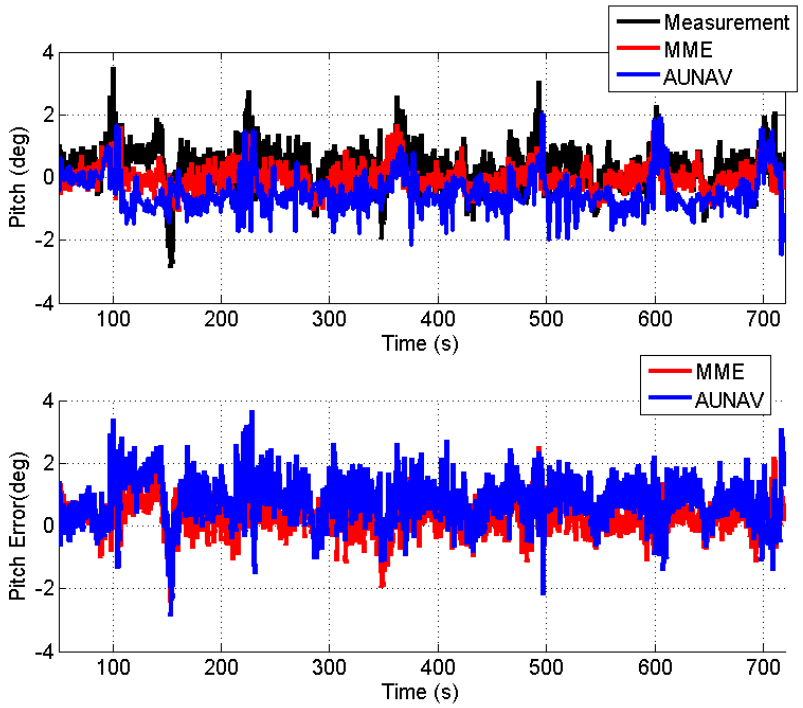

Figure 5.11: Comparison of Pitch Angle Estimation Convergence During Initialization Experiment.

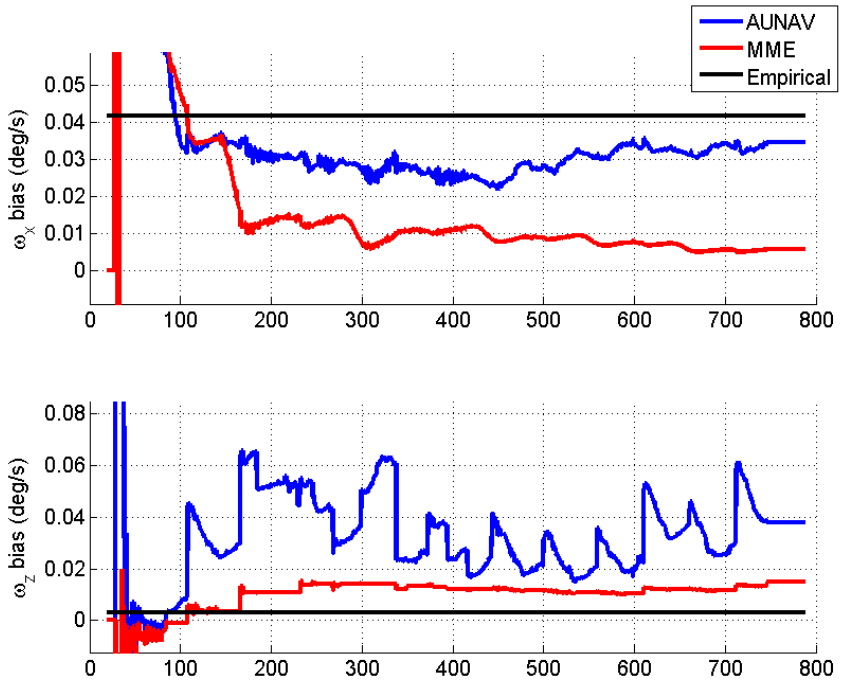

Figure 5.12: Comparison of Gyroscope Bias Estimation Convergence During Initialization Experiment.

filter. The estimate is too large because the MME estimator does not include any models of scale factor. Therefore these effects get lumped into the yaw rate bias estimate. These errors are not as large as the errors in the AUNAV estimate. The roll rate bias estimate does not converge to the empirically obtained reference, which is in contrast with the roll rate bias estimate produced from the AUNAV estimator. Since the estimate from the AUNAV estimator converges correctly, it is concluded that the error in the MME roll rate bias estimate is *not* due to unmodeled scale factor or misalignment, because these effects would also corrupt the AUNAV estimator. Rather, these errors arise from the fact that the lateral state filter does not have a direct measurement of the lateral velocity or roll. Recall that the lateral state filter uses an initial estimate of the lateral velocity, which is computed from the heading estimate and the GPS velocity as described in Section 2.2.4, as a measurement. This lack of a true measurement for any of the states of the lateral filter causes the error in the roll rate bias estimate.

5.2 Tracking

Figure 5.13 shows the trajectory of the dynamic test, while Figure 5.14 shows the yaw rate and the yaw constraint condition. Again, this is the same test from Section 4.2, consisting of driving in a straight line, turning uphill onto the skid pad, and finally performing a dynamic maneuver before returning to the starting location. The yaw constraint determination is the same for both algorithms. Figures 5.15 - 5.18 shows the results from this test, starting with the velocity estimates in Figure 5.15. The velocity estimates show good performance throughout the run, as would be expected due to the direct one second measurement updates. The velocity curves also offer insight into the quality of the other state estimates. Consistent drift in between updates, producing a sawtooth like appearance, would be sure evidence that either the bias states or the attitude states were inaccurate. It would likely indicate error in both, for reasons discussed in Sections 3.3 and 4.2. Some small drift in between updates is noticeable in the both the MME forward and down velocity estimates,

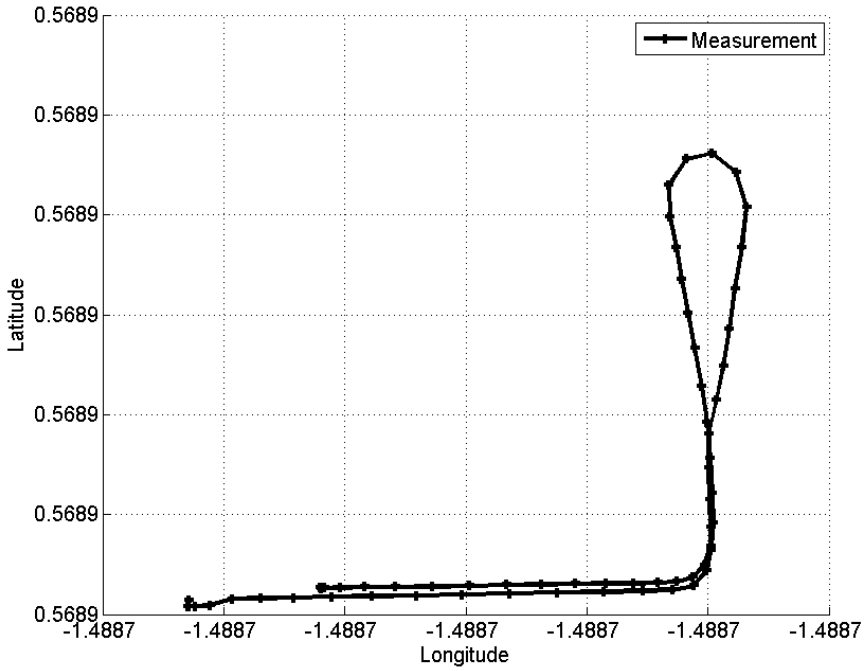

Figure 5.13: Vehicle Trajectory During Dynamic Test on NCAT Skid Pad.

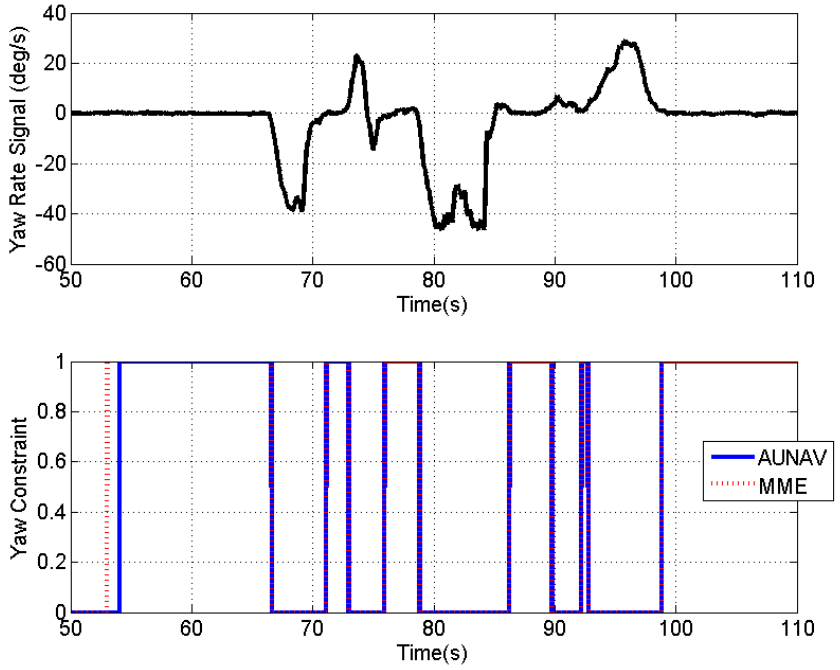

Figure 5.14: Profile of Dynamic Conditions During Dynamic Test on NCAT Skid Pad.

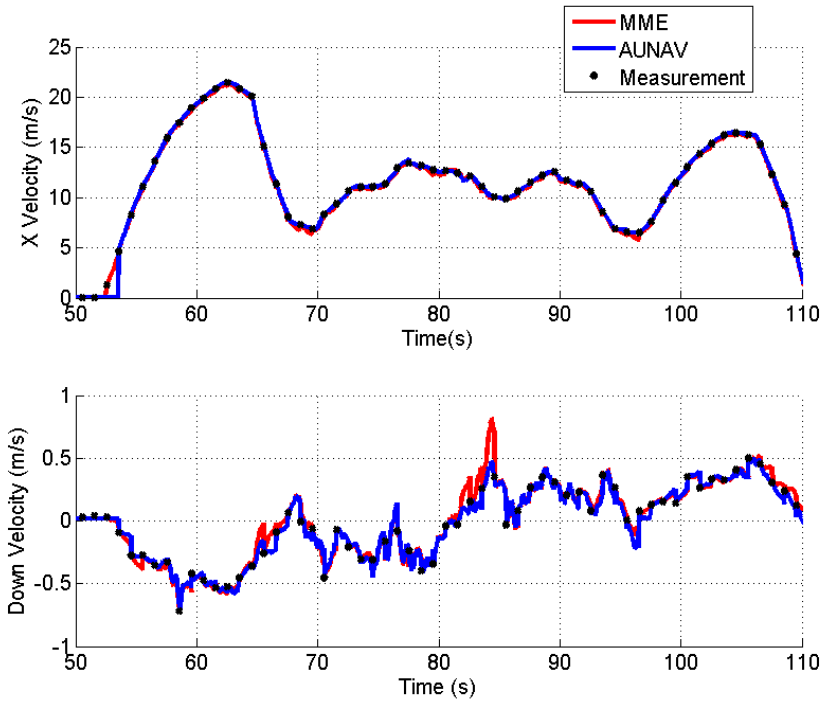

Figure 5.15: Comparison of Forward and Vertical Velocity Estimates During Dynamic Test.

but these are of a very small magnitude. Even though these errors are very small, they are noticeably larger than any errors in the AUNAV estimates. The errors in the down velocity occur during the high dynamic portions of the experiment, and are due to a combination of the unmodeled lever arm effects and the roll error which also occurs at this time.

Figure 5.16 shows the estimate of the roll angle along with the difference between it and the Septentrio reference. The MME roll angle estimate closely follows the reference throughout most of the run. The main exception to this occurs at $t \approx 84s$, where the error drifts off to close to 4 degrees and slowly converges back. It is highly likely that this error can be explained by precisely the same effect that causes the error in the sideslip estimates of both the AUNAV and MME estimators at the same time in the run. It was shown in Section 4.2 that unmodeled lever arm effects come into play under high dynamics via the difference between the velocity at the antenna and the velocity at the IMU. In the case of the modular filter, this difference is going to show up in the course measurement which is a component of the initial sideslip estimate, as in (2.39). The initial sideslip estimate serves to calculate the

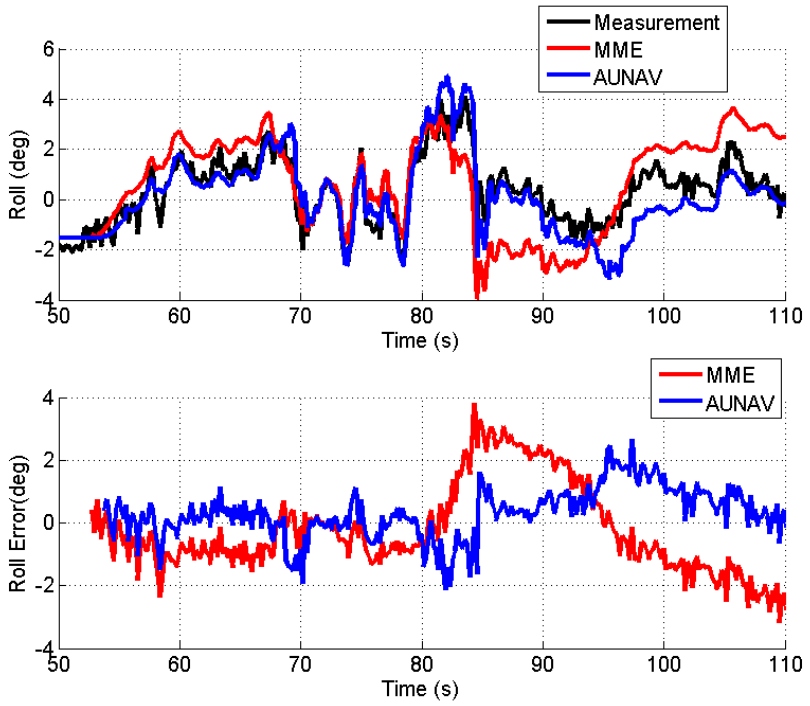

Figure 5.16: Comparison of Roll Angle Estimates During Dynamic Test.

input to the lateral state filter as in (2.51). The roll angle is a part of this filter, and thus it will be impacted by errors coming in through the measurement. The AUNAV estimate is much more accurate, by comparison, showing that the roll estimate of the AUNAV estimator is less subject to the lever arm effects than the MME.

Figure 5.17 shows the pitch estimate tracking the reference admirably. In fact the only instance of large error occurs at the same time where a problem arose in the roll estimate, likely due to the lever arm effects as well. It is also instructive to consider how the pitch estimate is calculated when analyzing these results. The pitch estimate for this filter is strictly an estimate of the road grade, calculated according to (2.26), and it completely neglects any effects of suspension pitch. The reference pitch angle is a measurement of the total pitch, so it includes the suspension pitch. The error plot shown in Figure 5.17 represents the difference between the pitch estimate and the reference pitch. This plot shows the same data as Figure 4.15 in Section 4.2. Recall the primary conclusion from the discussion there, which was that the Septentrio mounting angle is almost zero relative to the road grade.

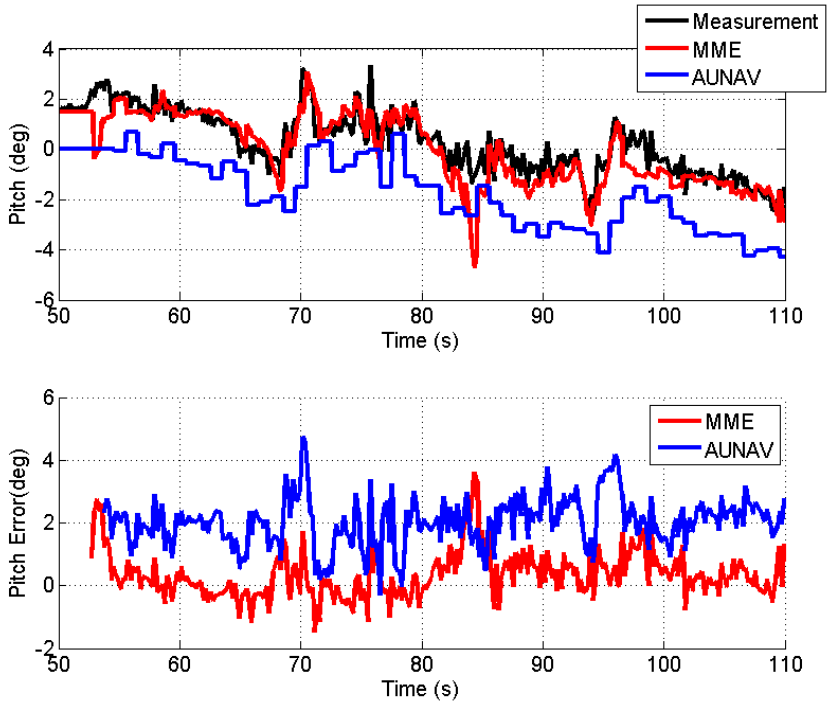

Figure 5.17: Comparison of Pitch Angle Estimates During Dynamic Test.

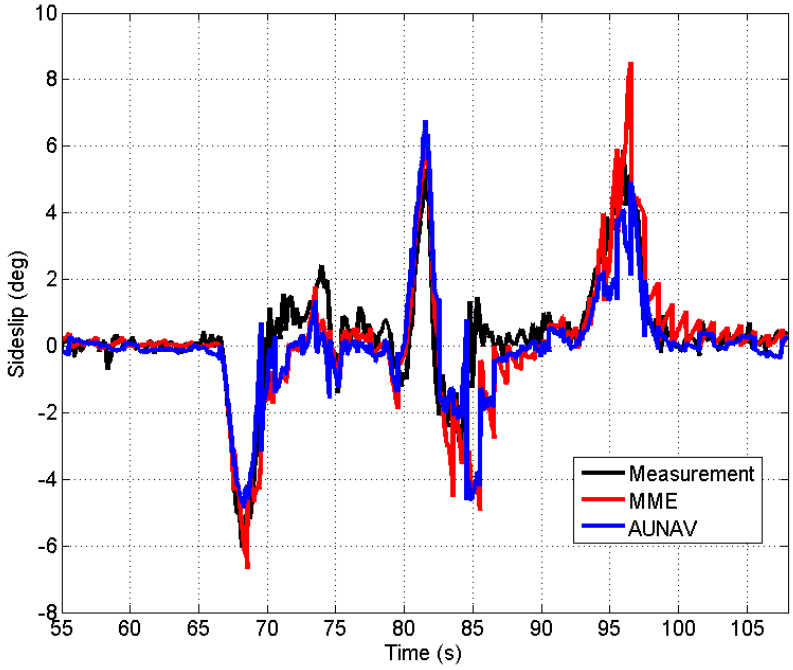

Figure 5.18: Comparison of Sideslip Angle Estimates During Dynamic Test.

The sideslip angle is plotted in Figure 5.18. Conclusions regarding the estimation quality shown here can be divided into two groups: the performance before $t \approx 84s$ and the performance after. Before this time, the filter outputs a very nice estimate of sideslip. Figure 5.18 shows the estimate following the reference smoothly and accurately during this time. This changes at $t \approx 84s$. Previously in this section the roll error that occurs at this time is discussed and attributed to the lever arm effects. The effects of this error show themselves here also. The sideslip estimate is a function of the forward velocity estimate and the lateral velocity estimate, per Equation (2.52). The lateral velocity estimate is coupled with the roll estimate in the lateral filter, thereby subjecting it to negative impact from any errors in the roll angle estimate. It has been mentioned several times now that an inaccurate roll angle estimate or an inaccurate estimate of the lateral accelerometer bias produces the exact same effect in the lateral velocity estimate, and that this effect is a sawtooth shaped structure. The reason for the sawtooth is simply that the filter is integrating the true acceleration plus some bias, either from roll and gravity or uncompensated sensor bias. This is what is happening in the sideslip estimate for the time $t > 84s$. The roll angle error grows, creating an artificial bias in the accelerometer, and the sawtooth shapes appear in the lateral velocity. As the roll angle begins to converge again, the magnitude of the saw teeth diminishes. The AUNAV estimator recovers more quickly from this problem than the MME does, and it produces in general a smoother estimate of sideslip than the MME.

5.3 Conclusion

The same experimental tests which were performed to validate the performance of the AUNAV estimator were used to validate modular filter. The tests were divided into two phases as in the case of the AUNAV estimator. Simulations and experimental data were used to validate the initialization phase of the modular filter. The experimental data comes from the same run used to validate the AUNAV estimator initialization, and the simulation was designed in Carsim to mimic the real world experiment. The results from the simulations

and real data are in agreement and show that given dynamic conditions over time the modular filter is able to accurately separate the bias and level angles despite large initial errors in the bias state. The MME estimator performs well during the dynamic estimation phase of the experiment, except for the instance where the lever arm effects come into play. The modular filter is found to be more susceptible to these effects than the AUNAV estimator. In short, the modular filter provides good estimation performance for ESC applications.

The AUNAV estimator performs better than the MME estimator by most comparisons. The roll angle estimates from the AUNAV estimator were found to be more accurate than the MME estimates in both tests. There is little performance difference when looking at the forward and vertical accelerometer biases, but the AUNAV estimator converges faster when estimating the lateral accelerometer bias. The velocity estimates are almost qualitatively the same, although the AUNAV estimator is observed to be slightly more accurate. The pitch angle estimates are also comparable, with the main difference being the suspension pitch. In this case the MME estimate is probably preferable over the AUNAV estimate, because it is always clear exactly what it is outputting. The MME, in steady state, always produces an unbiased estimate of the road grade. The AUNAV estimator, by contrast, may not have converged to the true pitch value. It was demonstrated in this thesis that the AUNAV pitch estimate is *close* to the true pitch value, and that correspondingly the forward bias estimate is *close* to the true value, but some ambiguity remains. This is of course due to the lack of the pitch rate gyroscope. There is no ambiguity in the road grade estimate from the MME estimator, making it the preferable pitch estimate on vehicles which have negligible suspension pitch. The AUNAV estimator strongly out performs the MME in regards to roll rate bias estimation, as shown in Section 5.1. On the other hand, the yaw rate bias is more accurately estimated by the MME estimator. This is because this estimate is only a function of the yaw rate signal and the course measurement. Therefore, misalignment errors in the yaw direction do not affect the bias estimate. Unmodeled misalignment errors will affect

the bias estimate of the AUNAV filter, as described in Section 4.1, although including these effects in the AUNAV model could potentially mitigate these errors.

It is no surprise that the AUNAV estimator outperforms the MME. The reason is simply that the AUNAV estimator is based on a more complete model of the underlying kinematic relationships. The more accurate the model is, the more accurate the estimator will be. This obvious conclusion actually draws a parallel with model based state estimation methods. The performance of those methods is also primarily a function of the model accuracy. The difference, however, is that the model parameters for the the kinematic AUNAV and MME estimators are measurable and practically unchanging. This is the primary advantage of the AUNAV and MME estimators. Finally, although the MME estimator is capable of accurately estimating important vehicle states for ESC systems, it is outperformed by the AUNAV estimator.

Chapter 6

Other GPS Applications: Tire Radius Estimation, Tire Pressure Monitoring, and Steering Misalignment Detection

6.1 Introduction

GPS information can be also be used to serve vehicle functions other than state estimation and navigation. In this chapter, GPS information is used to estimate the effective radius of the tires, to detect pressure drops in the tires, and to detect steering misalignment. The state estimates from the AUNAV estimator can also be used to aid in these processes. For example, road grade information is important for the pressure change detection system, because large road grades affect the tire loading and may impact the radius estimate. Furthermore, an accurate yaw rate signal is required for the steering misalignment detection algorithm presented here. The yaw rate gyroscope bias estimate produced from the state estimator can be used to remove the bias from the yaw rate signal in order to provide the needed level of accuracy. In short, this chapter will show how both direct GPS information and the state estimates produced from the AUNAV estimator can be used to aid in other vehicle safety functions beyond dynamic state estimation.

6.2 Tire Rolling Radius

Combining GPS with on board ESC sensors can aid in estimating the tire's rolling radius, which is an important parameter for various vehicle systems. An approach based on linear estimation techniques is presented in [31], where the authors combine GPS and wheel speed information with a vehicle model to obtain estimates of the rolling radius and the longitudinal stiffness of the *driven* wheels. In [11] the authors expand on the work done in [31]

by considering the adverse properties of linear estimation methods. Finding that noise in the measurement matrix (as opposed to merely having noise in the measurement vector) causes a bias in the linear least squares estimate, the authors propose using nonlinear optimization methods to solve the problem. Improvements to the nonlinear estimation strategy are given in [12], [13]. The approaches in all of these are for finding the longitudinal stiffness and rolling radius of the *driven* wheels. However, the rolling radius of the *undriven* wheels are estimated and discussed also. Equation (6.1) is the simple, important governing relationship which is exploited to produce the estimate of the undriven rolling radius.

$$V_x = R_{eff}\omega \tag{6.1}$$

Equation (6.1) describes the relationship between the forward velocity V_x, the rolling radius R_{eff}, and the wheel speed ω under the important condition that there is no slip. In [13] the radius estimate is obtained by solving Equation (6.1) in matrix (batch) form. For reference, Equation (6.2) expands the relationship to deal with slip by including the forward slip ratio σ.

$$V_x = \frac{R_{eff}\omega}{1-\sigma} \tag{6.2}$$

All of the above references focus on methods for estimating the driven wheel radius together with the longitudinal stiffness. These estimates are useful in themselves, however the authors also focus on how the stiffness in particular might be an indicator of the tire-road friction coefficient. This follows the hypothesis raised in [24], [25] that the longitudinal stiffness value could be an accurate predictor of the friction even at low slip values. In pursuing this investigation, the authors find that many things influence the stiffness estimate including tire pressure [13]. Therefore much discussion is given to how the stiffness might also be an indicator of tire pressure. In this thesis the focus is on how the estimate of the rolling radius might accurately indicate tire pressure or pressure changes, as opposed to the stiffness. While

minor consideration is given to this idea in [11], [12], and [13], this section of this thesis will explore how the radius estimate alone might be a successful predictor of tire pressure.

If the rolling radius estimate of the undriven wheel can successfully predict tire pressure, the method might be expanded to the driven wheels using relative radius estimation methods. Therefore consider the case where there is no slip, then Equation (6.1) can be used to form a linear Kalman filter. In this case the state estimate is held constant in between the recurring measurements of V_x/ω. For comparison, this is simply the recursive version of the approach shown in [13] for the undriven wheel radius. Equations (6.3) and (6.4) show the state vector and the measurement vector for the filter.

$$\hat{X} = \begin{bmatrix} R_{eff-fl} & R_{eff-fr} & R_{eff-rl} & R_{eff-rr} \end{bmatrix}^T \tag{6.3}$$

$$Y = \begin{bmatrix} \frac{V_{x-fl}}{\omega_{fl}} & \frac{V_{x-fr}}{\omega_{fr}} & \frac{V_{x-rl}}{\omega_{rl}} & \frac{V_{x-rr}}{\omega_{rr}} \end{bmatrix}^T + \underline{\gamma} \tag{6.4}$$

In (6.4), $\underline{\gamma}$ represents the measurement noise vector. The filter operates as a normal Kalman filter according to Equations (2.1-2.3). Since the wheel speeds are by definition the rotational speeds at the tire, the translational speeds also must be at the tire. These are denoted in (6.4) by $V_{x-fl}, V_{x-fr}, V_{x-rl}, V_{x-rr}$. Calculating these velocities is straightforward, and Equation (6.5) shows how this is done.

$$\begin{aligned} V_{x-fl} = V_{x-rl} &= V_x + \frac{1}{2} l_{tw} \omega_z \\ V_{x-fr} = V_{x-rr} &= V_x - \frac{1}{2} l_{tw} \omega_z \end{aligned} \tag{6.5}$$

In (6.5), ω_z is the yaw rate, and l_{tw} represents the track width.

Figure 6.1 shows the estimation performance in simulation. Carsim was used to simulate a simple scenario where a front wheel drive sedan drives in a straight line for 10 minutes. The estimate converges to the true value over time. Note that a small bias is present in the

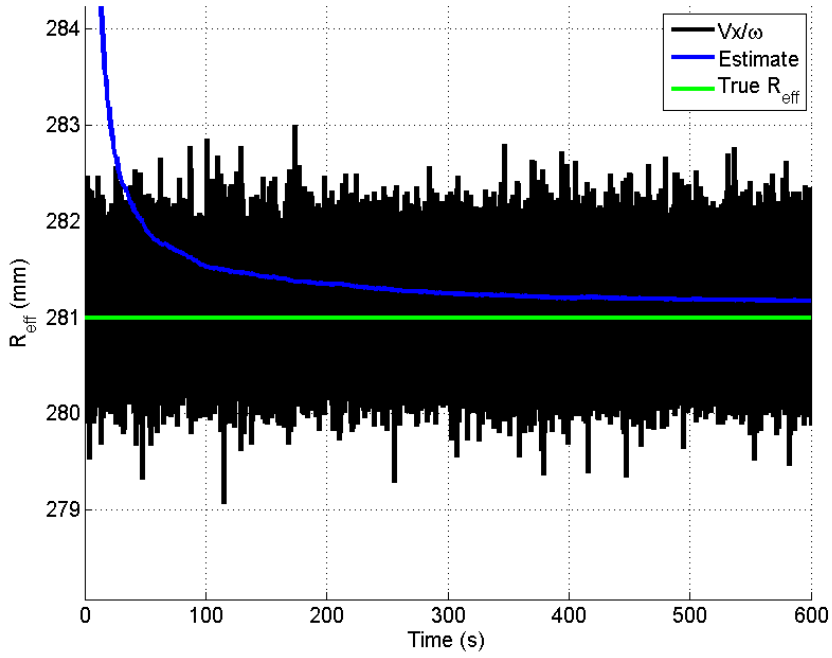

Figure 6.1: Undriven Wheel Effective Radius Estimate Convergence in Simulation.

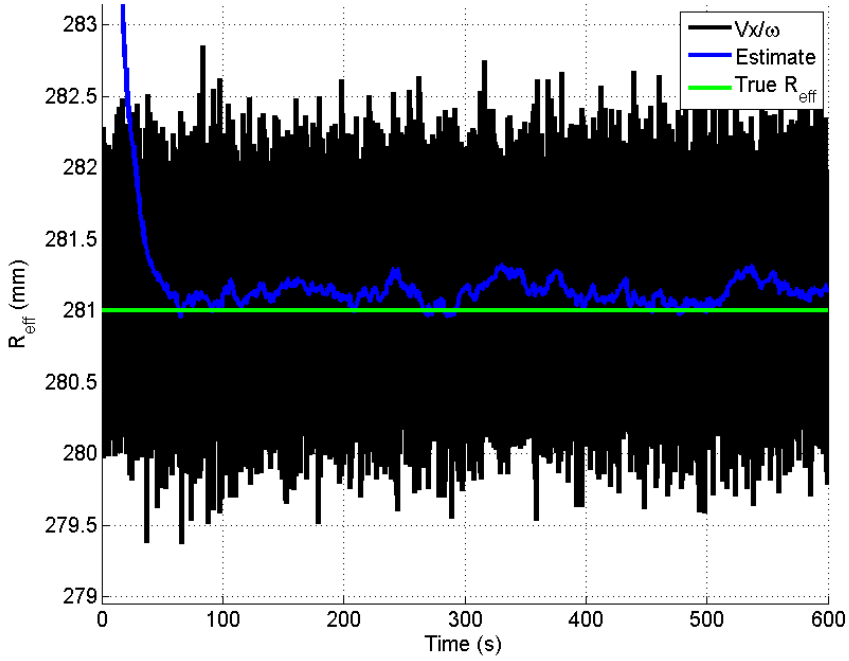

Figure 6.2: Undriven Wheel Effective Radius Estimate Convergence with Faster Tuning in Simulation.

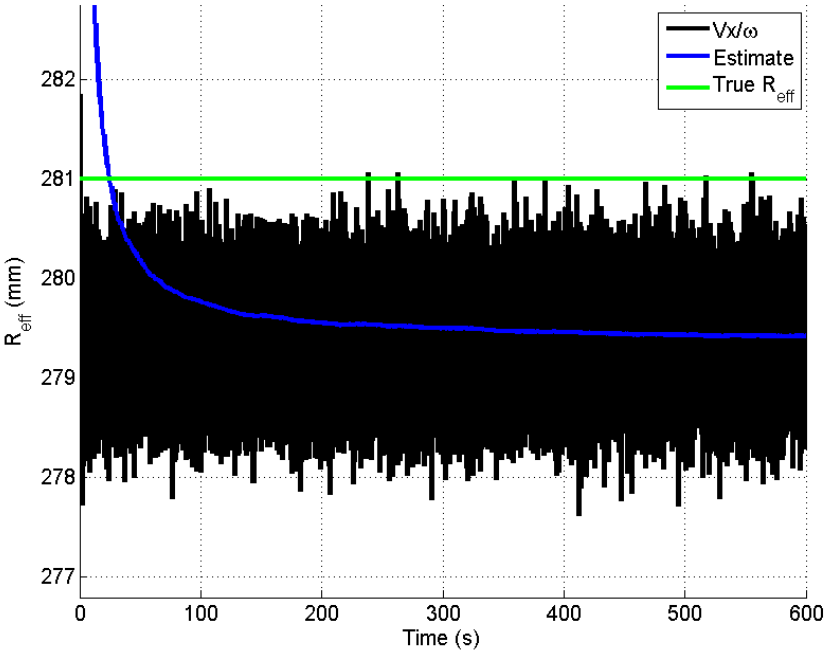

Figure 6.3: Driven Wheel Effective Radius Estimate Convergence in Simulation.

estimator. This is because even for the undriven wheel there exists a very tiny steady state slip of around 0.04% (according to Carsim's models). Yet the bias from this is miniscule, at only 0.1mm. The convergence speed is a function of the tuning, although there is a trade off between speed and smoothness. A faster tuning offers quicker convergence. This is also seen in Figure 6.2, where the tuning has been "sped up" for the same run. That is, the process noise covariance matrix (Q) value is larger. The sensor noise covariance matrix (R) value is held constant. In this case the filter converges more quickly, but it is not as smooth. Later discussions will focus on the importance of smoothness. Figure 6.3 shows the radius estimate of a driven wheel. The transient performance is exactly the same, (since the slip is constant), but the bias is larger because the slip is larger for the driven wheel. It violates the no slip assumption stated previously, and the bias is around 1.5mm. Figure 6.4 shows the radius estimate of the undriven wheel during a simulation of the vehicle driving in a figure 8 pattern at much lower speeds, demonstrating that the filter is able to operate under turning conditions.

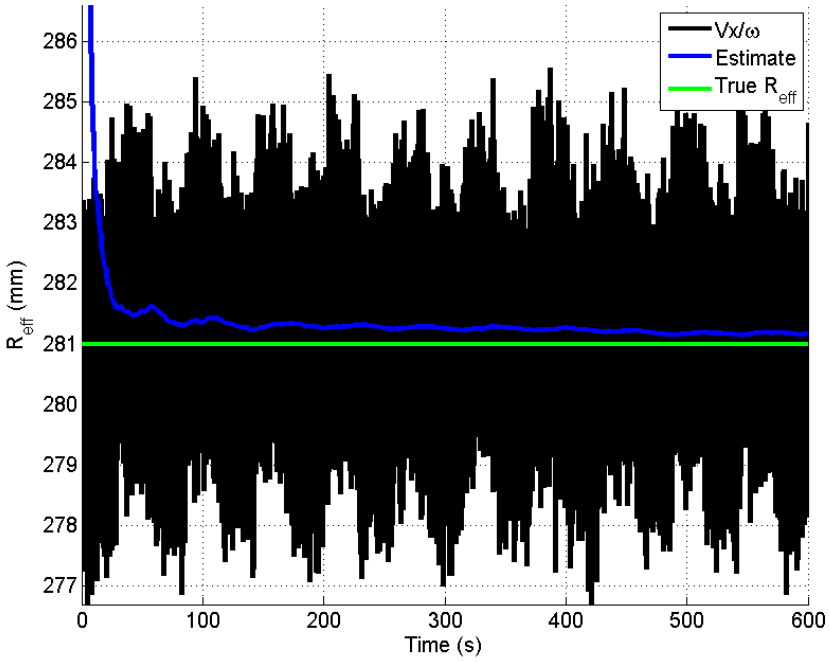

Figure 6.4: Undriven Wheel Effective Radius Estimate Convergence During Figure-8 Turning Maneuvers in Simulation.

6.3 Tire Pressure Monitoring

It is hypothesized that the effective rolling radius is a function of the tire pressure [11], [12], [13]. Since it has now been shown that GPS offers a way to very accurately estimate the effective radius, it is further hypothesized in this thesis that this estimate can be used to estimate the tire pressure and detect changes. In order to validate this hypothesis, four sets of data were collected on the loop at NCAT. The vehicle was driven at a nearly constant speed of 50mph for 45 minutes. Before and after each run, the pressure of all four tires was recorded. The G35 is a rear wheel drive vehicle, so the front tires are the undriven wheels. The front left and rear right tires were chosen as the variable tires, and they were run with a different pressure each time. While the other two tires were maintained at 33psi, the front left and rear right tires were set to 36, 32, 28, and 24 psi for each respective run. Figure 6.5 shows the estimates for the front left tire for each run. It should be noted that only the data from the straightaways is processed in the estimate. This is to limit any effects of weight transfer

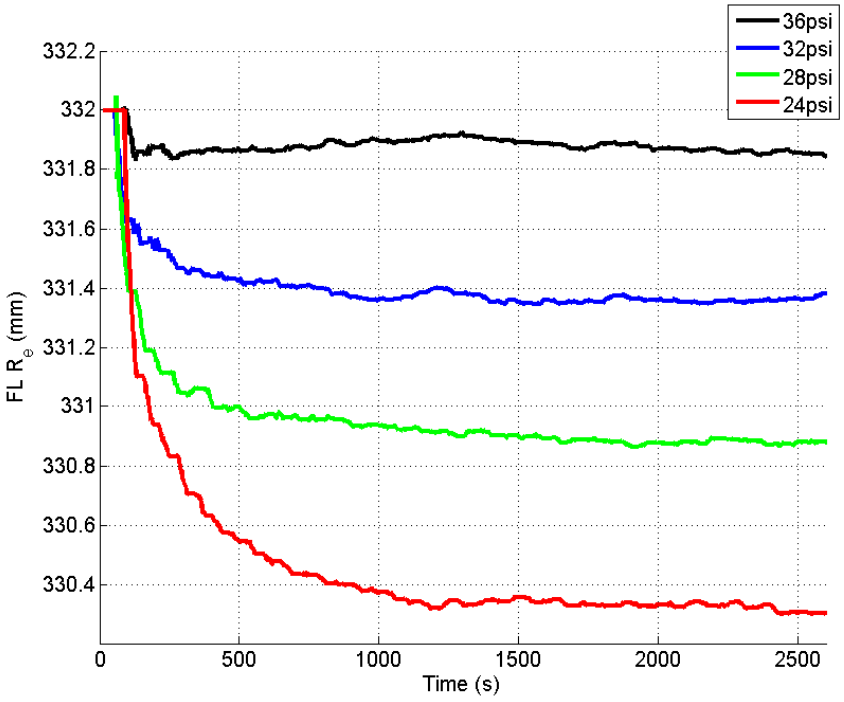

Figure 6.5: Front Left Tire Radius Estimates for Different Pressures.

on the tires, as the steeply banked turns might result in small tire deformations. Recall that the radius is specifically the *effective* or *rolling* radius, which means it is extremely difficult measure the true value. Yet it is known that this value lies between the loaded and unloaded radii [34]. The radius estimates were all found to lie within the measured loaded and unloaded radii as expected.

It can be seen that changes in tire pressure *do* result in changes in the rolling radius. The first pressure drop from 36 to 32 psi results in a radius change of 0.5mm. This might seem too small to track, but the estimator is clearly able to do so given enough data to average out the noise. This is because of the accuracy of the GPS velocity measurements, which are unbiased and have noise characteristics of approximately 5 $^{cm}/_{s}$ (1σ). The second drop from 32 to 28 psi results in a change of 0.5mm. The last drop from 28psi to 24 psi results in a change of 0.6mm. Figure 6.6 shows the estimates from the rear right tire. It is observed from this graph that the pressure drops correspond with decreases in the radius estimate of 0.5mm, 0.4mm, and 0.6mm. Recall that the method presented in this thesis

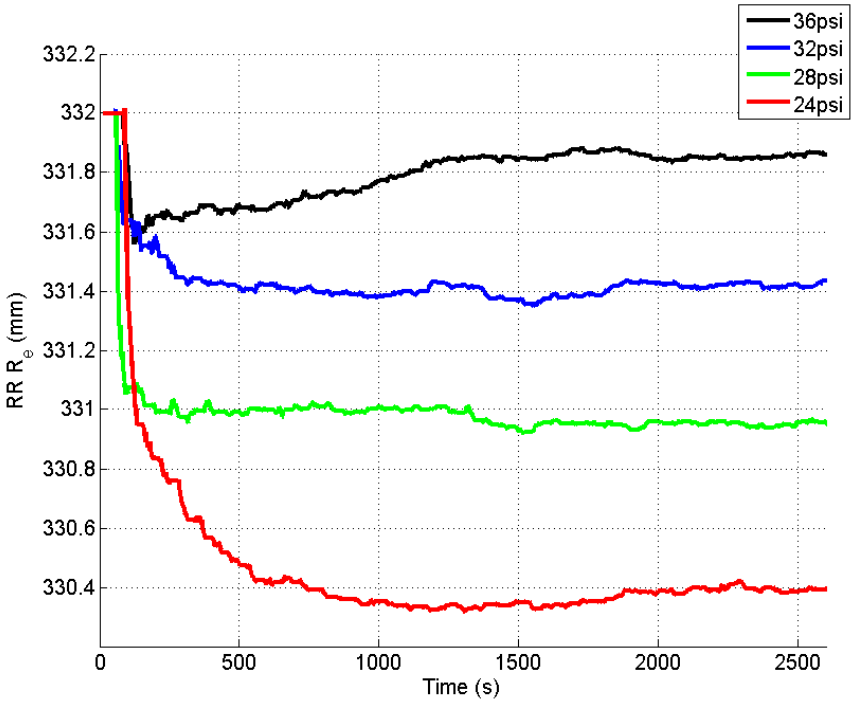

Figure 6.6: Rear Right Tire Radius Estimates for Different Pressures.

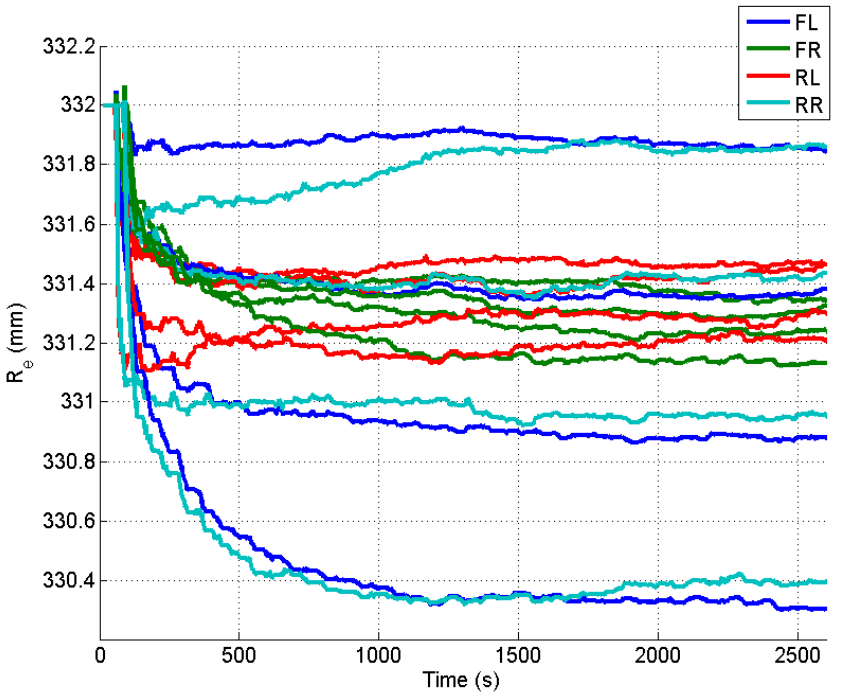

Figure 6.7: Estimates for All Four Tires for All Four Experiments.

results in a biased estimate when it is applied to the driven wheels. However, this bias does not affect the *change* in radius that results from the pressure. Therefore although the radius estimates of the driven wheels are not as accurate as those of the undriven wheels, they are still able to be used for pressure loss detection.

Looking at Figures 6.5 and 6.6 alone would suggest that these pressure ranges constitute an easily dividable space for a lookup table or other classifier. Indeed the gaps between signals is far larger than the noise, which is a good sign. However the repeatability of the estimates must be considered. Speaking qualitatively, if there is much variation between the radius estimates of a tire with a constant pressure, then it will be very difficult to isolate radius changes which correspond to pressure changes. Figure 6.7 shows the estimates for all four tires for all runs. Recall that the front right and rear left tires were maintained near 33psi for all four runs. It appears that all of the estimates for the control tires lie within the range 331.1mm and 331.5mm. Furthermore, the estimates for the front left and rear right tires lies within this range for the 32psi run. This indicates a degree of consistency across all four tires. From Figure 6.7 it can be inferred that if the radius estimate is within this range, then the tire pressure lies within 30psi to 34psi. Similar regions can be constructed for the other pressure ranges. The space from 331.1mm to 330.7mm corresponds to a pressure between 30psi and 26psi. Below 330.7mm, the tire pressure can be considered to be less than 26psi. Clearly these conclusions only apply to this particular data, and much more data would be needed to construct a reliable tire pressure predictor. However, Figures 6.5 - 6.7 show that indirect tire pressure monitoring is possible with GPS.

6.4 Steering Misalignment Detection

Steering misalignment is a problem because of the accelerated and irregular wear that it causes on the tires. Large misalignments are obvious and easily noticed by the driver, but smaller misalignment problems can be far more subtle. Even small misalignments have negative effects on tire wear and gas mileage, therefore it would be advantageous if the

vehicle were able to detect small misalignments and inform the driver, and GPS provides the ability to do this.

The phrase "steering misalignment" is used here to describe any condition which causes the steering system to produce an improper steer angle at the tires for a given steering wheel input. This may be caused by a variety of problems in the suspension or steering systems, but from a general standpoint the end effect is the same. The premise of the misalignment detection strategy presented in this thesis is that if the vehicle is driving straight, then the steer angle ought to be zero degrees. This is not exactly true, however, because of the steering effects of road crown. Most roads generally have a very small side slope in order to improve water drainage, and this slope will cause a very slight offset from zero in the steer angle. In this thesis, all tests were done on the same section of roadway, so the effects of road crown will be the same for all experiments. Therefore the road crown effects will be ignored, although the accurate estimate of the vehicle roll angle produced by the AUNAV estimator could be used to detect non standard road crowns.

The steering misalignment detection algorithm functions as follows. It requires the steering wheel angle (SWA), an accurate yaw rate signal, noise characteristics for the yaw rate sensor, and a baseline steer angle. The baseline steer angle is the steady state steer angle for straight driving when the vehicle is perfectly aligned, which can be ascertained the first time a production vehicle is driven. If there is no road crown, then the baseline steer angle will be zero. For the experiments in this thesis, the baseline steer angle was found to be approximately 2.25 degrees. The assumption is made that the Infiniti G35 was perfectly aligned before the experiments were made. Since road crown is assumed to be constant, the baseline steer angle is adjusted to be zero degrees, meaning that a constant 2.25 degrees is subtracted from all SWA measurements. This is because it is the *relative* steer angle compared to the baseline which is important, not necessarily the absolute steer angle. When it is determined that the vehicle is driving straight, as described in Section 2.3.2, a moving average (100 second window) of the SWA measurement is taken. The yaw rate gyroscope

bias estimate from the AUNAV filter is subtracted from the raw yaw rate signal during the process of determining straight driving conditions. If the average steer angle is offset from zero, when the vehicle is driving straight, then it is concluded that there is a misalignment problem.

Four experiments were conducted to validate the steering misalignment detection strategy. These tests were conducted using the Infiniti G35, which is instrumented as described at the beginning of Chapter 4. Each test consisted of driving 10 minutes north on I-85 towards Atlanta while maintaining a steady speed and with minimal lane changes. The first test was conducted without making any modifications to the vehicle, assuming that the alignment at that time was very good. Thus the alignment for the first test was taken as the baseline. For each subsequent test, misalignment error was manually introduced by adjusting the toe angle of the front right tire. The toe was adjusted by screwing the tie road in or out of the sleeve by some amount. In this experiment, the tie road was screwed into the sleeve, thereby shortening the overall length. Since this linkage is located behind the center of the tire, the change resulted in a toe out situation on the front right tire. The tie road was adjusted by making two turns on the screw for the first experimental run, 1/2 turn for the second, and 1 full turn for the final run. Figure 6.8 shows the raw SWA measurements and the corresponding moving averages for the experiments. The baseline steer angle of 2.25 degrees is also observed in this plot. Notice that the magnitude of the steer angle increases with the amount of alignment adjustments made. Figure 6.9 shows the moving averages of the steering wheel angles, the yaw rate signals for each run, and the yaw rate detection threshold for straight driving determination. It can be clearly seen that adjustments to the toe angle of the tire results in non zero steering angles when the vehicle is driving straight. Furthermore, the magnitude of the offset clearly corresponds with the degree to which the toe angle is adjusted. The yaw rate threshold for straight driving detection is shown as a reference. A threshold of 3σ was used, where the assumed standard deviation of the noise on the yaw rate signal is $\sigma = 0.5^{deg}/_{s}$. The yaw rate signals have been improved by subtracting

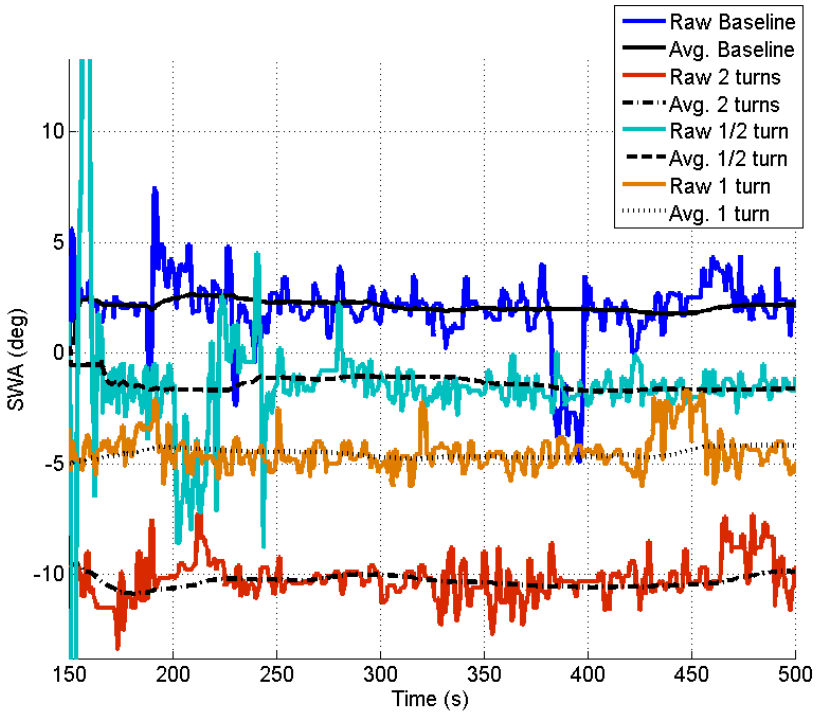

Figure 6.8: Steer Angle Signals for All Four Misalignment Experiments.

the estimates of the yaw rate bias provided by the AUNAV filter. In this case only a marginal improvement was obtained because of the high quality of the Crossbow gyroscopes. However less expensive sensors will have larger biases, making it necessary to use the bias estimate in order to be able to detect straight driving. Based on Figure 6.9 it is concluded that it is possible to detect even slight misalignments using GPS enhanced inertial sensor signals and the steering wheel angle.

6.5 Conclusions

It has been shown that GPS information can be combined with wheel speeds to very accurately estimate the rolling radius of the undriven wheels. The algorithm operates on the assumption that there is no slip. This assumption is valid for the undriven wheels except in the case of braking. This does not present an obstacle for the algorithm however, because the braking signals are readily available on the CAN bus and can be used as a toggle on

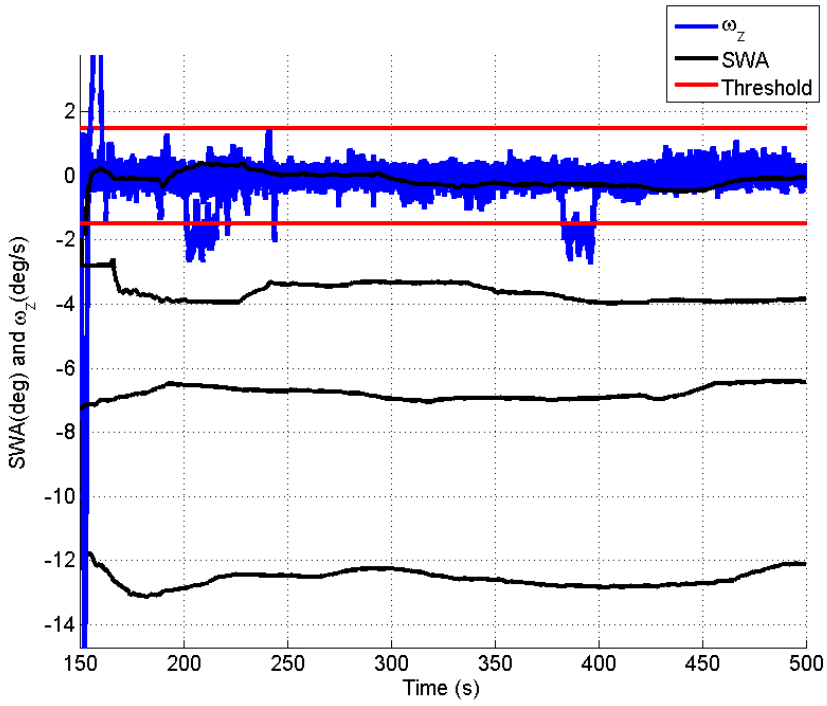

Figure 6.9: Averaged Steer Angle Signals and Yaw Rate Signals For All Four Misalignment Experiments.

and off switch for the estimator. It has also been shown that the rolling radius varies with tire pressure, presenting an opportunity for indirect tire pressure monitoring. That is, GPS offers the possibility of detecting pressure losses without actually needing pressure sensors in the tire. The TREAD act requires tire pressure monitoring systems (TPMS) to notify the driver if the tire pressure has dropped by 25% of the recommended cold pressure [1]. The recommended pressure for the G35 is 33psi, which means that any TPMS must be able to detect pressure drops of 8.25 psi. The results from Figure 6.7 show that this is extremely possible with the rolling radius estimate. Such a system would require many more data sets for analysis to provide statistically significant detection and alarm thresholds. However, the operation is simple: if the radius estimate drops below a certain value, then the system interprets that to correspond with a drop of a certain pressure. Given the data presently available, the design would be as follows. If the radius for one of the front tires drops below 331.1mm, then a drop of 4psi (12%) can be safely inferred. This is well within the accuracy

requirements of the TREAD act. Again, more data needs to be collected to improve the statistics of the detection algorithm, but the concept is shown feasible here. Future work would also deal with studying and isolating the effects of pressure on the radius. For instance, the effects of vertical loading conditions on the radius estimate needs to be studied, because these effects must be isolated from the pressure detection algorithm. Yet it has been shown that GPS can be used to accurately estimate the vehicle mass online [3]. If the mass is known, and more importantly if the change in mass is known, this information can be included in the process of inferring the pressure from the radius. Tire wear might also present a problem. If the change in radius due to tire wear is in the same range as the change in radius due to pressure, it may be very difficult to separate these effects. Yet for now it is concluded that the GPS aided estimate of the undriven wheel rolling radius offers great potential in determining pressure loss without needing tire pressure sensors. It has also been shown that the yaw rate signal, aided with the bias estimate from the AUNAV estimator, can be used with the steer angle sensor to detect steering misalignment. Data was collected on portions of I-85 north with various alignment adjustments to the steering linkage. The data confirmed the hypothesis that misalignment conditions result in steady state steer angle offsets when the vehicle is driving straight. Therefore if the vehicle is known to be driving straight via the yaw rate information and if there is a non zero steady state steering wheel angle, then it is accurately inferred that there is an alignment problem.

Chapter 7

Conclusions and Future Work

This thesis has shown the advantages of combining the increasingly common single antenna GPS systems with sensors which are already present on board vehicles equipped with ESC and RSC systems. First, the Automotive Navigation (AUNAV) estimator was presented. This estimator integrates GPS information with information from a 5 degree of freedom inertial sensor cluster to produce accurate estimates of the vehicle sideslip, attitude, velocity, and position along with accurate estimates of the inertial sensor biases. An analysis of the observability of the AUNAV estimator was also presented, showing that this estimator is fully observable under certain dynamic conditions. The AUNAV estimator was then contrasted with the Modified Modular Estimator (MME), which is a different approach to integrating the same sensors. Finally, methods of estimating the effective tire radius, detecting tire pressure changes, and detecting steering misalignment were presented. These methods all take advantage of GPS information and information provided by the AUNAV estimator.

The AUNAV estimator is a specialized form of the generic loosely coupled GPS/INS integration filter that is shown in [19]. It consists of the GPS and IMU functioning as two independent navigation systems, with the GPS serving as truth, and one extended Kalman filter. The states of the EKF are the errors in the INS solution, so the measurements used by the filter are the differences between the GPS and INS solutions. The estimates of the INS solution errors are added back to the original INS solutions and the corrected solution is used as the final estimate. The EKF also produces estimates of the inertial sensor biases, and these are used in the INS processing stage to improve the solutions. The AUNAV estimator differs in its design from the loosely coupled filter in two ways. First, the AUNAV estimator

does not use a pitch rate gyroscope because these are not available on modern commercial vehicles. Second, the AUNAV estimator uses the GPS course measurement as a measurement of heading (yaw) when the vehicle is driving straight. This is made possible because when the course and heading are equivalent when there is no sideslip. The AUNAV estimator also differs from the loosely coupled filter in that the AUNAV estimator is primarily used to estimate sideslip, while the loosely coupled filter is primarily analyzed for its positioning capabilities.

A study of the observability of the AUNAV estimator was conducted. Many authors have analytically determined the observability of the loosely coupled filter under various dynamic scenarios, making the important case that the observability is contingent on the dynamics. It was hypothesized in this work that the same conclusions apply for the AUNAV estimator, and this was tested using an observability test for linear time varying systems and simulations. The observability tests confirmed the hypothesis that the AUNAV estimator has the same observability characteristics as the loosely coupled filter. The AUNAV estimator requires acceleration changes in order to be made fully observable, but combinations of the attitude states and the inertial biases remain observable when the vehicle is not performing any acceleration. This means that dynamics are required in order for the estimator to be able to separate the attitude states from the bias states. Simulations were performed which confirmed the results of the observability checks. It was also found that the roll rate gyroscope bias converges when there are no dynamics. The yaw measurement gives the AUNAV estimator an observability advantage over the loosely coupled filter by making the yaw angle and the yaw rate gyroscope bias observable when the vehicle is driving straight.

The performance of the AUNAV estimator was demonstrated using experimental data collected with an Infiniti G35 sedan which is instrumented with various GPS and inertial sensors, along with an attitude determination system. It was shown that the AUNAV estimator is able accurately estimate the sideslip, roll, and pitch during dynamic maneuvers. It was also shown that the estimator is able to accurately separate the attitude states and

the biases when dynamic maneuvers are performed over time, confirming the observability analysis. Finally, it was shown that the AUNAV estimator is able to accurately estimate sideslip which is increasing a rates as slow a 1.5 $^{deg}/_{s}$.

The performance of the AUNAV estimator was also compared with the performance of the MME estimator using the same data from the previously experiments. The MME estimator is based on the approach presented in [6]. It consists of several independent GPS/INS filters which estimate states along the different axes. The roll angle must be additionally high pass filtered from a lumped state in the lateral filter. The pitch angle is obtained by estimating the road grade as in [4]. The sideslip angle is computed from the lateral and forward velocity estimates. It was found that in general the AUNAV estimator outperforms the MME estimator in almost all cases. This is because the AUNAV estimator uses a more complete model of the kinematic relationships than the MME does. However, the MME does a better job of estimating the yaw rate gyroscope bias. This is because its simplified model is more isolated from the effects of misalignment than the AUNAV estimator. In fact, the MME heading filter is not subject to yaw angle misalignment effects at all. The AUNAV filter combines all yaw rate errors, including misalignment, into the bias estimate, making it a less accurate estimate of the pure sensor bias.

A simple method of using GPS and wheel speed sensors to estimate the absolute effective rolling radius of the wheels was also shown. The method was demonstrated in simulation to be unbiased if the no slip assumption holds. While the estimate for the driven wheels is biased due to wheel slip, the estimate for the undriven wheels can serve as an absolute radius for any relative calculations for the driven wheels. There is no reference method or "truth" measurement for the effective radius with which to compare the estimate in real experiments, yet it is known that the effective radius lies within the loaded and unloaded radii. Several runs were performed on the G35 on the NCAT loop, and the radius estimates for all four tires were within this range for all runs. This served as a sanity check, showing that while there may not be an exact truth for comparison, the estimates from the experiment are reasonable.

A method for inferring the tire pressure on the basis of the rolling radius estimate is also proposed. Tests were conducted showing that the rolling radius estimate does vary according to tire pressure, and that these changes are distinguishable from the noise on the estimates. This leads to the conclusion that a simple hypothesis tester can be used to detect pressure drops of half of the percentage required by the TREAD act [1]. Future work requires many more data runs to improve the statistics of such a detector and to investigate other sources of change in the radius estimate. This is because the radius is only a good indicator of pressure if the changes resulting from pressure changes can be isolated.

A method of using the yaw rate bias estimate from the AUNAV estimator, the yaw rate signal, and the steering wheel angle to detect steering misalignment was presented. This method assumes that an aligned vehicle will have a zero steady state steer angle when driving straight if there is no road crown. For this thesis, the effects of road crown are ignored, although they could potentially be incorporating using the roll angle estimate provided by the AUNAV estimator. Experiments were conducted by altering the toe angle of the front right tire of the G35 and collecting data along a 10 minute section of I-85 north. It was found that tire misalignment results in a non zero steady state steer angle when driving straight, and that the corrected yaw rate signal and the steer angle signal can be used to accurately detect even subtle steering misalignments. If a non zero steady state steering angle is present when the corrected yaw angle is zero, then a flag is set indicating a misalignment condition.

There are many important avenues of future work stemming from this thesis. First, a more detailed analysis on the sensitivity of the AUNAV estimator to sensor quality needs to be done. The bias stability of the inertial sensors in particular plays an important role in the overall observability and operation of the filter. It was claimed in this thesis that once the biases are properly identified, the estimator can sustain accurate estimates of the level angles and the biases for a period of time even when the vehicle is not undergoing any dynamics. The bias stability is the main factor determining how long this period of time is. Furthermore, the length of this period of time is an extremely important system parameter

of the AUNAV estimator, and it would have to be accurately known in order to safely implement the estimator in commercial vehicles. Sensor quality also affects the amount of time that the estimator can sustain accurate state estimates during a GPS outage. This is another very important system parameter because GPS outages are ultimately inevitable. Another avenue of future work would be to investigate other methods of determining when the vehicle is driving straight. One potential method would be to run a model switching algorithm. The switching would be between an instance of the AUNAV filter which assumes straight driving and one which does not. It currently remains to be seen whether there are any advantages in using such a method or any other method beyond the current detection strategy. The effects of the bias stability on the straight driving detection method also requires additional study. It might be possible for the bias to cause the filter to never detect periods of straight driving, making it even more difficult to estimate the bias. This could potentially cause rapid degradation in the filter performance.

More data is required to improve the tire pressure detection algorithm, in order to make the results statistically significant and to more accurately quantify the accuracy and resolution. Another current problem is that there is no way to establish a baseline radius for the tire. The absolute radius can be accurately estimated for the undriven wheels, but there is no way of knowing as a baseline what pressure a given radius corresponds to. A solution to this problem would be essential to real world implementation. The effects of tire loading on the radius estimate also needs to be studied, as the pressure effects have to be isolated from these. The radius estimates have a very fine resolution, and it is likely that large loading variations could cause true deformations in the tires within the order of magnitude of the estimation resolution. The steering misalignment algorithm would also benefit from more data for the same reasons as the tire radius estimation algorithm. That is, more data is required in order to statistically set detection thresholds. The effects of road crown also need to be investigated. It is possible for non standard road crowns to cause the system to

falsely detect a steering misalignment. The roll angle estimate could be used to mitigate this problem, but this has not yet been studied.

This thesis has shown the benefits of combining single antenna GPS information with the signals from standard ESC/RSC sensor clusters. These results can be applied *today*, from the standpoint of sensor availability, to any vehicle equipped with a navigation package and ESC/RSC. Navigation packages are becoming increasingly ubiquitous on passenger vehicles, making these solutions extremely cost effective. No additional sensors are required. Furthermore these solutions offers the advantage of robustness to vehicle parameters, since they does not use any such parameters. Neither changes in vehicle loading, new sets of tires, nor after market sway bar installations will cause any problems for the systems presented here. In conclusion, the solutions offered here present great opportunity to improve vehicle safety without one cent of additional materials cost, and it is all made possible with the Global Positioning System.

Bibliography

[1] National Highway Traffic Safety Administration. Federal motor vehicle safety standards; tire pressure monitoring systems; controls and displays. Technical report, Department of Transportation, http://www.nhtsa.gov/cars/rules/rulings/tirepresfinal/index.html, 2000.

[2] Rusty Anderson and David M. Bevly. Estimation of slip angles using a model based estimator and GPS. In *Proceedings of the American Control Conference*, volume 3, pages 2122–2127 vol.3, June-2 July 2004.

[3] Hong S. Bae, Jihan Ryu, and J. Christian Gerdes. Road grade and vehicle parameter estimation for longitudinal control using GPS. In *Proceedings of the IEEE Conference on Intelligent Transportation Systems*, August 2001.

[4] H.S. Bae and J. Christian Gerdes. Command modification using input shaping for automated highway systems with heavy trucks. In *Proceedings of the American Control Conference*, volume 1, pages 54–59 vol.1, June 2003.

[5] V.L. Bageshwar, D. Gebre-Egziabher, W.L. Garrard, and T.T. Georgiou. A Stochastic Observability Test for Discrete-Time Kalman Filters. *JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS*, 32(4), 2009.

[6] David M. Bevly. GPS: A low-cost velocity sensor for correcting inertial sensor errors on ground vehicles. *Journal of Dynamic Systems, Measurement, and Control*, 126:255–264, June 2004.

[7] David M. Bevly, Robert Daily, and William Travis. Estimation of critical tire parameters using GPS based sideslip measurements. In *Proceedings of the SAE Dynamics, Stability, and Controls Conference*, number 2006-01-1965, pages 87–94, February 2006.

[8] David M. Bevly, Robert Sheridan, and J. Christian Gerdes. Integrating INS sensors with GPS velocity measurements for continuous estimation of vehicle sideslip and tire cornering stiffness. In *Proceedings of the American Control Conference*, June 2001.

[9] D.M. Bevly, J. Ryu, and J.C. Gerdes. Integrating INS sensors with GPS measurements for continuous estimation of vehicle sideslip, roll, and tire cornering stiffness. *Intelligent Transportation Systems, IEEE Transactions on*, 7(4):483–493, Dec. 2006.

[10] F.X. Cao, D.K. Yang, A.G. Xu, J. Ma, W.D. Xiao, C.L. Law, K.V. Ling, and H.C. Chua. Low cost sins/gps integration for land vehicle navigation. In *Intelligent Transportation Systems, 2002. Proceedings. The IEEE 5th International Conference on*, pages 910–913, 2002.

[11] C.R. Carlson and J.C. Gerdes. Identifying tire pressure variation by nonlinear estimation of longitudinal stiffness and effective radius. In *Proceedings of AVEC 2002 6th International Symposium of Advanced Vehicle Control*, 2002.

[12] C.R. Carlson and J.C. Gerdes. Nonlinear estimation of longitudinal tire slip under several driving conditions. In *American Control Conference, 2003. Proceedings of the 2003*, volume 6, pages 4975–4980. IEEE, 2003.

[13] C.R. Carlson and J.C. Gerdes. Consistent nonlinear estimation of longitudinal tire stiffness and effective radius. *Control Systems Technology, IEEE Transactions on*, 13(6):1010–1020, 2005.

[14] J.L. Crassidis and J.L. Junkins. *Optimal Estimation of Dynamic Systems*, volume 2. Chapman & Hall, 2004.

[15] Robert Daily, William Travis, and David M. Bevly. Cascaded observers to improve lateral vehicle state and tire parameter estimates. *International Journal of Vehicle Autonomous Systems*, 5:230–255, 2007.

[16] G. Dissanayake, S. Sukkarieh, E. Nebot, and H. Durrant-Whyte. The aiding of a low-cost strapdown inertial measurement unit usingvehicle model constraints for land vehicle applications. *IEEE Transactions on Robotics and Automation*, 17(5):731–747, 2001.

[17] NHTSAs National Center for Statistics and Analysis. Traffic safety facts 2008 data. Technical report, National Highway Transportation Safety Administration, http://www-nrd.nhtsa.dot.gov/Pubs/811162.PDF, 2008.

[18] J. Gao, MG Petovello, and ME Cannon. Development of precise GPS/INS/wheel speed sensor/yaw rate sensor integrated vehicular positioning system. In *Proceedings of the National Technical Meeting of the Institute of Navigation (ION NTM06*, volume 2, pages 780–792, 2006.

[19] S. Gleason and D. Gebre-Egziabher, editors. *GNSS Applications and Methods*. Artech House Publishers, 685 Canton Street Norwood, MA 02062, 2009.

[20] S. Godha and ME Cannon. Integration of DGPS with a low cost MEMSbased Inertial Measurement Unit (IMU) for land vehicle navigation application. In *Proceedings of the 18th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS05)*, pages 333–345, 2005.

[21] D. Goshen-Meskin and I.Y. Bar-Itzhack. Observability analysis of piece-wise constant systems. i. theory. *Aerospace and Electronic Systems, IEEE Transactions on*, 28(4):1056 –1067, October 1992.

[22] D. Goshen-Meskin and I.Y. Bar-Itzhack. Observability analysis of piece-wise constant systems. ii. application to inertial navigation in-flight alignment [military applications]. *Aerospace and Electronic Systems, IEEE Transactions on*, 28(4):1068 –1075, October 1992.

[23] P.D. Groves and Ebooks Corporation. *Principles of GNSS, inertial, and multisensor integrated navigation systems*. Artech House, 2008.

[24] F. Gustafsson. Slip-based tire-road friction estimation* 1. *Automatica*, 33(6):1087–1099, 1997.

[25] F. Gustafsson, M. Drevo, M. Lofgren, N. Persson, and H. Quicklund. Virtual sensors of tire pressure and road friction. 2001.

[26] Sinpyo Hong, Man Hyung Lee, Ho-Hwan Chun, Sun-Hong Kwon, and J.L. Speyer. Observability of error states in gps/ins integration. *Vehicular Technology, IEEE Transactions on*, 54(2):731–743, March 2005.

[27] Sinpyo Hong, Man Hyung Lee, Ho-Hwan Chun, Sun-Hong Kwon, and J.L. Speyer. Experimental study on the estimation of lever arm in gps/ins. *Vehicular Technology, IEEE Transactions on*, 55(2):431–448, March 2006.

[28] Sinpyo Hong, Man Hyung Lee, Sun Hong Kwon, and Ho Hwan Chun. A car test for the estimation of gps/ins alignment errors. *Intelligent Transportation Systems, IEEE Transactions on*, 5(3):208–218, Sept. 2004.

[29] Henrik Jansson, Ermin Kozica, Per Sahlholm, and Karl Henrik Johansson. Improved road grade estimation using sensor fusion. In *Proceedings of the 12th Reglermöte in Stockholm, Sweden*, 2006.

[30] Peter Lingman and Bengt Schmidtbauer. Road slope and vehicle mass estimation using kalman filtering. In *Proceedings of the 19th IAVSD Symposium*, Copenhagen, Denmark, 2001.

[31] S.L. Miller, B. Youngberg, A. Millie, P. Schweizer, and J.C. Gerdes. Calculating longitudinal wheel slip and tire parameters using GPS velocity. In *American Control Conference, 2001. Proceedings of the 2001*, volume 3, pages 1800–1805. IEEE, 2001.

[32] NASA. Robot coordinate frames and motion direction. Online, May 2009.

[33] American Association of State Highway and Transportation Officials. *A policy on geometric design of highways and streets*. American Association of State Highway and Transportation Officials, 444 North Capitol Street, N.W, Suite 249, Washington, DC 20001, 5 edition, 2004.

[34] R. Rajamani. *Vehicle dynamics and control*. Springer, 2006.

[35] Ihnsoek Rhee, Mamoun F. Abdel-Hafez, and Jason L. Speyer. Observability of an integrated gps/ins during maneuvers. *Aerospace and Electronic Systems, IEEE Transactions on*, 40(2):526–535, April 2004. read.

[36] Jonathan Ryan, David Bevly, and Jianbo Lu. Robust sideslip estimation using GPS road grade sensing to replace a pitch rate sensor. In *2009 IEEE International Conference on Systems, Man, and Cybernetics*, pages 2026 –2031, oct. 2009.

[37] Jonathan Ryan, Jianbo Lu, and David Bevly. State estimation for vehicle stability control: A kinematic approach using only gps and vsc sensors. In *Proceedings of the ASME 2010 Dynamic Systems and Control Conference*, Sept. 2010.

[38] Jihan Ryu and J. Christian Gerdes. Estimation of vehicle roll and road bank angle. In *American Control Conference, 2004. Proceedings of the 2004*, volume 3, pages 2110–2115 vol.3, June-2 July 2004.

[39] Jihan Ryu and J. Christian Gerdes. Integrating inertial sensors with global positioning system (GPS) for vehicle dynamics control. *Journal of Dynamic Systems, Measurement, and Control*, 126(2):243–254, 2004.

[40] Jihan Ryu, Eric J. Rossetter, and J. Christian Gerdes. Vehicle sideslip and roll parameter estimation using GPS. In *Proceedings of the AVEC International Symposium on Advanced Vehicle Control*, 2002.

[41] Per Sahlholm and Karl Henrik Johansson. Road grade estimation for look-ahead vehicle control. In *Proceedings of the 17th IFAC World Congress*, Seoul, Korea, 2008.

[42] L. M. Silverman and H. E. Meadows. Controllability and observability in time-variable linear systems. *SIAM Journal on Control*, 5(1):64–73, 1967.

[43] R.F. Stengel. *Optimal control and estimation*. Dover Pubns, 1994.

[44] H. Eric Tseng, Li Xu, and Davor Hrovat. Estimation of land vehicle roll and pitch angles. *Vehicle System Dynamics*, 45:433–443, May 2007.

[45] Yunchun Yang and J.A. Farrell. Magnetometer and differential carrier phase GPS-aided INS for advanced vehicle control. *Robotics and Automation, IEEE Transactions on*, 19(2):269–282, Apr 2003.