

Detection and Removal of Erroneous GPS Signals Using Angle of Arrival

by

Thomas L. Bitner

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
December 14, 2013

Keywords: Relative Positioning, Spoofing, GPS/INS

Copyright 2013 by Thomas L. Bitner

Approved by:

David Bevly, Chair, Professor of Mechanical Engineering
Dan Marghitu, Professor of Mechanical Engineering
Song-yul Choe, Professor of Mechanical Engineering

Abstract

In this thesis, an algorithm for detecting multipath and spoofed GPS signals based on signal angle of arrival (AOA) is developed. As the first step in determining the AOA of the signals, a multi-antenna array of GPS antennas is used to determine the attitude of a test vehicle by calculating precise relative positioning vectors (RPVs) between the antennas. The RPVs are calculated with the real-time kinematic (RTK) positioning algorithm, which allows for an RPV error on the order of a centimeter. The precision of the RPVs allows for sub-degree accuracy of the attitude angles.

A GPS/INS extended Kalman filter is used to propagate the attitude estimate between GPS measurements. The propagated attitude estimate allows for estimation of the RPVs between the antennas without using GPS measurements. To check an incoming set of GPS measurements, an expected AOA with respect to the antennas is computed using the estimated RPVs and the known unit vectors to the GPS satellites. The actual AOAs may be estimated using the incoming GPS measurements and then compared with the expected AOAs. If the difference between the expected and estimated AOA for a signal is not within a specified threshold, the signal may be rejected as a faulty signal. Single-differenced pseudo-range measurements and single-differenced carrier phase residuals are explored as alternate metrics for determining faulty signals.

Finally, AOA and the alternate metrics are experimentally tested for their abilities to detect multipath signals and are compared with each other. The single-differenced carrier phase residual proved to be the most reliable multipath detection metric when used with calculated RPVs, having a near perfect success rate with the collected data. Using the carrier phase residual metric with attitude-generated RPVs proved to be less effective with a success rate of approximately 50%. The AOA approach had a similar success rate of about 50%,

with worse results in instances of repeated multipath. The single-differenced pseudorange metric proved unable to reliably detect multipath due to the short antenna baselines and relatively large error on pseudorange measurements. The experimental results are validated by comparing positioning solutions before and after rejecting the signals as well as observing the pseudorange measurements of signals detected as multipath.

Acknowledgments

First and foremost I would like to thank God, through whom all things are possible and from whom all blessings flow. I have Him to thank for all of the gifts that have allowed me to be successful in all that I do.

I would like to thank my graduate adviser, Dr. David Bevly, as well as my undergraduate advisers, Dr. Chad Rodekohr and Dr. Brian Beasley, and Dr. James Wanliss for their support and guidance throughout my studies. It would have been impossible to achieve all that I have without their help.

I would like to thank the members of the GPS and Vehicle Dynamics Lab for their help and support during my time in Auburn. It has truly been a pleasure to work with so many talented engineers.

I would like to thank my parents for raising me and for instilling in me many of the qualities that make me who I am. I would also like to thank my family in general for their support throughout my time in college and graduate school. Although the past six years have been full of challenges, I have always had the great comfort of knowing that so many people are so proud of what I am doing.

Finally, I would like to thank Christina Clemens for her loving support in all that I do. She has been especially supportive during my time in graduate school, for which I am eternally grateful.

Table of Contents

Abstract	ii
Acknowledgments	iv
List of Tables	viii
List of Figures	ix
1 Introduction	1
1.1 Motivation	2
1.2 Previous Work	2
1.3 Contributions	4
1.4 Thesis Outline	4
2 The Global Positioning System	6
2.1 Signal Structure	6
2.2 Measurements	8
2.2.1 Pseudorange Measurement	8
2.2.2 Carrier Phase Measurement	9
2.3 Multipath Error	10
2.4 Spoofing	11
2.4.1 Spoofing Scenarios	11
2.5 Current Anti-Spoofing Measures	12
3 Attitude Determination	14
3.1 Differential Carrier Phase Positioning	14
3.1.1 Calculating Float Ambiguities	17
3.1.1.1 Time Update	17
3.1.1.2 Measurement Update	18

3.1.1.3	Implementation	20
3.1.2	LAMBDA method	21
3.1.3	Calculating RPVs	24
3.1.4	Single-differenced cycle ambiguities	24
3.2	Rotating from ECEF to ENU	26
3.3	Calculating Attitude	28
4	GPS/INS Integration	33
4.1	GPS/INS Kalman Filter	33
4.2	IMU Errors	36
4.2.1	Bias Error	37
4.2.2	Random Error	37
5	Faulty Signal Detection Algorithm	42
5.1	Initialization Stage	42
5.2	Check Stage	43
5.2.1	Calculate AOA	43
5.2.2	Alternate AOA Checking Strategies	48
5.2.2.1	Single-Differenced Pseudorange	48
5.2.2.2	Single-Differenced Carrier Phase	49
5.3	Update stage	62
5.3.1	Choosing the correct RPV	63
5.3.1.1	RPV magnitude	64
5.3.1.2	Change in RPV	65
5.3.1.3	Expected RPV	65
5.3.2	Calculating Single-Differenced Ambiguities	66
6	Experimental Results	67
6.1	Experimental Equipment	67
6.2	Testing Locations	68

6.3	Performance of Faulty Signal Rejection Methods	69
6.3.1	Calculate AOA	69
6.3.2	Single-Differenced Pseudorange	70
6.3.3	Carrier Phase Residual	71
6.3.4	Calculating Single-Differenced Ambiguities	72
6.4	Results Validation	72
6.5	Effects of Satellite Geometry on Positioning	75
7	Conclusions and Future Work	77
7.1	Conclusions	77
7.2	Future Work	78
	Bibliography	80
A	Determination of Position	82
B	Kalman Filter	85
B.1	System	85
B.2	Time Update	86
B.3	Measurement Update	87

List of Tables

3.1	Parameters for use in Equations (3.16) through (3.19)	20
6.1	Typical attitude accuracy for Septentrio PolaRx2e receiver. Roll accuracy assumes a 90 degree separation between the antennas.	67

List of Figures

1.1	Block diagram overview of the signal rejection system. Text in parentheses denotes the appendix or chapter and section which contains information about each block.	5
2.1	Components of a GPS signal	7
2.2	Diagram of multipath error due to signal reflection from building.	10
3.1	ECEF and ENU coordinates	28
3.2	Antenna configuration shown in vehicle coordinate frame. View is top-down onto the vehicle. Figure recreated from figure in Septentrio PolarEx manual.	30
3.3	Euler angle sequence (recreated from Septentrio manual) [1]	32
4.1	Unfiltered stationary gyro measurements over 3 axes. Red lines denote mean measurement.	38
4.2	Gyro data filtered at several bandwidths by low pass filter.	39
4.3	Monte Carlo simulation of random walk. 3σ envelope shown in black.	41
5.1	2-D example of AOA. Solid red lines depict LOS to GPS satellite.	45
5.2	Expected AOA vs. apparent AOA error for a given single-differenced carrier phase measurement error. Note that each curve is evaluated up until the point at which the argument for inverse cosine does not yield a real angle.	47
5.3	Plot of baseline length vs. number of potential AOAs.	52
5.4	Attitude accuracy as a function of baseline length.	53
5.5	Graphical representation of multiple AOAs that fulfill the remainder requirement of Equations (5.15) and (5.16). The expected AOA is shown as a blue solid line, while incorrect AOAs are shown as red solid lines. Red and blue dashed lines represent error bounds based on attitude estimation error from attitude filter.	54

5.6	Expected AOA vs. expected error in L , plotted for multiple values total angle error. Notice that once the error reaches $\pm.5$ cycles, no signals can be rejected.	56
5.7	Representation of AOA error bounds projected onto a hemisphere for two orthogonal baselines. Antenna baseline 1 extends from the origin outward along the X-axis while baseline 2 extends from the origin outward along the Y axis. The pair of blue bands represent error bounds on an AOA for baseline 1 while the red bands represent error bounds on an AOA for baseline 2. The error bounds will be used to determine the probability of accepting a false signal.	58
5.8	Differential surface area dA as it relates to $d\theta_1$ and $d\theta_2$.	60
6.1	Infiniti G35 with Septentrio GPS antennas.	68
6.2	Perceived AOA error for a data run. Black circles depict AOA error from the predicted RPV while red circles depict AOA error from the calculated RPV.	70
6.3	Remainders of single-differenced carrier phase measurement. Red circles represent residuals generated from calculated RPV, black circles represent residuals generated from predicted RPV.	71
6.4	Comparison of position estimate with and without rejected signals.	73
6.5	The upper sub-plot shows the carrier phase residual measurement for an antenna pair during a data run. The portion of graph above the .1-cycle threshold indicates the detection of a faulty signal. The lower sub-plot is the pseudorange residual for one of the antennas, showing the measured pseudorange differs significantly from the expected pseudorange when a faulty signal is detected.	74
6.6	Comparison of positioning with and without rejected signals. In this particular case, rejecting the signals yields a worse positioning solution.	76

Nomenclature

Acronyms

AOA Angle of Arrival

C/A Code Coarse Acquisition Code

DOP Dilution of Precision

GDOP Geometrical DOP

GPS Global Positioning System

HDOP Horizontal DOP

IMU Inertial Measurement Unit

KF Kalman filter

LAMBDA Least squares AMBIGUITY Decorrelation Adjustment

LOS line of sight

MEMS micro-electromechanical system

NCAT National Center for Asphalt Technology

NNSS Navy Navigation Satellite System

P(Y) Code Encrypted Precision Code

PDOP Position DOP

PRN Pseudo Random Noise

RPV Relative Position Vector

TDOP Time DOP

VDOP Vertical DOP

WGS84 World Geodetic System 1984

ZOH Zero-order Hold

Notation

ϕ_i^k Carrier phase measurement from GPS satellite k to receiver i

ρ_i^k Pseudorange from GPS satellite k to receiver i

$\begin{bmatrix} (e_x)_i^k & (e_y)_i^k & (e_z)_i^k \end{bmatrix}$ Unit vector from receiver i to satellite k in ECEF coordinates

$\begin{bmatrix} (x^k - x_i) & (y^k - y_i) & (z^k - z_i) \end{bmatrix}^T$ Vector from receiver i to satellite k in ECEF coordinates

$\begin{bmatrix} \delta x_{1,2} & \delta y_{1,2} & \delta z_{1,2} \end{bmatrix}$ Relative position vector between antennas 1 and 2

$\Delta\phi_{1,2}^k$ Single-differenced carrier phase measurement

$\Delta\rho_{1,2}^k$ Single-differenced pseudorange measurement

ϵ_i^k Error on pseudorange or carrier phase measurement from GPS satellite k to receiver i

λ_{L1} L1 carrier wavelength (19.03 cm)

$\nabla\Delta\phi_{1,2}^k$ Double-differenced carrier phase measurement

$\nabla\Delta\rho_{1,2}^k$ Double-differenced pseudorange measurement

ν_k Measurement noise

ϕ	Roll
ψ	Heading
θ	Pitch
A_k	Discrete-time state transition matrix
B_k	Discrete-time input matrix
b_x, b_y, b_z	Gyro bias estimates
C/N_0	Carrier-to-noise Ratio
cdt^k	Satellite error for satellite k
cdt_i	Receiver clock error for receiver i
$cdt_{1,2}$	Relative clock error between antennas 1 and 2
g_x, g_y, g_z	Gyro measurements
I_i^k	Ionospheric error from GPS satellite k to receiver i
L	Left null space of geometry matrix
N_i^k	Integer number of cycle ambiguities on carrier phase measurement from GPS satellite k to receiver i
P_k	State estimate variance-covariance matrix
Q_k	Discrete-time process disturbance variance-covariance matrix
r_i^k	True range from GPS satellite k to receiver i
R_k	Measurement noise variance-covariance matrix
T_i^k	Tropospheric error from GPS satellite k to receiver i

u_k Discrete-time input

w_k Process noise

E, N, U East, North, and Up local coordinates

Chapter 1

Introduction

Since the proposal of the Global Positioning System, or GPS, in 1973, it has become the gold standard in outdoor positioning. However, it was not the first satellite-based navigation system. The Navy Navigation Satellite System (NNSS), or Transit, employed satellites which constantly broadcast a 400 MHz tone. By analyzing the changing Doppler shift of the satellite and broadcast ephemeris data, a user at a known altitude could calculate the range to the satellite at closest approach. While Transit allowed a user to calculate position to within a few hundred meters, the system did have several limiting factors. The requirement of known altitude and a necessary velocity correction made transit nearly useless to aircraft. Also, the potential for mutual interference of satellites limited the constellation to about five satellites worldwide, leaving windows of satellite unavailability of 35 to 100 minutes. Despite the flaws of the Transit system, its implementation showed that a satellite-based navigation system was plausible and paved the way for future satellite navigation systems.

By 1972, the Navy was working on the Timation constellation, which featured satellites with high-precision clocks. The Timation satellites were primarily intended to provide accurate time information around the Earth, which GPS satellites are still used for today. The approval for GPS (the system was actually named NAVSTAR, the Global Positioning System) came in late 1973 with the first operational prototype satellite launched in early 1978. Four satellites were approved for the first phase of the project, with two more granted to solve the problem of tracking boosters from test launches of the Trident missile. It was the approval of the two extra GPS satellites, instead of upgrading the Transit constellation, that would allow GPS to eventually grow into the system it is today, with 32 satellites [14].

1.1 Motivation

One of the biggest barriers to accurate GPS navigation is error caused by multipath, which occurs when a GPS signal reflects off of an object. Multipath is discussed and illustrated in Section 2.3, but for now it will suffice to say the reflection can cause an incorrect positioning solution. Another problem in GPS navigation, generally unrelated to multipath error, is known as “spoofing.” Spoofing is the act of generating a counterfeit GPS signal with the express intent of influencing the positioning solution of a GPS receiver. While most civilians will probably never face a real threat from spoofing, it is a serious threat to military users, and is discussed further in Section 2.4. In both cases, it is likely the GPS signal will be arriving at the GPS antenna from a direction other than the location of the GPS satellite.

This thesis uses a GPS/INS attitude system to check the angle of arrival (AOA) of incoming GPS signals; that is, which direction they are arriving from relative to vehicle GPS antennas. With a known attitude solution for the vehicle system and known positions of the GPS satellites, the expected AOA may be computed for each satellite and checked against the perceived angle of arrival of the signals. If the expected and perceived AOA do not match up to within a preset threshold, the GPS signal in question can be rejected and its signal excluded from GPS calculations.

1.2 Previous Work

While there has been much research in the area of anti-spoofing for GPS, surprisingly few people have researched the use of AOA as a spoofing or multipath detection technique. However, AOA and multiple antenna arrays are more commonly used for anti-jamming.

In 1999, Brown et. al. used a controlled radiation pattern antenna (CRPA), an antenna with multiple internal elements, on aircraft to detect the AOA of jamming signals [2]. The CRPA systems were able to create a null in the direction of the jammer in a manner similar to a beamformer, mitigating the jamming affects on the antenna. The creation of the null

also allowed the aircraft to detect the direction of the jammer, provided the aircraft attitude was known accurately. Using the CRPAs, two or more aircraft working in a team could find ground-based jammer locations with good accuracy.

In 2002 McDowell proposed a system for the mitigation of multipath using a beamforming array and proposed another system in 2007 to mitigate GPS spoofers and jammers [10, 11]. In both cases, the system used a complex beamforming array to identify potential jamming and spoofing threats, then produce a synthetic null in the direction of the jammer or spoofer. The null is created by multiplying each of the antenna inputs by a complex weight, which can be configured to boost or lower signal power in a given direction.

In 2009, Montgomery et. al. [12, 13] proposed a system in which a pre-surveyed stationary pair of GPS antennas was used to observe differenced carrier phase measurements, comparing the actual measurements to expected measurements. The work showed that under normal circumstances, the differenced carrier phase measurements were predictable and changed slowly with respect to time. While the antenna pair could not be spoofed under normal conditions, as it is illegal to broadcast spoofed GPS signals outdoors, the receiver setup was tested indoors. Instead of using an actual spoofer, Montgomery et. al. collected GPS signals outside and sent the signals inside, along with a broadcast antenna. The setup is fairly representative of a GPS spoofer as all of the GPS signals are coming from a central location (source). The results showed that, excluding noise, the differenced carrier phase measurements were the same for the signals associated with each satellite and that they were constant with respect to time. In [13], it is mentioned that the concepts could be used in conjunction with an IMU for a non-stationary array of antennas, although Montgomery et. al. make no attempt at testing their setup dynamically.

In this thesis, an anti-spoofing and anti-multipath algorithm is designed. The algorithm uses an IMU to propagate the attitude state between GPS measurements, allowing the incoming GPS measurements to be checked for faulty signals by use of angle of arrival (AOA) or other related metrics.

1.3 Contributions

The following contributions are made in this thesis:

- An algorithm is developed to check incoming GPS measurements in order to detect multipath error or spoofed GPS signals.
- Several metrics, including AOA, are tested and compared in their ability to detect faulty GPS measurements
- Potential errors in the metrics are quantified as a function of AOA, baseline length, attitude error, and signal error

1.4 Thesis Outline

This chapter has served as an introduction to the GPS system and the problem at hand. In Chapter 2, the structure of the GPS signal is discussed, followed by the GPS measurements used in this work. Multipath error and spoofing are also explained, followed by potential spoofing scenarios and an overview of types of spoofing countermeasures.

Chapter 3 presents a differential GPS positioning technique to compute precise relative position vectors (RPVs) between pairs of antennas. Once the RPVs have been calculated, they are rotated from the standard GPS positioning coordinate frame into a local coordinate frame to be used in calculating vehicle attitude. In Chapter 4, the IMU used in the work is then characterized and a kinematic Kalman to estimate vehicle attitude using IMU and GPS measurements is proposed.

Finally, the multipath and spoofing rejection algorithm is discussed in Chapter 5, including the calculation and checking of AOA. Alternate metrics closely related to AOA are also explored. Results are then presented on the performance of the algorithm in Chapter 6, while conclusions and suggestions for future work are made in Chapter 7. In the Appendices, standalone GPS positioning and Kalman filtering are discussed. An overview of the signal rejection system in the form of a block diagram may be seen in Figure 1.1.

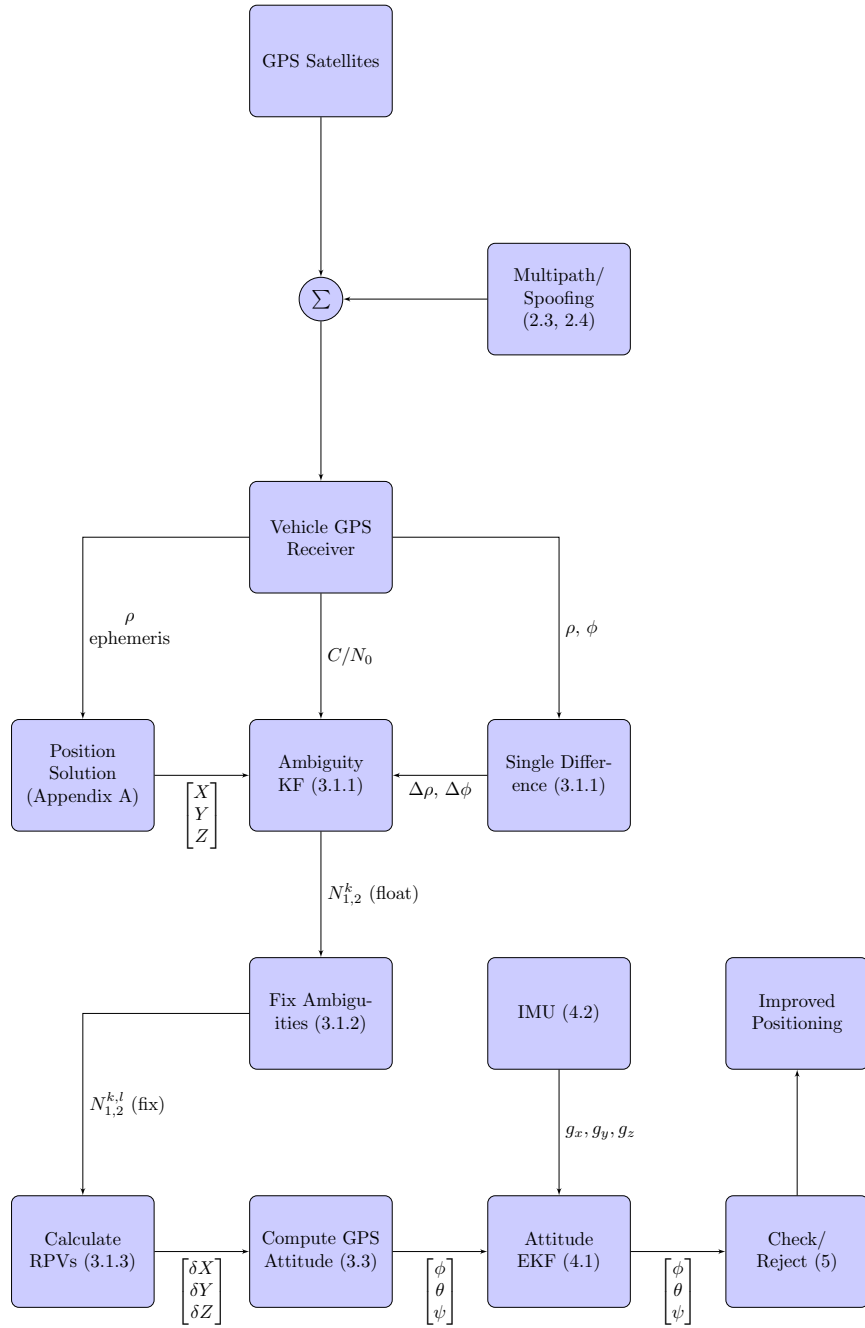


Figure 1.1: Block diagram overview of the signal rejection system. Text in parentheses denotes the appendix or chapter and section which contains information about each block.

Chapter 2
The Global Positioning System

2.1 Signal Structure

The GPS signal is made up of three signal components:

1. Carrier Wave
2. Navigation Data
3. Spreading Sequence

Refer to Figure 2.1 for a visual representation of the components of a GPS signal.

The first component, the carrier wave, is a radio wave responsible for carrying the other two signal components. There are two carrier frequencies, the L1 and L2 frequencies, in use on all GPS satellites. Additionally, an L5 signal is planned for future use. The L1, L2, and L5 frequencies are all scalar multiples of a common frequency, f_0 :

$$f_0 = 10.23 \text{ MHz} \tag{2.1}$$

$$f_{L1} = 154f_0 = 1575.42 \text{ MHz} \tag{2.2}$$

$$f_{L2} = 120f_0 = 1227.60 \text{ MHz} \tag{2.3}$$

$$f_{L5} = 115f_0 = 1176.54 \text{ MHz} \tag{2.4}$$

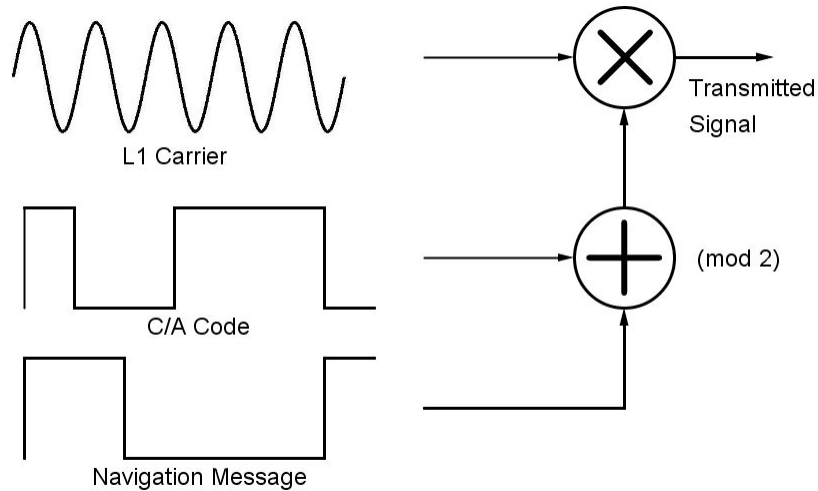


Figure 2.1: Components of a GPS signal

The second component, the navigation data, contains information on the orbit parameters and health of the GPS satellites as well as the transmission time of those parameters. It is transmitted at a rate of 50 bits per second. The third component, the spreading sequence, is made up of pseudo random noise (PRN) – a sequence of bits that appears to be random but is in fact generated by a known process. There are two such sequences for each satellite: a coarse acquisition (C/A) code and an encrypted precision code, the P(Y) code. Each satellite has a unique C/A code, sometimes referred to as a Gold code (after Robert Gold), associated with it. These Gold codes have correlation properties that allow all of the GPS satellites to transmit at the same frequency without interfering with each other. Specifically, a Gold code will have a high correlation value with a copy of the same code if the two are perfectly in phase, or low correlation if the two are out of phase or if it is cross-correlated with another Gold code. Those same correlation properties also allow receivers to track GPS signals with impressive carrier-to-noise ratios (C/N_0), even though GPS signals are *much weaker* than the ambient background noise.

2.2 Measurements

For the purpose of this work, there are three measurements (in addition to ephemeris data) of interest from each satellite: pseudorange measurements, carrier phase measurements, and carrier-to-noise ratio measurements. The pseudorange and carrier phase measurements are often compared to a tape measure or ruler from a GPS receiver to a GPS satellite, as both of them estimate the distance from the the receiver to the satellite. For the pseudorange, the ticks of the “ruler” are far apart, nearly 300 meters, corresponding to the distance between bit changes of the Gold codes. Conversely, the carrier phase measurements are based on the cycles of the L1 frequency, which repeats about every 19 centimeters. The cycle phase may generally be tracked to about 0.01 cycles, giving the measurement its precision. The downside to the carrier phase measurement is that it includes an ambiguous number of extra cycles between the user and satellite, so the measurement is useless on its own. Keeping with the ruler analogy, imagine a ruler from the receiver to a satellite, where the end of the ruler is an unknown distance before or past the satellite, but that distance is some multiple of 19 centimeters. Finally, the carrier-to-noise ratio is a measurement of GPS signal quality and is used in this work to calculate measurement covariance matrices for a Kalman filter, discussed in a later section.

2.2.1 Pseudorange Measurement

Pseudoranges may be modeled by the equation:

$$\rho_i^k = r_i^k + cdt_i + cdt^k + T_i^k + I_i^k + \epsilon_i^k \quad (2.5)$$

where i identifies a receiver, k identifies a satellite, r_i^k is the real range, or distance, from antenna i to satellite k , cdt_i and cdt^k are the receiver and satellite clock error (in units of meters) respectively, T_i^k and I_i^k are tropospheric and ionospheric errors respectively, and ϵ_i^k represents the remaining noise, modeled as zero mean Gaussian noise. The range is the

magnitude of the vector from the GPS receiver to the GPS satellite. The receiver clock error is caused by the fact that the clock of a receiver will not be perfectly synchronized with GPS time. It is shown in Appendix A that the receiver clock error may be solved at each instance of measurements, along with the GPS receiver position using measurements from at least four satellites. The satellite clock error is a very similar error and is caused by the inconsistency between a GPS satellite clock and true GPS time. Parameters for calculating an estimate of the clock error of each GPS satellite as a function of time are included in the ephemeris data. Tropospheric and ionospheric errors are present due to the signals passing through the atmosphere.

2.2.2 Carrier Phase Measurement

The carrier phase measurement model is similar to that of the pseudorange measurement:

$$\phi_i^k = r_i^k + cdt_i + cdt^k + T_i^k - I_i^k + \lambda_{L1}N_i^k + \epsilon_i^k \quad (2.6)$$

It should be noted that there are three differences between the pseudorange and carrier phase measurements. The first difference between the equations is the cycle ambiguity mentioned in Section 2.2, represented by the term N_i^k in Equation (2.6). The integer number of cycles is multiplied by the L1 carrier frequency wavelength ($\lambda_{L1} = 19.03$ cm) to express the distance in meters. The second difference is in the ϵ_i^k term, which still represents Gaussian noise, but with much lower variance than the Gaussian noise in Equation (2.5). Finally, the ionospheric error on the carrier phase measurement will have the same magnitude as the ionospheric error on the pseudorange, but will have the opposite effect. As mentioned in Section 2.2, carrier phase measurements are not useful alone; however, they may be used together with pseudorange measurements to provide a precise relative positioning solution between receivers, discussed in detail later.




Figure 2.2: Diagram of multipath error due to signal reflection from building.

2.3 Multipath Error

GPS receivers calculate a pseudorange to each satellite based on the time it takes for the signal to travel between the satellite and receiver. Using pseudoranges from multiple satellites, a position solution may be computed using a recursive least squares algorithm. If a GPS signal travels an indirect path to the GPS receiver, generally from being reflected off of a building or other obstacle, the signal takes a longer time to reach the receiver. Since the estimation of the range between a receiver and satellite is based on time, the extra travel time translates into a distance error in the range measurement. This error, known as multipath, is common in urban and foliage-rich environments and is the primary barrier to accurate positioning in those environments. An example of multipath is shown in Figure 2.2, where the receiver does not have a direct line of sight (LOS) to the GPS satellite supplying the signal. The path of the received GPS signal is shown in red.

2.4 Spoofing

In the GPS community, the term “spoofing” is used to describe the act of creating one or multiple counterfeit GPS signals with the intent of misleading the location of a GPS receiver. A GPS receiver that is successfully “spoofed” will accept the counterfeit GPS signal as an authentic one and use it in the positioning solution. The error of the resulting position solution will vary based on several factors, including the number of satellites in view, the geometry of those satellites, and the number of spoofed GPS signals accepted by the receiver. A spoofing attack that is well carried out could potentially fool the GPS receiver into accepting all of the spoofed GPS signals, allowing a skilled attacker to decide the perceived location of the GPS receiver.

While no commercially available spoofers exist, Humphreys et al. point out the hardware for a spoofer can be assembled from readily available inexpensive parts. The software for such a device could prove to be more complicated; however, the structure of the GPS signal is well known and the literature on software-defined receivers is plentiful [6].

2.4.1 Spoofing Scenarios

There are a plethora of potential GPS spoofing scenarios, many of them dangerous and all of them criminal. In [13], Montgomery et al. provide two such hypothetical scenarios in an article for *Inside GNSS* magazine. The first scenario involves spoofing a receiver used to synchronize timing between power grids, global trading networks, or communications networks. GPS receivers are well suited for this task, as a GPS receiver must have a very accurate estimate of time to function correctly. A successful spoofing attack on such a receiver could have far-reaching, disastrous consequences. Another hypothetical scenario would involve an operator of a fishing vessel who wishes to log a fictitious fishing trip in order to hide an illegal one. There are some places in the world where fishing in certain areas is restricted, so a fishing vessel might stand to gain a profit by fishing in one of these areas. Since many commercial fishing vessels are tracked by GPS in an attempt to prevent

these illegal voyages, a vessel could enter one of these restricted areas while spoofing its own GPS tracking system to avoid legal consequences [13]. Spoofing has many military implications too, especially in the case of military vehicles which are designed to operate autonomously. For example, Iran captured a CIA surveillance drone in December 2011. According to an interview with an Iranian engineer, the communications between the drone and remote operators were jammed, forcing the drone into an autonomous mode which relied on GPS for navigation. The drone was then spoofed and fooled into thinking it was landing at a US military base, only to land where it could be captured by Iranian forces [15]. While it is impossible to know if the drone was really spoofed or landed due to damage or a malfunction, the situation certainly serves as motivation for anti-spoofing measures.

2.5 Current Anti-Spoofing Measures

It is clear that spoofing poses a real threat to the integrity of GPS position solutions, so it is important to know measures may be taken to reduce or remove that threat. In 1995, Edwin L. Key published an internal memorandum for the MITRE corporation outlining countermeasures to detect GPS spoofing attempts. In the memorandum, Key recommended six techniques for detecting a spoofing attack [8]:

1. Amplitude Discrimination
2. Time-of-arrival Discrimination
3. Consistency of Navigation and Inertial Measurement Unit (IMU)
4. Polarization Discrimination
5. Angle-of-Discrimination
6. Cryptographic Authentication

Most anti-spoofing measures today still fall into one of these categories. This work is focused on the fifth category of techniques, which will be discussed in more detail in the rest of the

thesis. As the name suggests, the idea is to discriminate against, or reject, GPS signals which arrive from a direction other than that of the associated GPS satellite.

Most of the techniques identify spoofed GPS signals based on the characteristics of received signals. The first technique, amplitude discrimination, rejects GPS signals based on the power of the signals. Real GPS signals are very weak, as they are generated in space and must travel, at minimum, about 20 kilometers through the atmosphere; conversely, spoofed signals are generally much stronger. An uncharacteristically strong C/N_0 , or sudden change in C/N_0 , may be an indicator of spoofed signal being received.

Technique three takes a completely different approach to anti-spoofing by checking the results of the GPS solutions instead of the signals themselves. Although Key mentioned IMUs specifically, it may be expanded to include alternate navigation solutions, such as cell phone navigation. While this technique should be effective in theory, many alternate navigation solutions are less accurate than GPS, allowing at least a small window for a spoofing attack. An IMU in particular will drift over time, and so a high quality IMU would be necessary to effectively check for an altered GPS positioning solution.

Technique six takes yet another approach to the problem, essentially altering the structure of the GPS signal so that it can not be replicated. The P(Y) code, discussed in Section 2.1, is a perfect example, using a longer code in place of the shorter C/A code. Since the P(Y) code is known only to military users, would-be spoofers would be unable to replicate the L2 GPS signal. The author refers the reader to [7] and [19] for more detailed explanations of common anti-spoofing countermeasures.

Chapter 3

Attitude Determination

Knowing the attitude, or orientation, of the vehicle in a given coordinate frame is an important step in the ultimate goal of rejecting multipath or spoofed GPS signals. If the attitude of the vehicle is known, along with the positions of the GPS satellites, it is possible to calculate the angle of arrival of a GPS signal with respect to a pair of GPS antennas on the vehicle. To determine the full vehicle attitude (i.e. roll, pitch, and yaw) using GPS requires an array of three antennas mounted on a vehicle. Ideally, the array of antennas will line up with the coordinate axes of the vehicle. Differential GPS positioning techniques may then be used to precisely determine the RPV, allowing the attitude to be determined with simple geometry and coordinate rotations. The precision of the RPVs here is critical, as less precision in the RPVs will translate to less precision in the attitude estimates.

3.1 Differential Carrier Phase Positioning

Differential carrier phase positioning, also known as real-time kinematic (RTK) positioning, exploits the precision of carrier phase measurements to provide very accurate RPV estimates between two antennas. Before continuing with the derivation of the RTK calculations, it is helpful to rewrite Equations (2.5) and (2.6) by expanding both range terms:

$$\rho_i^k = \begin{bmatrix} (e_x)_i^k & (e_y)_i^k & (e_z)_i^k & 1 \end{bmatrix} \begin{bmatrix} x^k - x_i \\ y^k - y_i \\ z^k - z_i \\ cdt_i \end{bmatrix} + cdt^k + T_i^k + I_i^k + \epsilon_i^k \quad (3.1)$$

$$\phi_i^k = \begin{bmatrix} (e_x)_i^k & (e_y)_i^k & (e_z)_i^k & 1 \end{bmatrix} \begin{bmatrix} x^k - x_i \\ y^k - y_i \\ z^k - z_i \\ cdt_i \end{bmatrix} + cdt^k + T_i^k - I_i^k + \lambda_{L1} N_i^k + \epsilon_i^k \quad (3.2)$$

where $\begin{bmatrix} (e_x)_i^k & (e_y)_i^k & (e_z)_i^k \end{bmatrix}$ is the unit vector from receiver i to satellite k and $\begin{bmatrix} (x^k - x_i) & (y^k - y_i) & (z^k - z_i) \end{bmatrix}^T$ is the vector from the GPS receiver to the GPS satellite in ECEF coordinates. Before the next step, it is important to note two important facts about Equations (3.1) and (3.2). First, $\begin{bmatrix} (e_x)_i^k & (e_y)_i^k & (e_z)_i^k \end{bmatrix} \approx \begin{bmatrix} (e_x)_j^k & (e_y)_j^k & (e_z)_j^k \end{bmatrix}$ for two receivers that are close to each other. In fact, because the distance to a GPS satellite from either receiver is several orders of magnitude longer than the distance between two receivers (for this work), the two unit vectors may be considered equal. The second important fact is that the two atmospheric errors are strongly correlated by position; that is, they are nearly the same for two receivers that are close to each other. For receivers only a few meters apart, the errors will be nearly identical in the measurements of each receiver. Then, the pseudoranges and carrier phase measurements from receivers i and j for satellite k are subtracted, or differenced. The differenced measurements yield:

$$\rho_1^k - \rho_2^k = \Delta\rho_{1,2}^k = \begin{bmatrix} (e_x)_1^k & (e_y)_1^k & (e_z)_1^k & 1 \end{bmatrix} \begin{bmatrix} \delta x_{1,2} \\ \delta y_{1,2} \\ \delta z_{1,2} \\ cdt_{1,2} \end{bmatrix} + \epsilon_{1,2}^k \quad (3.3)$$

$$\phi_1^k - \phi_2^k = \Delta\phi_{1,2}^k = \begin{bmatrix} (e_x)_1^k & (e_y)_1^k & (e_z)_1^k & 1 \end{bmatrix} \begin{bmatrix} \delta x_{1,2} \\ \delta y_{1,2} \\ \delta z_{1,2} \\ cdt_{12} \end{bmatrix} + \lambda(N_{1,2}^k) + \epsilon_{1,2}^k \quad (3.4)$$

where $\begin{bmatrix} \delta x_{1,2} & \delta y_{1,2} & \delta z_{1,2} \end{bmatrix}$ is the RPV from antenna 1 to antenna 2, $cdt_{1,2} = cdt_1 - cdt_2$ and $\epsilon_{1,2}^k = \epsilon_1^k - \epsilon_2^k$. Equations (3.3) and (3.4) give the single-differenced pseudoranges and single-differenced carrier phase measurements, respectively. While many of the errors have been removed from the measurements, the relative clock bias still remains. The next step is to choose one satellite l - generally the satellite farthest from the horizon - and subtract the single-differenced pseudorange and carrier phase measurements associated with that satellite from the other single-differenced measurements, yielding the double-differenced pseudoranges and carrier phase measurements:

$$\nabla\Delta\rho_{1,2}^{k,l} = \begin{bmatrix} (e_x)_1^{k,l} & (e_y)_1^{k,l} & (e_z)_1^{k,l} \end{bmatrix} \begin{bmatrix} \delta x_{1,2} \\ \delta y_{1,2} \\ \delta z_{1,2} \end{bmatrix} + \epsilon_{1,2}^{k,l} \quad (3.5)$$

$$\nabla\Delta\phi_{1,2}^{k,l} = \begin{bmatrix} (e_x)_1^{k,l} & (e_y)_1^{k,l} & (e_z)_1^{k,l} \end{bmatrix} \begin{bmatrix} \delta x_{1,2} \\ \delta y_{1,2} \\ \delta z_{1,2} \end{bmatrix} + \lambda_{L1}(N_{1,2}^{k,l}) + \epsilon_{1,2}^{k,l} \quad (3.6)$$

where $\begin{bmatrix} (e_x)_1^{k,l} & (e_y)_1^{k,l} & (e_z)_1^{k,l} \end{bmatrix} = \begin{bmatrix} (e_x)_1^k & (e_y)_1^k & (e_z)_1^k \end{bmatrix} - \begin{bmatrix} (e_x)_1^l & (e_y)_1^l & (e_z)_1^l \end{bmatrix}$, $N_{1,2}^{k,l} = N_{1,2}^k - N_{1,2}^l$, and $\epsilon_{1,2}^{k,l} = \epsilon_{1,2}^k - \epsilon_{1,2}^l$. The double-differenced pseudorange and carrier phase measurements may be computed for $k = 1, \dots, n$ where n is the number of observed satellites and $k \neq l$. While the only error left in Equation (3.5) is the random error, its variance has now quadrupled from that of the ϵ_i^k term in Equation (2.5), which prevents the guaranteed precision desired for attitude determination. Equation (3.6) still maintains a high degree of precision, since the ϵ_i^k term in Equation (2.6) had a very small variance originally. However, the double-differenced integer ambiguities must be calculated before the double-differenced carrier phase measurements may be used to calculate an RPV. Calculating the double-differenced integer ambiguities is a two stage process. First, the ambiguities must be approximated, yielding non-integer (“float”) approximations of the ambiguities. The float

ambiguities are computed using a Kalman filter, which keeps a running approximation of the integer ambiguities. After each iteration of the Kalman filter, an attempt is made to “fix”, or round, the ambiguities to integers. Fixing the integers is attempted using the least-squares ambiguity decorrelation adjustment (LAMBDA) method, which uses the state estimate covariance matrix from the Kalman filter to round the integers “intelligently,” in order from least variance to most variance.

3.1.1 Calculating Float Ambiguities

As a first step, the known information must be put into a form that allows the ambiguities to be estimated. The singled-differenced pseudorange and carrier phase measurements may be represented as:

$$\begin{bmatrix} \Delta\rho_{1,2}^k \\ \vdots \\ \Delta\phi_{1,2}^k \\ \vdots \end{bmatrix} = \begin{bmatrix} (e_x)_1^k & (e_y)_1^k & (e_z)_1^k & 1 \\ \vdots & \vdots & \vdots & \vdots \\ (e_x)_1^k & (e_y)_1^k & (e_z)_1^k & 1 \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix} \begin{bmatrix} \delta x_{1,2} \\ \delta y_{1,2} \\ \delta z_{1,2} \\ cdt_{1,2} \end{bmatrix} + \lambda_{L1} \begin{bmatrix} 0 \\ \vdots \\ N_{1,2}^k \\ \vdots \end{bmatrix} + w \quad (3.7)$$

To estimate the ambiguities, it is necessary to have a model of the ambiguities that will fit the Kalman filter. The equations for implementing a Kalman filter and some background information on the filter may be found in Appendix B.

3.1.1.1 Time Update

For the time update step of the Kalman filter, it is necessary to have a model describing the state transition that fits the form of Equation (B.1). Once a GPS signal is being tracked, the carrier phase integer ambiguity should not change. Therefore, A_k is given by an $m \times m$ identity matrix, where m is the number of ambiguities:

$$A_k = I_{m \times m} \quad (3.8)$$

Also, since there is no input into the “system” of integer ambiguities, the $B_k u_k$ term of Equation (B.1) may be omitted completely. Strictly speaking, since there should be no variability in any of the integer ambiguities, Q_k should be an $m \times m$ matrix of zeroes; however, having a “perfect” model puts the filter at risk of going to sleep, as explained in Section B.3. To prevent the filter from going to sleep, Q_k may be artificially inflated to ensure that new measurements will always be taken into account. For this work, Q_k was chosen to be:

$$Q_k = I_{m \times m} \times 10^{-6} \quad (3.9)$$

It is important to note that because Q_k reflects process disturbance error that is not actually present in the system, P_k will not represent the real covariance of the integer ambiguity measurements. The state estimation error covariance matrix, P_k , is still useful as a relative comparison of the variance between the individual ambiguities.

3.1.1.2 Measurement Update

To make Equation (3.7) fit the form of Equation (B.2), it is necessary to remove the matrix of unit vectors, called the geometry matrix, and the RPV. It is possible to remove the geometry matrix term by multiplying Equation (3.7) by the left null space, L , of the geometry matrix; that is, a matrix L such that:

$$L \begin{bmatrix} (e_x)_1^k & (e_y)_1^k & (e_z)_1^k & 1 \\ \vdots & \vdots & \vdots & \vdots \\ (e_x)_1^k & (e_y)_1^k & (e_z)_1^k & 1 \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix} = 0 \quad (3.10)$$

Multiplying Equation (3.7) by L will yield:

$$L \begin{bmatrix} \Delta\rho_{1,2}^k \\ \vdots \\ \Delta\phi_{1,2}^k \\ \vdots \end{bmatrix} = L \begin{bmatrix} 0_{m \times m} \\ \lambda_{L1} I_{m \times m} \end{bmatrix} \begin{bmatrix} N_{1,2}^k \\ \vdots \end{bmatrix} + w \quad (3.11)$$

This fits the format of Equation (B.2), giving:

$$y_k = L \begin{bmatrix} \Delta\rho_{1,2}^k \\ \vdots \\ \Delta\phi_{1,2}^k \\ \vdots \end{bmatrix} \quad (3.12)$$

$$C_k = L \begin{bmatrix} 0_{m \times m} \\ \lambda_{L1} I_{m \times m} \end{bmatrix} \quad (3.13)$$

$$x_k = \begin{bmatrix} N_{1,2}^k \\ \vdots \end{bmatrix} \quad (3.14)$$

The computation of the measurement covariance matrix, R_k , is taken from [9]. R_k is given by:

$$R_k = L \begin{bmatrix} \sigma_{\rho,1}^2 + \sigma_{\rho,2}^2 & 0 & 0 & 0 \\ 0 & \ddots & 0 & 0 \\ 0 & 0 & \sigma_{\phi,1}^2 + \sigma_{\phi,2}^2 & 0 \\ 0 & 0 & 0 & \ddots \end{bmatrix} L^T \quad (3.15)$$

where σ_ρ and σ_ϕ are given by the Table 3.1 and the equations:

$$\sigma_\rho^2 = \sigma_{\rho atm}^2 + \sigma_{DLL}^2 \quad (3.16)$$

$$\sigma_{DLL} = \lambda_c \sqrt{\frac{4d^2 B_{n\rho}}{C/N_0} (2(1-d) + \frac{4d}{TC/N_0})} \quad (3.17)$$

$$\sigma_\phi^2 = \sigma_{\phi atm}^2 + \sigma_{PLL}^2 \quad (3.18)$$

$$\sigma_{PLL} = \frac{\lambda_{L1}}{2\pi} \sqrt{\frac{B_{n\phi}}{C/N_0} (1 + \frac{1}{TC/N_0})} \quad (3.19)$$

It is important to note that the C/N_0 measurement is generally reported in decibels, but should be converted to the absolute C/N_0 before use in Equations (3.17) and (3.19)

Table 3.1: Parameters for use in Equations (3.16) through (3.19)

Parameter	Description	Value (units)
$\sigma_{\rho atm}^2$	Variance due to atmospheric code delay	5.22 (m)
λ_c	Code chip width	293.05 (m)
d	Correlator spacing	0.5 (chips)
$B_{n\rho}$	Code loop noise bandwidth	2 (Hz)
T	Prediction integration time	2 (ms)
$\sigma_{\phi atm}^2$	Variance due to atmospheric carrier delay	0.03 (m)
λ_{L1}	Carrier wavelength	19.03 (cm)
$B_{n\phi}$	Carrier loop noise bandwidth	18 (Hz)

3.1.1.3 Implementation

To initialize the filter, an estimate of the the integer ambiguities must first be calculated. It is observed that subtracting the single-differenced pseudorange from the single-differenced carrier phase provides a good candidate:

$$\Delta\phi_{1,2}^k - \Delta\rho_{1,2}^k = \vec{e} \begin{bmatrix} \delta x_{1,2} \\ \delta y_{1,2} \\ \delta z_{1,2} \\ cdt_{1,2} \end{bmatrix} + \lambda(N_{1,2}^k) - \vec{e} \begin{bmatrix} \delta x_{1,2} \\ \delta y_{1,2} \\ \delta z_{1,2} \\ cdt_{1,2} \end{bmatrix} = \lambda_{L1}(N_{1,2}^k) \quad (3.20)$$

where $e_1^k = \begin{bmatrix} (e_x)_1^k & (e_y)_1^k & (e_z)_1^k & 1 \end{bmatrix}$. Note that the error terms have been omitted for clarity. Equation (3.20) may be simplified to yield the initial estimate of $N_{1,2}^k$:

$$N_{1,2}^k = \frac{\Delta\phi_{1,2}^k - \Delta\rho_{1,2}^k}{\lambda_{L1}} \quad (3.21)$$

The initial state covariance matrix is somewhat arbitrary based on desired filter performance, specifically how quickly it is desired that P_k should converge. From [9], it was decided that the state covariance matrix would be initialized as:

$$P_0 = \frac{1}{2} \times I_{m \times m} \quad (3.22)$$

After initialization, updating the filter is fairly straightforward using the steps outlined for a Kalman filter.

3.1.2 LAMBDA method

Once float estimates of the ambiguities have been calculated, the next step is attempting to get a set of fixed ambiguities. The most straightforward method of obtaining fixed ambiguities would be to merely round the float ambiguities; however, simple rounding does not account for the information contained in the state covariance matrix associated with the float ambiguities. A popular method that does include the state covariance matrix is the LAMBDA method[?]. As the name implies, the LAMBDA method uses a decorrelation transformation to decorrelate the error in the ambiguities so that they may be selected more carefully. More specifically, the transformation centers around a transformation matrix Z such that

$$z = Z^* a \quad (3.23)$$

$$Q_z = Z^* Q_a Z \quad (3.24)$$

where a represents the float cycle ambiguities, z represents the transformed float ambiguities, Q_a represents the state covariance matrix of the true cycle ambiguities, Q_z represents the covariance matrix of the transformed ambiguities, and Z^* denotes the conjugate transpose of Z . Ideally, the transformed ambiguities will be completely uncorrelated, making Q_z a diagonal matrix. To obtain Z , Q_a must first be represented by the equation:

$$Q_a = L^{-*} D^{-1} L^{-1} \quad (3.25)$$

An alternative representation is to take the inverse of both sides of the equation, giving:

$$Q_a^{-1} = L D L^* \quad (3.26)$$

where L is a lower triangular matrix with ones along the main diagonal, and D is a diagonal matrix. Z will be a matrix of integers and is intended to approximate L ; if $Z = L$:

$$Q_z = Z^* Q_a Z = Z^* L^{-*} D^{-1} L^{-1} Z = D^{-1} \quad (3.27)$$

which gives the intended result of Q_z being diagonal. The integer restriction on L means that usually $Z \neq L$ and the transformed ambiguities will not be perfectly uncorrelated; however, the new ambiguities will be much less correlated than the original ambiguities. For details on the implementation of the decorrelation algorithm, including how to calculate the matrices involved, the author refers the reader to [4].

Once the ambiguities are decorrelated, a search must be performed to determine the most probable fixed integer ambiguities. The most probable integer set will minimize the function:

$$(\hat{z} - z) Q_z^{-1} (\hat{z} - z) \leq \chi^2 \quad (3.28)$$

where \hat{z} is the vector of estimated transformed ambiguities, z is a vector of integers to be tested, $Q_{\hat{z}}^{-1}$ is the covariance matrix for \hat{z} , and χ^2 is a quantity used to control the volume of the search space for the integers. If $Q_{\hat{z}}^{-1}$ were diagonal (ambiguities perfectly decorrelated), the function would essentially be:

$$\sum_{i=1}^n \frac{(\hat{z}_i - z_i)^2}{q_{i,i}} < \chi^2 \quad (3.29)$$

where $q_{i,i}$ is the element in the i^{th} row and i^{th} column of $Q_{\hat{z}}$. Equation (3.29) shows the essence of the idea behind the integer search: error is weighted more heavily for more confident measurements, or the confident measurements are given “preference.”

There are several methods for choosing χ^2 , depending on how thorough of a search is desired. One method is to round all ambiguities except one to the nearest integer and the last ambiguity to the second nearest integer. The norm for this integer set (the left hand side of Equation (3.28)) may then be computed and set to χ^2 , ensuring that at least one integer set and probably a few of them will be contained within the bounds of χ^2 . Another technique to reduce the search space after an initial χ^2 has been selected is to reduce χ^2 to the norm of the first integer set found. Since the norm of an integer set is essentially a measurement of error between the fixed and float integers, there is no need to look for integer sets with a higher norm than the first set of integers [4].

Once a set of integers has been chosen, they must be transformed back into integer ambiguities for the differenced carrier phase equations. Recall that:

$$z = Z^* a \quad (3.30)$$

So the fixed integer set a may be calculated by multiplying z on the left side by the inverse of Z^* .

3.1.3 Calculating RPVs

Once the double-differenced integer ambiguities have been obtained through the LAMBDA method, the RPVs may be calculated. Rearranging Equation (3.6), it is seen that:

$$\nabla\Delta\phi_{1,2}^{k,l} - \lambda_{L1}(N_{1,2}^{k,l}) = \begin{bmatrix} (e_x)_1^{k,l} & (e_y)_1^{k,l} & (e_z)_1^{k,l} \end{bmatrix} \begin{bmatrix} \delta x_{1,2} \\ \delta y_{1,2} \\ \delta z_{1,2} \end{bmatrix} + \epsilon_{1,2}^{k,l} \quad (3.31)$$

Since the integer ambiguities are now known, calculating the relative positioning vector is given by a least squares solution:

$$(H^T H)^{-1} H^T (\nabla\Delta\phi_{1,2}^{k,l} - \lambda_{L1}(N_{1,2}^{k,l})) = \begin{bmatrix} \delta x_{1,2} \\ \delta y_{1,2} \\ \delta z_{1,2} \end{bmatrix} \quad (3.32)$$

$$H = \begin{bmatrix} (e_x)_1^{k,l} & (e_y)_1^{k,l} & (e_z)_1^{k,l} \\ \vdots & \vdots & \vdots \end{bmatrix} \quad (3.33)$$

3.1.4 Single-differenced cycle ambiguities

The most ideal situation for the task at hand would be to know the relative clock error between the two receivers and the integer cycle ambiguities for the single-differenced carrier phase measurements. In addition to a small increase in RPV precision, knowing the single-differenced cycle ambiguities makes the RPV calculation process more robust, as the ambiguities are all independent of one another and not connected to any one satellite.

Calculating the single-differenced ambiguities requires knowledge of the relative clock error between the two receivers, so the clock bias must be calculated first. It is important to note that for two separate GPS receivers with two separate clocks, the clock bias will vary with time and will be random, thus eliminating the possibility of calculating the single-differenced cycle ambiguities. However, it is possible to have two or more GPS receivers that

share a common oscillator. In the case of a shared oscillator, the receivers will experience a line bias instead of a clock bias. The manner in which a line bias affects the signals is the same as a clock bias, with the exception that the line bias is constant. To calculate an estimate of the clock error, a high precision RPV must first be calculated between the antennas of the two receivers. Once the high-precision RPV has been calculated, Equation (3.3) may be utilized to calculate a noisy estimate of the clock error. Rearranging Equation (3.3) to give the noisy clock error estimate yields:

$$\Delta\rho_{1,2}^k - \begin{bmatrix} (e_x)_1^k & (e_y)_1^k & (e_z)_1^k \end{bmatrix} \begin{bmatrix} \delta x_{1,2} \\ \delta y_{1,2} \\ \delta z_{1,2} \end{bmatrix} = cdt_{1,2} + e_{1,2}^k \quad (3.34)$$

As seen in Equation (3.34), a measurement of the clock error may be obtained by calculating the dot product of the unit vector to a satellite and the RPV between the receivers, then subtracting that dot product from the pseudorange associated with that satellite. The error on single-differenced pseudoranges is fairly substantial, so it is necessary to base the relative clock error estimate on a large number of measurements. The author saw good results by using an average of 1000-1500 measurements, although a weighted least squares or filter approach may be used – potentially relying on SNR or satellite elevation as an indicator of the quality of the measurement. Since there are often half a dozen or more satellites in view, especially in an open environment, it is possible to have a good estimate of the clock error using a couple hundred calculations of the RPV between the two receiver antennas.

Once the relative clock error has been determined, a similar calculation is carried out on the single-differenced carrier phase measurements to give the single-differenced ambiguities. Rearranging Equation (3.4) will give the single-differenced ambiguities:

$$\Delta\phi_{1,2}^k - \begin{bmatrix} (e_x)_1^i & (e_y)_1^i & (e_z)_1^i & 1 \end{bmatrix} \begin{bmatrix} \delta x_{1,2} \\ \delta y_{1,2} \\ \delta z_{1,2} \\ cdt_{1,2} \end{bmatrix} = \lambda(N_{1,2}^k) + e_{1,2}^k \quad (3.35)$$

Naturally there will be some error in the relative clock bias and RPV for this measurement, but as long as the combined errors do not exceed half of an L1 carrier wavelength, or about 9.5 cm, the ambiguities will be rounded to the correct value. Since it is known that the ambiguities must be integers, and the user now knows those integers, the small error of the pseudorange measurements may be used to correct the clock error. The clock error may be calculated by rearranging Equation (3.4) another way:

$$\phi_{1,2}^k - \begin{bmatrix} (e_x)_1^k & (e_y)_1^k & (e_z)_1^k \end{bmatrix} \begin{bmatrix} \delta x_{1,2} \\ \delta y_{1,2} \\ \delta z_{1,2} \end{bmatrix} - \lambda(N_{1,2}^k) = cdt_{1,2} + e_{1,2}^k \quad (3.36)$$

3.2 Rotating from ECEF to ENU

Before the attitude of the vehicle is calculated, the ECEF coordinates of the RPVs should be rotated into a local ENU frame, as vehicle attitude with respect to the ECEF frame has little intuitive meaning. To rotate ECEF coordinates into an ENU frame, two angles are needed – namely the latitude and longitude of the location of the ENU frame. If the Earth were a perfect sphere, calculating the latitude and longitude of any point on or above the surface of the sphere would be trivial. Latitude (ϕ) and longitude (λ) would be given by the equations:

$$\tan \lambda = \frac{Y_{ECEF}}{X_{ECEF}} \quad (3.37)$$

$$\tan \phi = \frac{Z_{ECEF}}{\sqrt{X_{ECEF}^2 + Y_{ECEF}^2}} \quad (3.38)$$

However, the Earth is not a sphere, but an ellipsoid, and is elongated outward in the X_{ECEF} and Y_{ECEF} directions. This elongation makes Equation (3.38) incorrect. In the field of GPS, the most commonly used model of the Earth is the World Geodetic System 1984 (WGS84). The following parameters may be found in [18] to find the latitude for an ECEF location:

$$a_e = 6,378,137 \text{ m} \quad (3.39)$$

$$b_e = 6,356,752 \text{ m} \quad (3.40)$$

$$e_p = .00335281 \quad (3.41)$$

$$L_{i+1} = L_c + e_p \sin(2L_i) \quad (3.42)$$

$$L_c = \text{atan}\left(\frac{Z_{ECEF}}{\sqrt{X_{ECEF}^2 + Y_{ECEF}^2}}\right) \quad (3.43)$$

It should be noted that Equation (3.42) is an iterative approach and should be carried out until the change between iterations is sufficiently small. It is recommended to initialize L_0 as L_c [18].

Once the latitude and longitude have been determined, converting from ECEF coordinates to ENU coordinates is a matter of simple rotations. In this thesis, the author has chosen to first rotate the ECEF coordinate frame about the Z-axis to line up the Y-axis with




Figure 3.1: ECEF and ENU coordinates

the intended East axis, then rotate about the rotated Y-axis to line up the rotated Z- and X-axes with the North- and Up-axes respectively, as shown in Figure 3.1. This rotation will work out well for this figure as the latitude and longitude may be used directly in this case without recalculating any angles. The coordinate rotations are shown by the equation:

$$\begin{bmatrix} U \\ E \\ N \end{bmatrix} = \begin{bmatrix} \cos \phi & 0 & \sin \phi \\ 0 & 1 & 0 \\ -\sin \phi & 0 & \cos \phi \end{bmatrix} \begin{bmatrix} \cos \lambda & \sin \lambda & 0 \\ -\sin \lambda & \cos \lambda & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} X_{ECEF} \\ Y_{ECEF} \\ Z_{ECEF} \end{bmatrix} \quad (3.44)$$

The coordinates in Equation (3.44) are out of order from the preferred East-North-Up order, but the order is easily rectified in an algorithm by reassigning the variables.

3.3 Calculating Attitude

Once the RPVs between the antennas have been determined, calculating the vehicle attitude becomes a matter of a few coordinate rotations. Figure 3.2 shows the configuration

of the antennas with respect to the vehicle coordinate frame, recreated from the manual for the receiver[1]. The X-axis is a horizontal axis and extends in the positive direction out of the passenger side of the vehicle. The Y-axis is a horizontal axis as well and extends in the positive direction out of the front of the car. Finally, the Z-axis is the vertical axis and extends in the positive direction out of the top of the car. As is shown in Figure 3.2, the antennas are situated such that the main antenna and first auxiliary antenna form a line along the Y-axis, while the main antenna and second auxiliary antenna form a line along the X-axis. It is important to note that this is not the only antenna configuration that could be used to determine vehicle attitude, but was chosen to make calculating the attitude easier. The antenna array could be rotated into some other position relative to the vehicle coordinate frame, as long as the relationship between the antenna configuration and vehicle axes is known. Additionally, the antennas do not need to be separated by 90 degrees. As the third attitude angle to be calculated (roll in this case) depends on the portion of the second baseline that is orthogonal to the first baseline, the 90 degree baseline will yield the best attitude estimate for a given baseline length compared to a larger or smaller angle.

Once the RPVs are in an ENU frame, the attitude of the vehicle may be calculated. As shown in Figure 3.3, heading is calculated first, followed by pitch, and roll is calculated last. Heading (ψ) may be calculated as:

$$\psi = \tan^{-1} \frac{E_{M \rightarrow A_1}}{N_{M \rightarrow A_1}} \quad (3.45)$$

where the $E_{M \rightarrow A_1}$ is the east component of the RPV from the main antenna to the auxiliary-1 antenna and $N_{M \rightarrow A_1}$ is the north component of the same RPV. In the following equations, $U_{M \rightarrow A_1}$ will be the up component of the RPV from the main antenna to the auxiliary-1 antenna and the subscript $M \rightarrow A_2$ will denote an RPV from the main antenna to the auxiliary-2 antenna. Once heading is computed, the ENU coordinates can be rotated by the negative heading angle about the U-axis, allowing pitch (θ) to then be calculated. Note

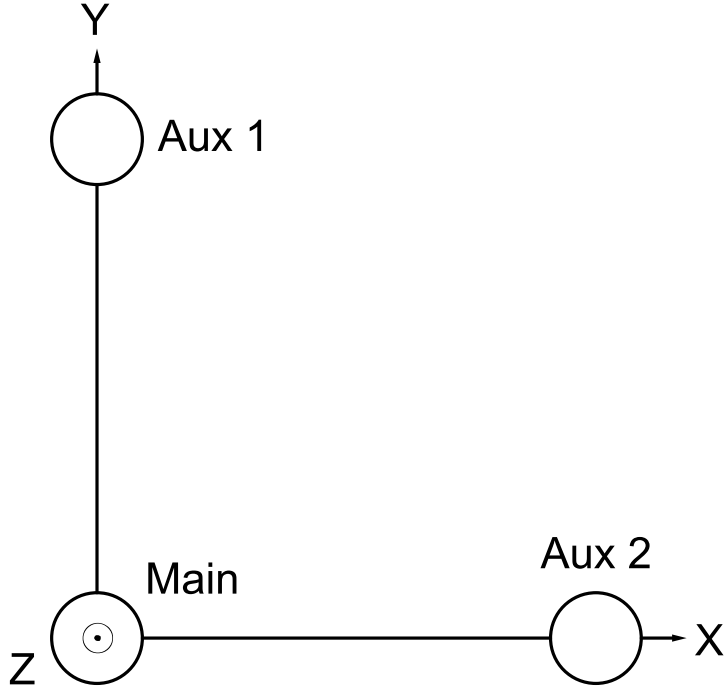


Figure 3.2: Antenna configuration shown in vehicle coordinate frame. View is top-down onto the vehicle. Figure recreated from figure in Septentrio PolarEx manual.

that the negative rotation is due to the fact that heading is calculated as a negative rotation about the the U-axis:

$$\begin{bmatrix} E'_{M \rightarrow A_1} \\ N'_{M \rightarrow A_1} \\ U'_{M \rightarrow A_1} \end{bmatrix} = \begin{bmatrix} \cos(-\psi) & \sin(-\psi) & 0 \\ -\sin(-\psi) & \cos(-\psi) & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} E_{M \rightarrow A_1} \\ N_{M \rightarrow A_1} \\ U_{M \rightarrow A_1} \end{bmatrix} \quad (3.46)$$

$$\theta = \tan^{-1} \frac{U'_{M \rightarrow A_1}}{N'_{M \rightarrow A_1}} \quad (3.47)$$

where $E'_{M \rightarrow A_1}$ denotes the east coordinate of the RPV once it has been rotated about the U-axis. There is no specific reason to make the heading a negative rotation, except that the GPS receiver used in this work derives it as a negative rotation. Keeping the derivation of the attitude measurements consistent is helpful in comparing them to the attitude measurements

of the receiver. Finally, the ENU coordinates of the second RPV will need to be rotated with the heading rotation and the pitch rotation so that roll (ϕ) can be calculated:

$$\begin{bmatrix} E''_{M \rightarrow A_2} \\ N''_{M \rightarrow A_2} \\ U''_{M \rightarrow A_2} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\theta) & \sin(\theta) \\ 0 & -\sin(\theta) & \cos(\theta) \end{bmatrix} \begin{bmatrix} \cos(\psi) & \sin(\psi) & 0 \\ -\sin(\psi) & \cos(\psi) & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} E_{M \rightarrow A_2} \\ N_{M \rightarrow A_2} \\ U_{M \rightarrow A_2} \end{bmatrix} \quad (3.48)$$

$$\phi = \tan^{-1} \frac{-U''_{M \rightarrow A_2}}{E''_{M \rightarrow A_2}} \quad (3.49)$$

where $E''_{M \rightarrow A_2}$ denotes the east coordinate of the RPV once it has been rotated about the U-axis and E-axis. If both RPVs have been rotated by all three attitude angles, the vehicle X-axis and Y-axis will be in line with the E-axis and N-axis, respectively.

There are other methods of calculating attitude, such as using a non-linear least squares solution or solving Wahba's problem. In both cases, a cost function is defined in which two sets of identical vectors in different frames are compared with a rotation matrix. The goal is to use the rotation matrix to minimize the cost function between the two sets of vectors. The primary reason the author does not use these techniques is that both of them yield a single rotation matrix, effectively the roll, pitch, and yaw rotation matrices combined. Since the author has an interest in knowing these attitude states individually, he has chosen to compute them separately instead of decomposing a single rotation matrix into three. The author refers the reader to [3] for more information on these alternative methods of computing attitude.




Figure 3.3: Euler angle sequence (recreated from Septentrio manual) [1]

Chapter 4

GPS/INS Integration

A large part of this work depends on having an accurate attitude estimate with which to check new GPS measurements. A more accurate attitude estimate will allow tighter tolerances on how much error is allowed in a GPS measurement before it may be rejected. However, a less accurate attitude estimate will mandate a larger window of GPS measurements that must be accepted, also allowing more room for error. In this chapter, a Kalman filter (KF) is designed to estimate the vehicle attitude using GPS and IMU measurements. Then, the error sources of the IMU used in this work are modeled and quantified.

4.1 GPS/INS Kalman Filter

To combine the GPS and IMU measurements, a kinematic KF was chosen. A kinematic KF differs from a regular Kalman filter in that it uses a kinematic model (i.e. the kinematic relationships between dynamic states), rather than a dynamic one, to model the system. Specifically, this Kalman filter will use gyro measurements as “inputs” to the system in order to advance the roll, pitch, and yaw estimates, while GPS attitude measurements will be the measurement updates for the filter. Since the states being estimated are non-linear, an extended KF (EKF) is used in place of the regular KF. The primary differences between the EKF and a standard KF are:

1. Since the states are non-linear, the state transition model is linearized about an operating point.
2. The state estimate covariance matrix is updated a little differently than with a standard KF.

To calculate the state space matrices needed for the EKF, the differential equations describing the change with respect to time are needed. They are:

$$\dot{\phi} = (g_x - b_x) \sin \phi \sin \theta + (g_y - b_y) \cos \theta - (g_z - b_z) \cos \phi \sin \theta \quad (4.1)$$

$$\dot{\theta} = (g_x - b_x) \cos \phi + (g_z - b_z) \sin \phi \quad (4.2)$$

$$\dot{\psi} = (g_x - b_x) \sin \phi \cos \theta - (g_y - b_y) \sin \theta - (g_z - b_z) \cos \theta \sin \phi \quad (4.3)$$

where $\dot{\phi}$, $\dot{\theta}$, and $\dot{\psi}$ are the changes in the attitude states with respect to time, g_x , g_y , and g_z are gyro measurements, and b_x , b_y , and b_z are gyro biases. The three attitude angles and three bias states will comprise the six states of the filter. The model is given by:

$$\dot{x} = Ax + B_u u + B_w w \quad (4.4)$$

$$\dot{x} = \begin{bmatrix} \dot{\phi} & \dot{\theta} & \dot{\psi} & \dot{b}_x & \dot{b}_y & \dot{b}_z \end{bmatrix}^T \quad (4.5)$$

$$x = \begin{bmatrix} \phi & \theta & \psi & b_x & b_y & b_z \end{bmatrix} \quad (4.6)$$

$$A = \begin{bmatrix} A_{1,1} & A_{1,2} \\ A_{2,1} & A_{2,2} \end{bmatrix} \quad (4.7)$$

$$A_{1,2} = \begin{bmatrix} -\sin \phi \sin \theta & -\cos \theta & \cos \phi \sin \theta \\ -\cos \phi & 0 & -\sin \phi \\ -\sin \phi \cos \theta & \sin \theta & \cos \phi \cos \theta \end{bmatrix} \quad (4.8)$$

$$A_{1,1} = A_{2,1} = A_{2,2} = 0_{3 \times 3} \quad (4.9)$$

$$B_u = \begin{bmatrix} -A_{1,2} \\ 0_{3 \times 3} \end{bmatrix} \quad (4.10)$$

$$u = \begin{bmatrix} g_x & g_y & g_z \end{bmatrix}^T \quad (4.11)$$

$$B_w = \begin{bmatrix} -A_{1,2} & 0_{3 \times 3} \\ 0_{3 \times 3} & I_{3 \times 3} \end{bmatrix} \quad (4.12)$$

$$y = Cx + \nu \quad (4.13)$$

$$C = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix} \quad (4.14)$$

where w is a 6-by-1 process noise vector and ν is a 3-by-1 measurement noise vector. Since the noise characteristics are estimated from sensor data, they will be defined with the discrete-time model. The model is discretized using a zero-order hold (ZOH) model with a sampling rate of 80 Hz (the measurement rate of the IMU, discussed later), giving the discrete-time state space model:

$$x_{k+1} = A_d x_k + B_{ud} u_k + B_{wd} w_k \quad (4.15)$$

$$A_d = I_{6 \times 6} + \frac{1}{80} A \quad (4.16)$$

$$B_{ud} = \frac{1}{80}B_u \quad (4.17)$$

$$B_{wd} = \frac{1}{80}B_w \quad (4.18)$$

$$y_k = Cx_k + \nu_k \quad (4.19)$$

$$w_k \sim N\left(0, \begin{bmatrix} (4 \times 10^{-5})I_{3 \times 3} & 0_{3 \times 3} \\ 0_{3 \times 3} & (10^{-8})I_{3 \times 3} \end{bmatrix}\right) \quad (4.20)$$

$$\nu_k \sim N\left(0, (2 \times 10^{-6})I_{3 \times 3}\right) \quad (4.21)$$

The values for the gyro noise and GPS attitude noise were taken from gyro and GPS measurements, while the uncertainty associated with the bias terms is chosen for desired filter performance. Detailed information about the EKF and the differences between an EKF and a KF may be found in [16].

4.2 IMU Errors

The IMU used in this work is the Crossbow IMU400-CC series. It is a 6-DOF micro-electromechanical system (MEMS) IMU, giving accelerometer and rate gyro measurements along each of three Cartesian coordinate axes. Since the scope of this work does not require the use of the accelerometer measurements, only the gyro measurements will be discussed. Using a simple sensor model found in [5], gyro output may be modeled by the equation:

$$g_r = r + c_r + b_r + w_{gyro} \quad (4.22)$$

where g_r is the gyro output, r is the true rotation rate of the gyro, c_r is a constant bias, b_r is a moving or walking bias, and w_{gyro} is random sensor noise, assumed to be normally distributed with zero mean. It is desirable to remove as much error as possible, as integrating rotation rate bias over time will produce error in the attitude estimates. Figure 4.1 shows raw gyro measurements from a stationary data set.

4.2.1 Bias Error

The specification sheet for the IMU400-CC series shows that each gyro measurement may experience a bias of up to ± 1 deg/s. As shown in Equation (4.22), the measurement bias has a constant component and a drifting component. Over a long static data set, the constant bias may be estimated by taking the mean of the measurements throughout the data set. For the data set in Figure 4.1, the largest calculated mean is just under .004 rad/s, or about .23 deg/s, well within the range given in the specification sheet for the sensor.

To observe a walking bias visually in the data, it is necessary to filter the data to remove the higher frequency noise. Figure 4.2 shows the data from the Y gyro in Figure 4.1 after low-pass filtering at several bandwidths. The first panel of Figure 4.2 starts to show the movement of the bias, although panels two and three are progressively more helpful in observing the bias over time. It is evident from the data that there is a detectable walking bias, although it seems to be bounded to within a range of about .001 rad/s, or .057 deg/s. Filtered data from the X and Z data sets showed walking biases with similar bounds. Over half a second, a bias of .057 deg/s has the potential to cause about .03 degrees of error in angular position.

4.2.2 Random Error

The random error portion of the gyro measurement may be modeled as a zero mean Gaussian noise. The standard deviation of the normal distribution is normally given in a specification sheet as a function of sampling rate as the “random walk” parameter. For

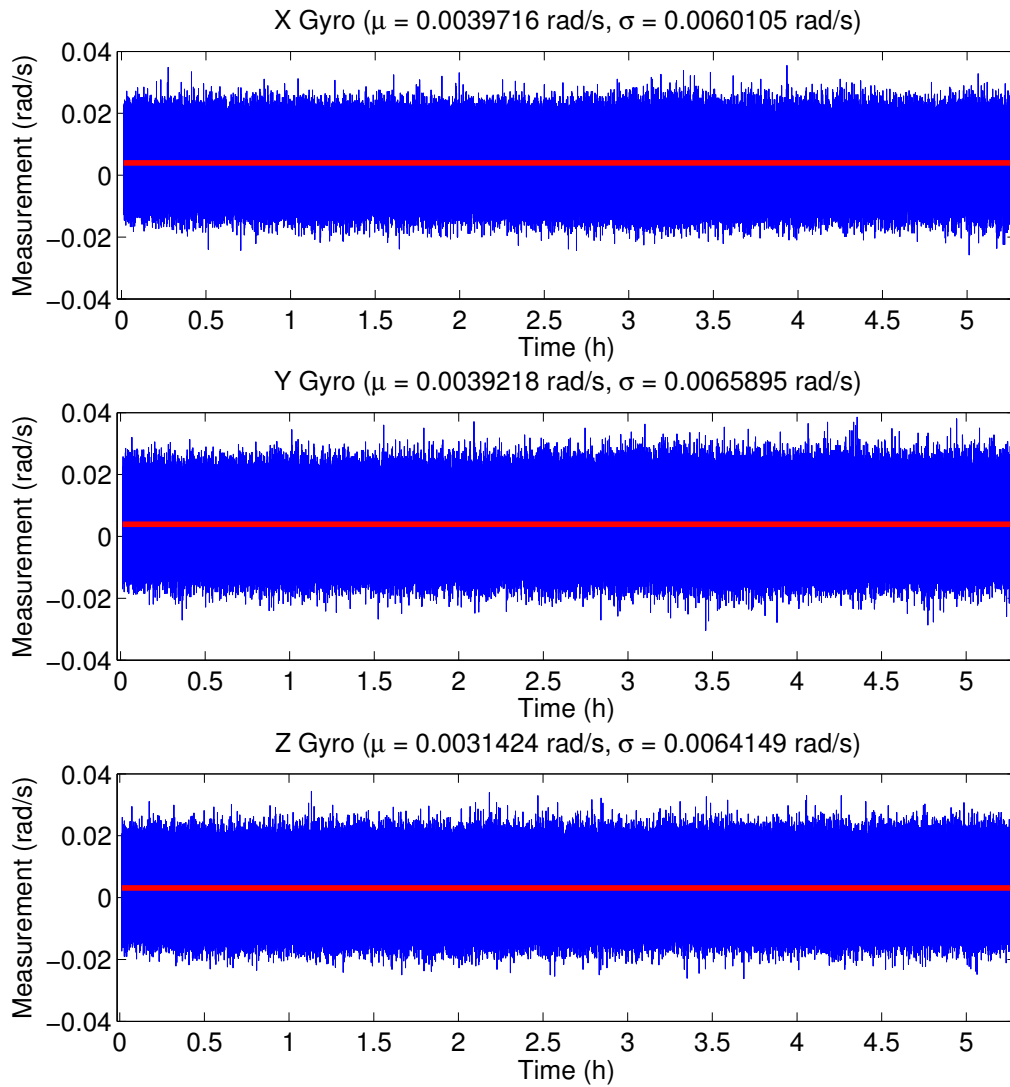


Figure 4.1: Unfiltered stationary gyro measurements over 3 axes. Red lines denote mean measurement.




Figure 4.2: Gyro data filtered at several bandwidths by low pass filter.

the Crossbow IMU400CC-200, the parameter is given as $< 4.5deg/\sqrt{hr}$. Computing the parameter in more useful units gives:

$$4.5 \frac{deg}{\sqrt{hr}} = .075 \frac{deg}{\sqrt{s}} = .075 \frac{\frac{deg}{s}}{\sqrt{hz}} \quad (4.23)$$

Multiplying the random walk parameter by the square root of the sampling frequency will give an upper bound on what to expect for the standard deviation of the random noise, giving:

$$\sigma = .075 \times \sqrt{80} = .671 \frac{deg}{s} = .0117 \frac{rad}{s} \quad (4.24)$$

In the data sets shown in Figure 4.1, the standard deviation of the sampled data sets is around the range of $\sigma = .006 - .007$ rad/s, so the experimental data is within the bounds expected from information in the specification sheet. Since the random error is completely non-deterministic, and thus cannot be compensated for, it is helpful to know the worst case error from integrating the random error. Monte Carlo simulations are a useful tool in such a situation, as it is simple to simulate a random walk of angular position by numerically integrating the generated random noise. The standard deviation calculated for Equation (4.24) was used in the Monte Carlo simulation. The results of the Monte Carlo simulations are shown in Figure 4.3, along with an envelope representing three standard deviations of the data at each time step. The envelope reaches approximately $\pm .16$ degrees at .5 seconds, which is the sampling period of the GPS receiver used in this work. From the 3σ envelope, it can be determined that about 99.7% of the random walks will result in attitude errors around or less than .16 degrees, while the largest simulated random walk was still less than .25 degrees. Overall, the Monte Carlo simulations show an acceptable amount of error for the purposes of this work.

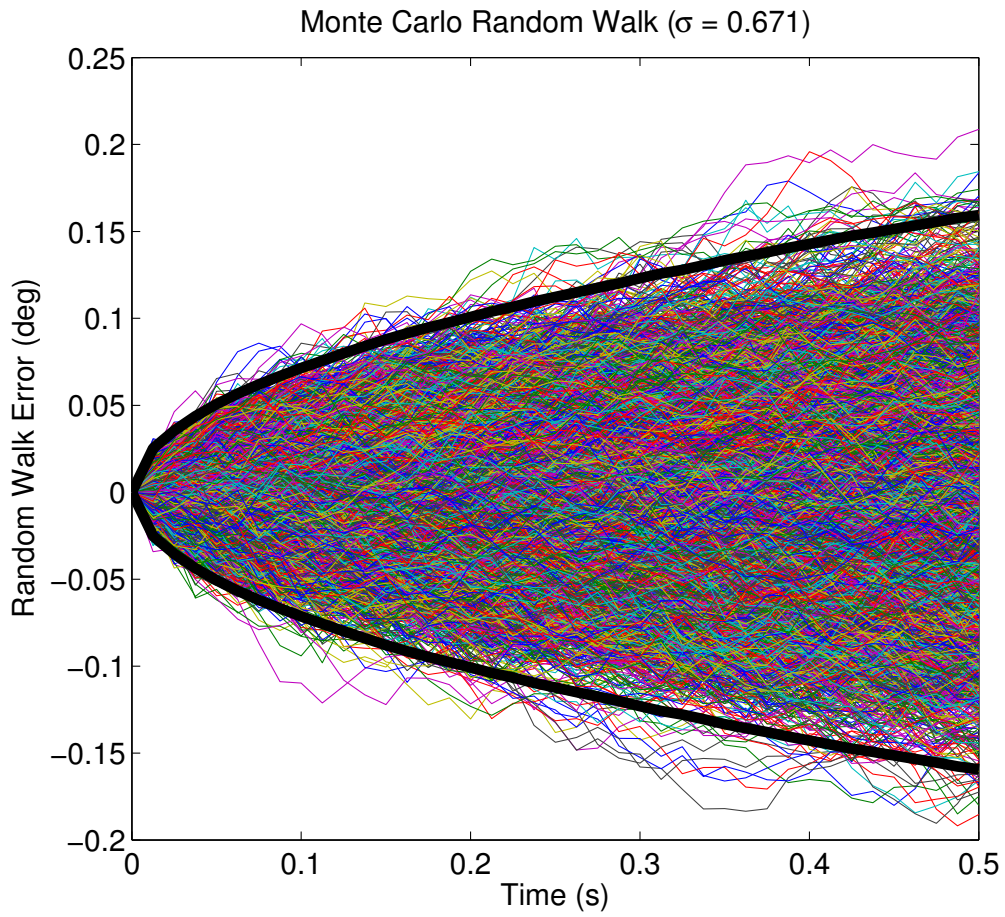


Figure 4.3: Monte Carlo simulation of random walk. 3σ envelope shown in black.

Chapter 5

Faulty Signal Detection Algorithm

The algorithm for checking for multipath or spoofed signals may be separated into three stages. The first stage is an initialization stage, during which the integer ambiguities, relative clock biases, and gyro biases are computed. Once the initialization stage is completed, RPVs and attitude may be computed using the integer ambiguities and relative clock biases while the estimated gyro biases may be used in the KF. Once the initialization stage is complete, the algorithm repeatedly alternates between a checking stage and an update stage. In the checking stage, the attitude estimate is propagated by integrating IMU measurements until a new set of GPS measurements is available. The new GPS measurements are then checked for an appropriate angle of arrival. Once the new GPS signals have been checked and the faulty signals rejected, the update stage uses the GPS measurements to calculate new positions and attitude, allowing the attitude KF to be updated with the GPS attitude estimate.

5.1 Initialization Stage

For the initialization stage, it is important to have the vehicle in an open area so that the initial integer ambiguity fixes, RPVs, and attitude estimates are correct. The specific steps for finding the integer ambiguities, RPVs, and attitude estimates were outlined in Sections 3.1 and 3.3. The initialization stage is also used to initialize the bias in the gyro measurements. If the vehicle is stationary, $r = 0$ in Equation (4.22), giving:

$$g_r = c_r + b_r + w_{gyro} \tag{5.1}$$

While the algorithm still has the potential to detect multipath or spoofed signals with double-differenced carrier phase measurements, the algorithm is much less robust using the double differences because the double differenced ambiguities are all linked to a common satellite. In the event that the common satellite for the double-differenced measurements is lost, all of the double-differenced ambiguities for that pair of antennas would be lost. Freedom from reliance on any one satellite is not the only advantage of resolving the single differenced ambiguities: the double-differenced ambiguities may be calculated directly from the single differenced ambiguities without the need for a Kalman filter or the LAMBDA method. It is important to keep the correct set of double differenced ambiguities so that the relative clock bias may be periodically checked via the double differenced RPV. The single differenced RPV should not be used to check the clock bias, as the single differenced RPV depends on the clock bias.

5.2 Check Stage

During the check stage, the attitude Kalman filter is updated using gyro measurements until a new set of GPS measurements is recorded. Once the new GPS measurements are recorded, the known baseline of the antennas, the attitude estimates, and the satellite positions are used to predict an expected AOA for each satellite. The integer ambiguities and new carrier phase measurements may be used to calculate the actual AOA for each satellite. The expected and actual AOA for each satellite are then compared, which results in a rejection of the measurements of a satellite if the difference in AOA exceeds a pre-determined threshold. The exact threshold at which a signal may be rejected relies on two factors which will be discussed at the end of Section 5.2.1.

5.2.1 Calculate AOA

A convenient way of calculating the angle of arrival of a GPS signal is to use a property of the dot product between two vectors, namely:

$$\cos \theta = \frac{\vec{a} \cdot \vec{b}}{|\vec{a}||\vec{b}|} \quad (5.2)$$

Equation (5.2) may be manipulated to give θ directly, yielding:

$$\theta = \cos^{-1} \left(\frac{\vec{a} \cdot \vec{b}}{|\vec{a}||\vec{b}|} \right) \quad (5.3)$$

In the context of this work, \vec{a} is a the vector from an antenna to a satellite and \vec{b} is the relative position vector between two antennas. Another useful property of the dot product is the property of scalar multiplication, namely:

$$c\vec{a} \cdot \vec{b} = c(\vec{a} \cdot \vec{b}) \quad (5.4)$$

where c is a scalar [17]. From Equation (5.4), we may rearrange the term inside the parentheses in Equation (5.3), giving:

$$\frac{\vec{a} \cdot \vec{b}}{|\vec{a}||\vec{b}|} = \frac{1}{|\vec{a}|}\vec{a} \cdot \frac{1}{|\vec{b}|}\vec{b} \quad (5.5)$$

We see that Equation (5.5) represents the dot product of the unit vectors of \vec{a} and \vec{b} . Using unit vectors is convenient, as the geometry matrix used in several parts of the algorithm already contains unit vectors to every satellite in view. The position vector between two antennas must be computed regularly for attitude determination of the vehicle, so the unit vector between the antennas is readily available as well. With the appropriate unit vectors, it is trivial to compute the expected AOA for each GPS signal given an expected RPV.

Mathematically, calculating the AOA of an actual signal requires only some elementary geometry. The reader will notice from Figure 5.1 that the vector \vec{r} between the two antennas and the difference L in distance to the satellite from the two antennas form two legs of a right triangle. We may then conclude that the AOA, θ , may be obtained by a property of the right triangle:




Figure 5.1: 2-D example of AOA. Solid red lines depict LOS to GPS satellite.

$$\cos \theta = \frac{L}{|\vec{r}|} \rightarrow \theta = \cos^{-1} \left(\frac{L}{|\vec{r}|} \right) \quad (5.6)$$

The length L is simple to obtain if the relative receiver clock error and single-differenced ambiguities are known. An estimate of L may be computed for each signal from the single-differenced carrier phase measurement:

$$\hat{L} = \Delta\phi_{1,2}^k - \lambda_{L1}(N_{1,2}^k) - (c dt_2 - c dt_1) \quad (5.7)$$

If the single-differenced ambiguities or relative receiver clock bias are unknown, there is no way to directly compute L . However, it is possible to calculate a related quantity, L' , which may be computed from the double-differenced carrier phase measurement. The ideal L' may be computed by:

$$L' = \begin{bmatrix} (\hat{x}_1^k - \hat{x}_1^l) & (\hat{y}_1^k - \hat{y}_1^l) & (\hat{z}_1^k - \hat{z}_1^l) \end{bmatrix} \begin{bmatrix} x_2 - x_1 \\ y_2 - y_1 \\ z_2 - z_1 \end{bmatrix} \quad (5.8)$$

This quantity may be checked by subtracting the double-differenced ambiguity term from each double-differenced carrier phase measurement:

$$\hat{L}' = \nabla \Delta \phi_{1,2}^{k,l} - \lambda_{L1} (N_{1,2}^{k,l}) \quad (5.9)$$

The double-differenced approach is less robust than the single-differenced approach as all of the double-differenced integer ambiguities rely on a common satellite. In addition, \hat{L}' does not provide an AOA, so it is a little more difficult to decide on an appropriate error tolerance. However, even with the downsides, using the double-differenced ambiguities is better than having no error detection at all.

To establish a threshold on acceptable AOA calculations, two contributing factors must be considered. The first, and most obvious, is the accuracy of the attitude estimate. Since the RPV used in checking new GPS signals is generated by the attitude estimate, it makes sense that the estimated AOA will be affected by the error in the RPV that is caused by attitude error. The second contributing factor to be considered is the noise on the single-differenced carrier phase measurement. Since \hat{L} is generated from the single-differenced carrier phase measurement, any error on the single-differenced carrier phase measurement will cause error in the estimated AOA of a signal. For example, suppose an expected value of $L = 0.8$ meters for a one-meter baseline. In that case:

$$\theta = \cos^{-1}\left(\frac{.8}{1}\right) = 36.9^\circ \quad (5.10)$$

Suppose then that there is an error of .02 meters on the single-differenced carrier phase measurement. With the .02 meter error, the estimated AOA is:




Figure 5.2: Expected AOA vs. apparent AOA error for a given single-differenced carrier phase measurement error. Note that each curve is evaluated up until the point at which the argument for inverse cosine does not yield a real angle.

$$\theta = \cos^{-1}\left(\frac{.82}{1}\right) = 34.9^\circ \text{ or } \theta = \cos^{-1}\left(\frac{.78}{1}\right) = 38.7^\circ \quad (5.11)$$

With an error of a little more than .1 carrier phase cycle, the measured AOA is thrown off by as much as two degrees. Figure 5.2 shows apparent AOA error as a function of expected AOA and various amounts of single-differenced carrier phase measurement error. Generally speaking, a longer baseline will be less affected by the single-differenced carrier phase error, and signals with a low or high expected AOA will require a wider threshold than those with an expected AOA closer to 90 degrees.

5.2.2 Alternate AOA Checking Strategies

The method for calculating the AOA in Section 5.2.1 is a good one, but the method relies on knowing the integer ambiguity, which is not always known. Specifically, the ambiguity will not be known for the case of a newly acquired satellite or one that has recently been rejected, making the AOA calculation in Section 5.2.1 impossible. It is still desirable to have some idea of the AOA of a signal though, so there are two more approaches that offer a reduced level of detection of AOA using less information.

5.2.2.1 Single-Differenced Pseudorange

The first alternate approach uses the single-differenced pseudorange measurements in place of the single-differenced carrier phase measurements. The reader will recall from Equations (3.3) and (3.4) that the single-differenced pseudorange and single-differenced carrier phase measurements will give the same quantity if the carrier phase measurement may be corrected with the cycle ambiguity, except that the pseudorange measurement is much more noisy. Then the single-differenced carrier phase measurement and integer ambiguity terms in Equation (5.7) may be replaced by the single-differenced pseudorange measurement, yielding:

$$\hat{L} = \Delta\rho_{1,2}^k - (cdt_2 - cdt_1) \quad (5.12)$$

This approach does have the drawback of the relatively high noise on the single-differenced carrier phase measurement. It is not uncommon for the error in the single-differenced pseudorange to be in excess of a meter, making this approach less effective for shorter baselines.

As an example, a one-meter baseline can have values of L between -1 and 1 meter. If the ideal L for a given satellite is .75 meters, then values of $-.25$ to 1.75 meter must be accepted for \hat{L} to account for the potential error on the single-differenced carrier phase measurements. In this particular example, the wide threshold has the potential to allow AOAs spanning over 100 degrees to be allowed. For the baselines used in this work (.94 meters and 1.91 meters),

or any baselines able to be implemented on a car, this method is not reliable enough to be used as the primary multipath/spoofing rejection algorithm.

5.2.2.2 Single-Differenced Carrier Phase

The second alternate approach utilizes the single-differenced carrier phase measurement without the knowledge of the single-differenced carrier ambiguities. If Equation (5.7) is divided by λ_{L1} and only the remainder is kept, it is seen that:

$$\text{rem}\left(\frac{\hat{L}}{\lambda_{L1}}\right) = \text{rem}\left(\frac{\Delta\phi_{1,2}^k - \lambda_{L1}(N_{1,2}^k) - (cdt_2 - cdt_1)}{\lambda_{L1}}\right) = \text{rem}\left(\frac{\Delta\phi_{1,2}^k - (cdt_2 - cdt_1)}{\lambda_{L1}}\right) \quad (5.13)$$

An estimate of \hat{L} may be calculated using Equations (5.2) and (5.6). Since this estimate is based on the RPV calculated from the IMU attitude solution, bounds may be set on acceptable values for \hat{L} based on confidence in the attitude solution. Alternatively, the ideal and calculated remainders may be compared by calculating the difference, or residual, between them and calculating acceptable error bounds on the residual.

It is interesting to note that unlike the method in Section 5.2.2.1, shorter baselines present some benefit in this method. By observing that:

$$-|\hat{r}| < L < |\hat{r}| \quad (5.14)$$

and

$$L = n\lambda_{L1} + R \quad (5.15)$$

for some integer n and remainder R where:

$$0 \leq R < \lambda_{L1} \quad (5.16)$$

then the number of potential AOAs, m , that will fulfill the remainder requirement will be:

$$\begin{aligned} m &= \frac{2|\hat{r}|}{\lambda_{L1}} + 1 & \text{if } R = 0 \\ m &= \lceil (\frac{2|\hat{r}|}{\lambda_{L1}}) \rceil & \text{if } R \neq 0 \end{aligned} \tag{5.17}$$

Note also that Equation (5.15) may be used to calculate all values of L that satisfy the remainder in Equation (5.13). The AOA associated with each value of L may be calculated using Equation (5.6). Figure 5.3 gives a visual representation of Equation (5.17), with an extra potential AOA being added for each additional half carrier cycle of baseline length. It is important to note that m in Equation (5.17) assumes multipath and spoofing signals may only come from above the vehicle. If it is assumed signals may come from below the vehicle as well (i.e. an airplane under spoofing attack), the number is doubled. If it is assumed that signals may come from above the vehicle and $|\hat{r}| < \frac{\lambda_{L1}}{2}$, then there is only one AOA that will fulfill the remainder condition. Alternatively, if $|\hat{r}|$ is large, there will be many AOAs that could fulfill the remainder requirement. There is a trade-off to the baseline length, however, as a longer baseline will give a more accurate attitude estimate on the GPS updates. Figure 5.4 shows the expected attitude error as a function of baseline length. Interestingly enough, the improved attitude accuracy offsets the increase in number of potential AOAs almost exactly. For instance, a baseline that is just under $.5\lambda_{L1}$ will have approximately 6 degrees of attitude error and one potential AOA, while a baseline of just under λ_{L1} will have approximately 3 degrees of attitude error and two potential AOAs. If it is assumed the attitude performance from the IMU is perfect, the choice between baseline lengths comes down to preference as to whether the error is represented by fewer AOAs with large error bounds or more AOAs with smaller bounds. However, practically speaking, there is an advantage to smaller baselines as any attitude error will increase the error bounds on all potential AOAs, which detracts from the attitude accuracy advantage of the longer baselines. Figure 5.5 shows a graphical representation of a case where there are four values of L , and thus four AOAs, that satisfy the remainder requirement in Equations (5.15) and

(5.16). All of the faulty signals that could fulfill the remainder requirement will have an L value that differs by an integer number of cycle ambiguities from the L of the correct AOA, as the difference in the length will not contribute to the remainder of the L value. The blue solid line represents the correct line of sight, and thus AOA, to a particular satellite, while the red solid lines represent incorrect AOAs. The dashed lines represent error bounds, which are based on attitude estimation error in the attitude filter and noise in the carrier phase measurement. The area outside of the error bounds represents AOAs that can be confidently rejected as multipath/spoofing signals. Having a correct attitude solution is especially important with this approach, as an increase in attitude error will mandate larger bounds on L , which in turn will effectively increase the error bounds on all potential AOAs. In the case of having more than a few potential AOAs, even a small attitude error could make the detection of faulty signals difficult.

The range of R that is acceptable depends not only attitude error but on the expected angle of arrival. Generally speaking, the error bounds on R are worst at 90 degrees for a given attitude error and are minimized nearing zero and 180 degrees. For example, suppose for a one-meter baseline that thresholds on R should be computed to take into account two degrees of error. Then:

$$L = 1 \cos(90^\circ) = 0 \tag{5.18}$$

Accounting for two degrees of error:

$$L = 1 \cos(88^\circ) = .0349 \text{ m} = .183 \text{ cycles} \tag{5.19}$$

In this case, an allowance must be made for $\pm .0349$ m, or $\pm .183$ carrier cycles, for a total range of .366 cycles. Since R represents the part of a whole cycle, roughly a third of all signals could potentially be accepted. Conversely, suppose the case for an AOA of 45 degrees is explored. Then:

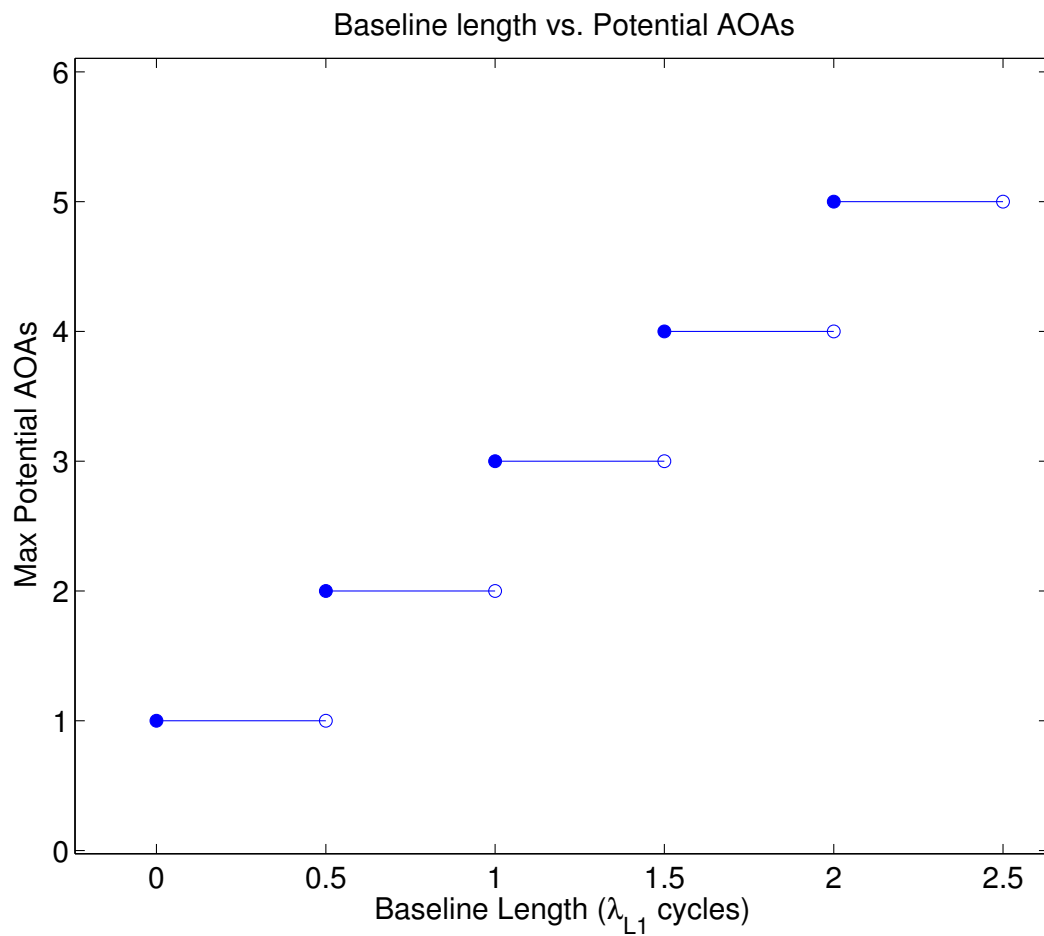


Figure 5.3: Plot of baseline length vs. number of potential AOAs.

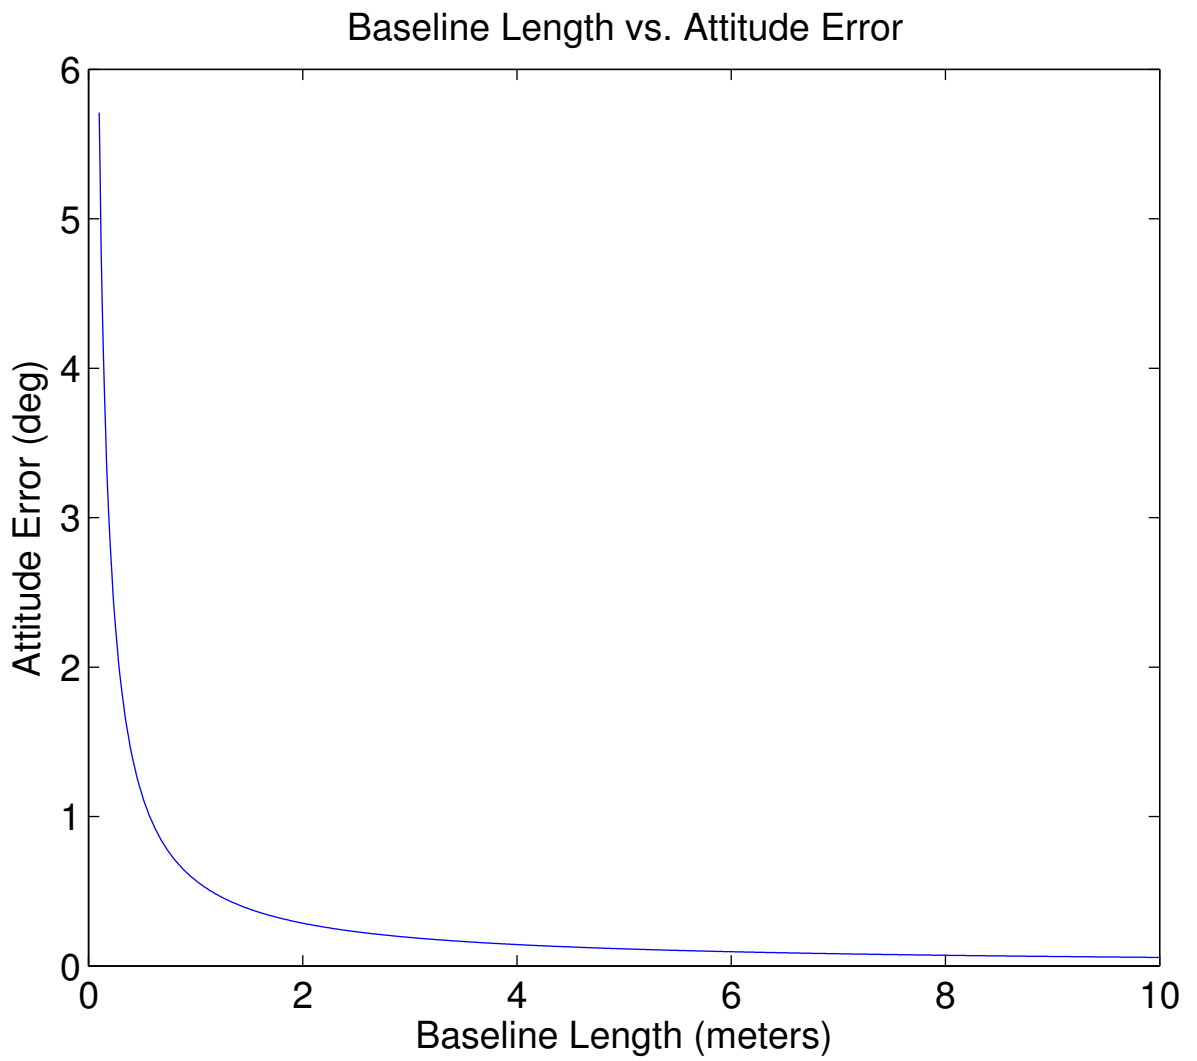


Figure 5.4: Attitude accuracy as a function of baseline length.




Figure 5.5: Graphical representation of multiple AOAs that fulfill the remainder requirement of Equations (5.15) and (5.16). The expected AOA is shown as a blue solid line, while incorrect AOAs are shown as red solid lines. Red and blue dashed lines represent error bounds based on attitude estimation error from attitude filter.

$$L = 1 \cos(45^\circ) = .707 \quad (5.20)$$

Accounting for two degrees of error:

$$L = 1 \cos(47^\circ) = .682 \quad (5.21)$$

In this second case, an allowance must only be made for $\pm .025$ m, or $\pm .131$ carrier cycles, for a total range of .262 cycles. This represents an improvement over the 90 degree case, as roughly 74 percent of signals may be rejected, instead of 63 percent with the 90 degree case. Figure 5.6 shows the relationship between expected AOA, error bounds for R , and attitude error for the two baselines used in this work. The vertical black lines represent practical boundaries of AOA, as most satellites within 15 degrees of elevation from the horizon are generally not visible or ignored due to higher levels of atmospheric error. The horizontal black dashed line represents $\pm .5$ cycles, at which point no signals can be rejected. Comparing the two plots, the shorter baseline has a clear advantage over the larger baseline in terms of allowable error.

The carrier phase residual approach is strengthened by the presence of two pairs of antennas that are orthogonal to each other. Figure 5.5 is best expanded to a three-dimensional representation by projecting the AOA bounds onto a hemisphere around the baseline. The AOA bounds will form concentric bands about the RPV of the antenna, similar to latitude lines on a globe. The 3-D representation of the AOA error bounds may be seen in Figure 5.7. The blue lines in Figure 5.7 represent a set of bounds for which a signal would be accepted by antenna pair 1, while the red lines represent a set of bounds for which a signal would be accepted by antenna pair 2. Since rejection of a faulty signal only requires the signal to be detected by one pair of antennas, a faulty signal must fall into one of the intersections of the bounds of antenna pairs 1 *and* 2 to escape detection by both antennas.




Figure 5.6: Expected AOA vs. expected error in L , plotted for multiple values total angle error. Notice that once the error reaches ± 0.5 cycles, no signals can be rejected.

Figure 5.7 also serves to compare several coordinate systems that will be useful in determining the probability of accepting a false signal as genuine. First, a standard Cartesian coordinate frame is defined with baseline 1 extending in the direction of the positive X-axis and baseline 2 extending in the direction of the positive Y-axis. The X and Y axes are in the horizontal plane with the Z-axis pointing up. Next, it is helpful to define spherical coordinate systems about the antenna baselines. θ_1 is the AOA for baseline 1 and denotes the angle between the positive X-axis and a vector extending away from the origin. θ_2 serves the same purpose for the Y-axis or baseline 2. For a vector extending outward from the origin, ϕ_1 is the right-handed rotation of the vector about the X-axis. If a vector is projected into the YZ-plane, $\phi_1 = 0$ corresponds to a vector projected onto the Y-axis only while $\phi_1 = \frac{\pi}{2}$ corresponds to a vector projected onto the z axis only. ϕ_2 denotes a vector projected in the same manner onto the XZ plane, with $\phi_2 = 0$ on the negative X-axis only and $\phi_2 = \frac{\pi}{2}$ on the Z-axis only. Finally, r_h denotes distance from the origin and is the same for both sets of coordinates. Note that ϕ as it is used here as a variable for spherical coordinates should not be confused with the carrier phase measurement.

For an AOA θ_i given by L_i that fulfills the remainder requirement, the error bounds on θ_i will be determined by signal noise and attitude error. The error bounds may be expressed as angles $\theta_{i,1}$ and $\theta_{i,2}$ or as lengths $L_{i,1}$ and $L_{i,2}$, both of which may be used to calculate the surface area of the hemisphere for which signals will be accepted. To find the surface area of the hemisphere for which incoming signals will be accepted by a single baseline, the surface area encompassed by each pair of bounds must be calculated first. Then the individual surface areas will be summed to give the total surface area of the hemisphere encompassed by the residual bounds. For n angles of arrival that satisfy the remainder requirement, the total surface area of the hemisphere encompassed by the residual bounds is given by:

$$\sum_{i=1}^n \int_0^\pi \int_{\theta_{i,1}}^{\theta_{i,2}} r_h^2 \sin \theta \, d\theta d\phi = r_h^2 \sum_{i=1}^n \pi (\cos \theta_{i,1} - \cos \theta_{i,2}) = r_h^2 \sum_{i=1}^n r_h^2 \pi \frac{L_{i,2} - L_{i,1}}{|\vec{r}|} \quad (5.22)$$




Figure 5.7: Representation of AOA error bounds projected onto a hemisphere for two orthogonal baselines. Antenna baseline 1 extends from the origin outward along the X-axis while baseline 2 extends from the origin outward along the Y axis. The pair of blue bands represent error bounds on an AOA for baseline 1 while the red bands represent error bounds on an AOA for baseline 2. The error bounds will be used to determine the probability of accepting a false signal.

where θ is the AOA, ϕ is rotation about the antenna baseline, r_h is the radius of the hemisphere, and $|\vec{r}|$ is the antenna baseline length. Also, $L_{i,2} > L_{i,1}$ so that the quantity $L_{i,2} - L_{i,1}$ will be positive and yield a positive surface area calculation. There is no need to establish a value for r_h as it will cancel out in Equation (5.25). Bounds on $\theta_{i,1}$, $\theta_{i,2}$, $L_{i,1}$, and $L_{i,2}$ are:

$$0 < \theta_{i,1}, \theta_{i,2} < \pi \quad (5.23)$$

$$-|\vec{r}| < L_{i,1}, L_{i,2} < |\vec{r}| \quad (5.24)$$

The probability of accepting a false signal P_{false} using one antenna baseline may then be given by dividing the surface area within the residual bounds by the total surface area of the hemisphere:

$$P_{false} = \left(r_h^2 \pi \sum_{i=1}^n \frac{L_{i,2} - L_{i,1}}{|\vec{r}|} \right) \frac{1}{2\pi r_h^2} = \frac{1}{2|\vec{r}|} \sum_{i=1}^n L_{k,2} - L_{k,1} \quad (5.25)$$

Note that the expression for P_{false} is derived assuming that all incoming signals will come from above the antenna baseline. If it is assumed that signals may also come from below the baseline, the AOA bounds will be projected onto a sphere instead of a hemisphere. As the bottom half of the sphere will mirror the top half, the expressions for total surface area and total area enclosed by the error bounds will both double, giving the same end result for P_{false} . To simplify the expression for P_{false} , the error on L_k may be expressed in terms of L1 carrier wavelengths: $L_{k,1}, L_{k,2} = L_k \pm a\lambda_{L1}$ where $0 \leq a \leq 0.5$. In that case:

$$L_{k,2} - L_{k,1} = (L_k + a\lambda_{L1}) - (L_k - a\lambda_{L1}) = 2a\lambda_{L1} \quad (5.26)$$

as long as $L_k + a\lambda_{L1} < |\vec{r}|$ and $L_k - a\lambda_{L1} > -|\vec{r}|$ for all L_k . The baseline length may also be expressed in terms of carrier wavelengths: $|\vec{r}| = b\lambda_{L1}$ where $\frac{n-1}{2} < b \leq \frac{n}{2}$. Equation (5.25) may then be rewritten in a much more compact format:




image of

Figure 5.8: Differential surface area dA as it relates to $d\theta_1$ and $d\theta_2$.

$$P_{false} = \frac{2na\lambda_{L1}}{2b\lambda_{L1}} = \frac{na}{b} \quad (5.27)$$

In cases where b approaches $\frac{n}{2}$ or b is large ($n \gg 1$), P_{false} reduces further to

$$P_{false} \approx 2a \quad (5.28)$$

As an example, suppose the error bounds on L_k are determined to be $\pm .1\lambda_{L1}$. If the requirements for Equation (5.28) are met, $P_{false} \approx .2$, or 20%.

As was mentioned earlier, the real strength of the carrier phase residual approach lies in using two orthogonal antenna baselines. If two sets of antennas are used to check incoming GPS measurements, the incoming signals must pass the residual test for both antennas simultaneously in order to be accepted. In Figure 5.7, passing the the residual test for both antennas at once is represented by the intersection of the error bounds from the two orthogonal sets of antennas. To calculate the cumulative area encompassed by these intersections, it is important to define dA in terms of θ_1 and θ_2 . In general, dA will be in the shape of a trapezoid as seen in Figure 5.8. From the figure, it can be seen that the area dA is given by:

$$dA = (r_h d\theta_1)(r_h d\theta_2) \frac{1}{\sin \theta_s} = (r_h d\theta_1)(r_h d\theta_2) \frac{1}{\sqrt{1 - \cos^2 \theta_s}} \quad (5.29)$$

where θ_s is one of the interior angles of the trapezoid.

Before attempting to solve for θ_s , it is helpful have expressions for the X and Y coordinates in terms of θ_1 , ϕ_1 , θ_2 , and ϕ_2 :

$$X = r_h \cos \theta_1 = -r_h \sin \theta_2 \cos \phi_2 \quad (5.30)$$

$$Y = r_h \cos \phi_2 = r_h \sin \theta_1 \cos \phi_1 \quad (5.31)$$

It is known that $\hat{\phi}_1$ is parallel to lines denoting the error bounds for θ_1 . Likewise, $\hat{\phi}_2$ is parallel to the error bounds for θ_2 . As such, the angle between the sides may be found by the dot product of $\hat{\phi}_1$ and $\hat{\phi}_2$. Equations for $\hat{\phi}_1$ and $\hat{\phi}_2$ are given by:

$$\hat{\phi}_1 = -\sin \phi_1 \hat{y} + \cos \phi_1 \hat{z} \quad (5.32)$$

$$\hat{\phi}_2 = \sin \phi_2 \hat{x} + \cos \phi_2 \hat{z} \quad (5.33)$$

Taking the arccosine of the dot product divided by the product of the magnitude of the vectors gives:

$$\theta_s = \arccos\left(\frac{\hat{\phi}_1 \cdot \hat{\phi}_2}{|\hat{\phi}_1||\hat{\phi}_2|}\right) = \arccos(\cos \phi_1 \cos \phi_2) \quad (5.34)$$

Substituting Equation (5.34) into Equation (5.29) yields:

$$\begin{aligned} dA &= (r_h d\theta_1)(r_h d\theta_2) \frac{1}{\sqrt{1 - \cos^2(\arccos(\cos \phi_1 \cos \phi_2))}} \\ &= (r_h d\theta_1)(r_h d\theta_2) \frac{1}{\sqrt{1 - (\cos \phi_1 \cos \phi_2)^2}} \end{aligned} \quad (5.35)$$

The $(\cos \phi_1 \cos \phi_2)$ term can be expressed in terms of θ_1 and θ_2 from Equations (5.30) and (5.31), giving:

$$da = \frac{r_h^2}{\sqrt{1 - (\cot \theta_1 \cot \theta_2)^2}} d\theta_1 d\theta_2 \quad (5.36)$$

The probability of accepting a false signal may then be computed by integrating over the individual intersections of AOA error bounds, summing the results, and dividing by the total surface area of a hemisphere:

$$\begin{aligned} P_{false} &= \frac{1}{2\pi} \sum_{i=1}^m \sum_{j=1}^n \int_{\theta_{i,1}}^{\theta_{i,2}} \int_{\theta_{j,1}}^{\theta_{j,2}} \frac{1}{\sqrt{1 - (\cot \theta_1 \cot \theta_2)^2}} d\theta_1 d\theta_2 \quad (5.37) \\ &= \frac{1}{6\pi} \sum_{i=1}^m \sum_{j=1}^n (1 - (\cot \theta_1 \cot \theta_2)^2)^{\frac{3}{2}} \sin \theta_1 \sin \theta_2 \cos \theta_1 \cos \theta_2 \Big|_{\theta_{i,1}}^{\theta_{i,2}} \Big|_{\theta_{j,1}}^{\theta_{j,2}} \end{aligned}$$

for $\cos^2 \theta_1 + \cos^2 \theta_2 \leq 1$. The constraint on θ_1 and θ_2 serves to ensure that the AOAs are integrated only where they intersect. For example, suppose $\theta_1 = \theta_2 = \frac{\pi}{6}$. Then $\cos^2(\frac{\pi}{6}) + \cos^2(\frac{\pi}{6}) = \frac{3}{2}$, indicating θ_1 and θ_2 do not intersect.

Another strength of the carrier phase residual method comes from the low likelihood that several authentic signals will have the same carrier phase remainders. In a single-spoofed attack, all of the spoofed signals will originate from the same location, which will cause the carrier phase remainders to be the same for all of the spoofed signals. The advantage of the carrier phase remainders being the same is that even if one of the signals happens to be within the acceptable bounds, the others will not be within the correct bounds and will be detected. In the case that multiple signals with the same carrier phase remainder are detected, the signal that happens to fall within the correct bounds may be rejected as well.

5.3 Update stage

The update stage occurs after a set of GPS measurements has been checked for AOA and the erroneous measurements have been rejected. During the update stage, a new GPS position, new RPs, and new GPS attitude solutions are computed, and the attitude Kalman

filter is updated with the new GPS attitude solution. Once all of the GPS calculations have been made, the algorithm returns to the checking stage until a new set of GPS measurements is received.

5.3.1 Choosing the correct RPV

Calculating a correct RPV between antenna pairs is crucial to erroneous signal detection. If the RPV between a pair of antennas is wrong, the attitude estimate will be wrong and the erroneous signal detection will cease to function properly. As is mentioned in the AOA section, it is preferable to know the single-differenced integer ambiguities as double-differenced ambiguities have the misfortune of all being linked to a single satellite. However, RPVs must be computed using double-differenced ambiguities regularly so that the clock bias inherent in the single differenced ambiguities may be estimated at each time step.

In Section 3.1, an approach for estimating double-differenced integer ambiguities was outlined which relied on a KF and a specialized rounding algorithm. The algorithm has the advantage of using GPS signal quality to weigh the measurements while a cost function is used to compare different sets of integer ambiguities to the float ambiguities. However, a major disadvantage of the KF approach is that the float ambiguities of newly acquired satellites may take some time to converge to values close to the correct integer ambiguity. A similar problem, mentioned in the previous paragraph, occurs if the base satellite used for double differencing is lost or reacquired, as all of the float ambiguity estimates must be reset.

An alternative course of action is to use single-differenced integer ambiguities to generate double-differenced integer ambiguities. To calculate double-differenced ambiguities from single-differenced ones, the single-differenced ambiguity of the base satellite is simply subtracted from the other single-differenced ambiguities. Perhaps the biggest advantage of

generating the double-differenced ambiguities from single-differenced ambiguities is the flexibility in choosing the base satellite in the double-differencing, as opposed to the the KF method which requires the base satellite to be chosen when the integers are fixed.

With two separate sets of double-differenced integer ambiguity estimates available, it is important to establish a metric to compare them so that the better RPV estimate may be chosen. This work investigated three different metrics to compare RPV estimates: RPV magnitude, change in RPV from previous RPV estimate, and comparison with expected RPV.

5.3.1.1 RPV magnitude

The simplest of the three metrics is checking the magnitude of the calculated RPVs against a measured baseline. The most obvious potential problem with the RPV magnitude approach is that the baseline between each set of antennas must be known accurately prior to using the system, but in most cases the baselines would be easily measurable. The more serious problem, and chief reason for not using this approach, is that fixing the integer ambiguities incorrectly can still yield an RPV with an acceptable magnitude. In fact, carrier phase errors may cause an incorrect RPV to be closer in magnitude to the correct baseline than the correct RPV in some cases.

Even if the RPV magnitude is not dependable for choosing the best RPV, it can be used to detect wrong RPVs. RPVs calculated with RTK positioning are generally accurate to within one centimeter, so a threshold may be established to eliminate RPV candidates outside of the threshold. The author used a threshold of ± 2 centimeters in this work. The threshold was chosen on the basis of RTK accuracy as well as potential error on the author's part in measuring the baseline from the center of each antenna.

5.3.1.2 Change in RPV

Another potential metric for comparing RPV candidates is to compare each of the RPV candidates to the previous RPV, then compare the magnitude of the difference between the previous and current vectors. The idea behind this approach is that an incorrect RPV might be obviously different from the previous RPV. The flaw with this approach is that it does not account for the dynamics involved in a non-stationary system. In the case of slow rotating dynamics, such as those encountered during very slow driving or driving along flat interstate, this approach may be helpful as the RPV is not expected to change much between measurements. However, in the course of normal city driving, it is not uncommon for the RPV to change drastically, especially during a sharp turn. Since the limitations on this approach stem from the comparison of the current RPV candidates to an old candidate, the next logical step is to compare the RPV candidates to the predicted RPV for the current measurements.

5.3.1.3 Expected RPV

The third approach for comparing RPV candidates is to compare the candidates with an expected RPV based on the known association between the RPs and the attitude of the vehicle. The predicted RPV, which is calculated as part of the AOA determination, may be subtracted from the RPV candidates. The magnitudes of the difference may then be compared to determine which RPV candidate is closer to the expected RPV. Under ideal circumstances, the predicted RPV will be representative of the true RPV for a given set of measurements.

Additionally, a threshold may be determined to reject RPs that are very different from the predicted RPV. The threshold will be dependent on the expected accuracy of the RPV as well as the accuracy of the attitude determination. For this work, the author used thresholds of .1 meter and .2 meters for the baselines of .94 meters and 2.14 meters,

respectively. These thresholds account for a worst case scenario of one centimeter of RPV error and approximately five degrees of separation of the predicted and candidate RPVs.

5.3.2 Calculating Single-Differenced Ambiguities

The method for calculating single-differenced cycle ambiguities is covered in Section 3.1.4, but it is important to mention here that calculating the ambiguities for each new set of GPS measurements can serve as a check against multipath signals as well. When an RPV is used to estimate single-differenced ambiguities from the single-differenced carrier phase measurement, the calculated ambiguities will not be integers, but should be very close because of the small amount of error on the carrier phase measurements. The author has observed in this work that the difference between the estimated and rounded ambiguity for a given single-differenced carrier phase measurement is generally within ± 1 cycle. Henceforth, this error will be referred to as the carrier-phase ambiguity residual error.

In concept, this method is the same idea as the signal checking method outlined in Section 5.2.2.2. The difference between the two is that the approach in Section 5.2.2.2 uses an RPV generated by an attitude estimate while calculating the ambiguities directly uses an RPV generated by the double-differenced ambiguities and carrier phase measurements. The result is that the latter method provides a much better estimate of the RPV, and as such allows for smaller thresholds on the acceptable residuals. Since the error on an RPV calculated by RTK positioning is about one centimeter at most, the result is a worst case AOA error of about .6 degrees and .3 degrees for .94-meter and 1.94-meter baselines respectively. In both cases, the error induced by the angular error accounts for an error of ± 0.05 cycles at most. In this particular case, the threshold is determined more by the noise on the signal than by the angular error associated with AOA.

Chapter 6

Experimental Results

6.1 Experimental Equipment

Experimental results were collected using a Septentrio PolaRx2e GPS receiver. The Septentrio receiver is commercially available as an attitude system and may be configured to output raw GPS measurements such as pseudoranges, carrier phase measurements, and carrier-to-noise ratios, as well as derived measurements such as position, velocity, and attitude. The availability of both raw and derived measurements make the Septentrio receiver a good choice for this work, as having both measurements allows the derived measurements of this work to be compared against those of the Septentrio receiver. Typical attitude accuracy from the receiver may be found in Table 6.1, recreated from the user manual for the receiver [1]. Note that the roll accuracy figures in Table 6.1 assume a 90 degree separation of the antennas, as seen in Figure 3.2. Roll accuracy figures for other antenna separation angles may be found in the manual but were omitted in Table 6.1 as they have no bearing on the experimental results presented in this thesis.

Baseline Length (m)	Heading Accuracy ($^{\circ}$)	Pitch Accuracy ($^{\circ}$)	Roll Accuracy ($^{\circ}$)
1	0.3	0.6	0.6
3	0.1	0.2	0.2
10	0.03	0.06	0.06

Table 6.1: Typical attitude accuracy for Septentrio PolaRx2e receiver. Roll accuracy assumes a 90 degree separation between the antennas.

The GPS antennas were mounted on an Infiniti G35 using a modified equipment/luggage rack on the roof of the car, as seen in Figure 6.1. The baseline from the main antenna to the auxiliary-1 antenna measures .94 meters while the baseline from the main antenna to



Figure 6.1: Infiniti G35 with Septentrio GPS antennas.

the auxiliary-2 antenna measures 1.905 meters. Given the baselines, it is expected that the heading and pitch accuracy will be a little worse than those shown for a 1 meter baseline in Table 6.1 – about 0.32 degrees for the heading and 0.64 degrees for the pitch. It is expected that the roll measurements will be accurate to about 0.3 degrees.

6.2 Testing Locations

The data for this work was taken primarily at two locations: an open field near the Auburn University solar house and at the National Center for Asphalt Technology (NCAT) test track in Opelika, AL. The first location was chosen primarily because it is a benign GPS environment, namely in that it is a large open space that ensures there is little to interfere with GPS signals. At the time the data was taken, there were also a large metal cargo container and small shed with a tin roof located at one end of the field, allowing for instances of multipath that are controlled (i.e., the car should not experience multipath until it is driven near the cargo container or shed). The second location, the NCAT test track, was chosen as a location that would have a regular occurrence of multipath, but not so much that

the attitude solution would be lost. In addition, the skidpad located near the track provides an area with a clear view of the sky to initialize the attitude solution. The two locations also allowed for a comparison of performance under different dynamics. The Auburn location allowed for much faster turning and had a somewhat rugged terrain, while the data taken at the NCAT test track was taken exclusively on smooth, paved surfaces with very wide turns.

6.3 Performance of Faulty Signal Rejection Methods

6.3.1 Calculate AOA

In the data runs taken at the NCAT test track, detecting multipath through the AOA proved challenging. In a run with approximately 20 instances of multipath, the AOA algorithm detected eight instances of multipath and one false detection. A data run in the field near the solar house had worse results as it was difficult to discern outliers in the data. It is likely that the higher dynamics (and thus higher attitude error) of the field data, as opposed to the track data, were at fault in masking some of the would-be outliers. Another potential cause for the worse performance could be the source of multipath signals in each of the data sets. In the track data, the multipath occurred randomly as a result of trees bordering both sides of the track and therefore were not likely to generate the same multipath error consecutively for several measurement intervals. On the other hand, multipath generated in the field data sets was due to stationary objects with sharply defined, flat sides, which should generate the same error at consecutive time intervals for a stationary vehicle near the objects. Once a signal is rejected, its integer ambiguity must be reset and thus the AOA can not be calculated until a new integer ambiguity is fixed. The loss of integer ambiguity means that the AOA can not be computed and thus explains at least some of the poor performance. Figure 6.2 shows the perceived AOA error for a separate data run with high dynamics, but no multipath. The black circles depict error based on the RPVs generated by the attitude estimates to check incoming signals, while the red circles depict error based on the calculated RPVs. As can be seen from the figure, the error on the signals

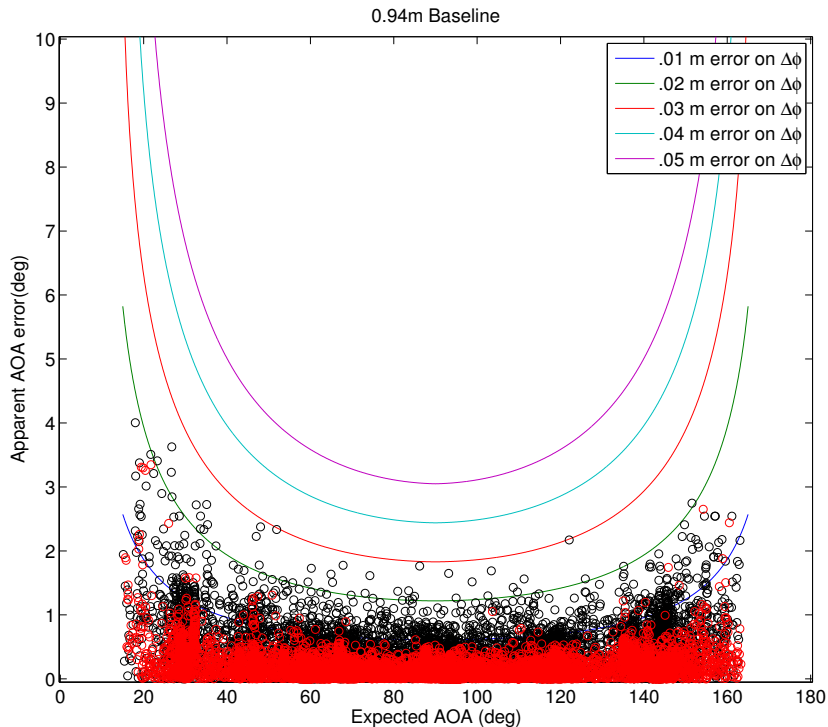


Figure 6.2: Perceived AOA error for a data run. Black circles depict AOA error from the predicted RPV while red circles depict AOA error from the calculated RPV.

from the real RPV is consistently lower than that on the RPV generated by attitude estimates. The upper bounds of both errors appear to be shifted by about a degree, which is reasonable considering the attitude error on the predicted RPVs. Ideally, the performance of the attitude-generated RPV will approach the performance of the real RPV with more accurate attitude estimation.

6.3.2 Single-Differenced Pseudorange

Detection of multipath signals using the single-differenced pseudorange measurements proved to be very unreliable. As is stated in Section 5.2.2.1, the error bounds on a single-differenced pseudorange measurement are far too large to be useful for detecting pseudoranges for the antenna baselines in this work.




Figure 6.3: Remainders of single-differenced carrier phase measurement. Red circles represent residuals generated from calculated RPV, black circles represent residuals generated from predicted RPV.

6.3.3 Carrier Phase Residual

The carrier phase residual method, explained in Section 5.2.2.2, showed performance similar to that of the calculation of AOA at the NCAT test track, with nine of 20 instances of multipath detected and a few erroneous detections. The performance was similar in the test data taken from the field near the AU solar house. Figure 6.3 shows the residuals generated from the single-differenced carrier phase measurements for a high-dynamic data set without multipath. Red circles represent remainders generated from calculated RPV while black circles represent residuals generated from predicted RPV. The results are similar to those seen in Figure 6.2, with the measurements generated by the calculated RPV much better than those from the predicted RPV. Once again, the performance related to the predicted RPV is expected to approach that of the predicted RPV with improved attitude estimation.

6.3.4 Calculating Single-Differenced Ambiguities

It was mentioned in Section 5.3.2 that calculating the single-differenced ambiguities and checking the difference between the calculated and rounded ambiguities could be used to check incoming signals. The author found this approach to yield the most consistent results in the experimental data sets – routinely picking out the most multipath signals and making the fewest erroneous rejections of signals. In the NCAT test track data set mentioned above, the single-differenced carrier phase residuals using the calculated RPV achieved nearly perfect performance in rejecting multipath. Some of the improved positioning in the field data was still not ideal but was certainly an improvement over the solution without rejecting signals. The quality of faulty signal detection from this method is a direct result of the precision of the RPs calculated with carrier phase measurements, which is shown in Figures 6.2 and 6.3.

6.4 Results Validation

A few methods were used to decide if the signals rejected from data sets were actually multipath signals. Since the primary reason for rejecting multipath is to get a better positioning solution, it stands to reason that comparing the positions of the receiver with and without the multipath signals should show a better position without the multipath signals. Figure 6.4 shows the calculated positions for a section of the field data run. The normal positioning solution is plotted in red while the positioning solution without the rejected signals is in green. The vehicle was stopped near the tin shed at the spot of the noisy position estimate, so multipath was to be expected at that location. The positioning solution without the detected multipath signals is shown to be much more reasonable than the positioning with the signals. In most of the data observed, removing the multipath signals improved the positioning solution visibly when compared to the normal positioning solution. However, it was also found that the positioning solution could potentially be made worse by removing



Figure 6.4: Comparison of position estimate with and without rejected signals.

the multipath signals. The circumstances under which the positioning was made worse by rejecting signals will be discussed more in Section 6.5.

Another metric for validation of the rejected signals is to check the pseudoranges for a change at the time of signal rejection. Figure 6.5 shows the difference between the expected and measured pseudorange for a particular data run on the skidpad as well as the carrier phase residual for that signal. The obvious change in pseudorange in Figure 6.5 corresponds with the rise in the carrier phase residual above the .1-cycle threshold, which is to be expected if the signal takes an indirect path to the receiver.

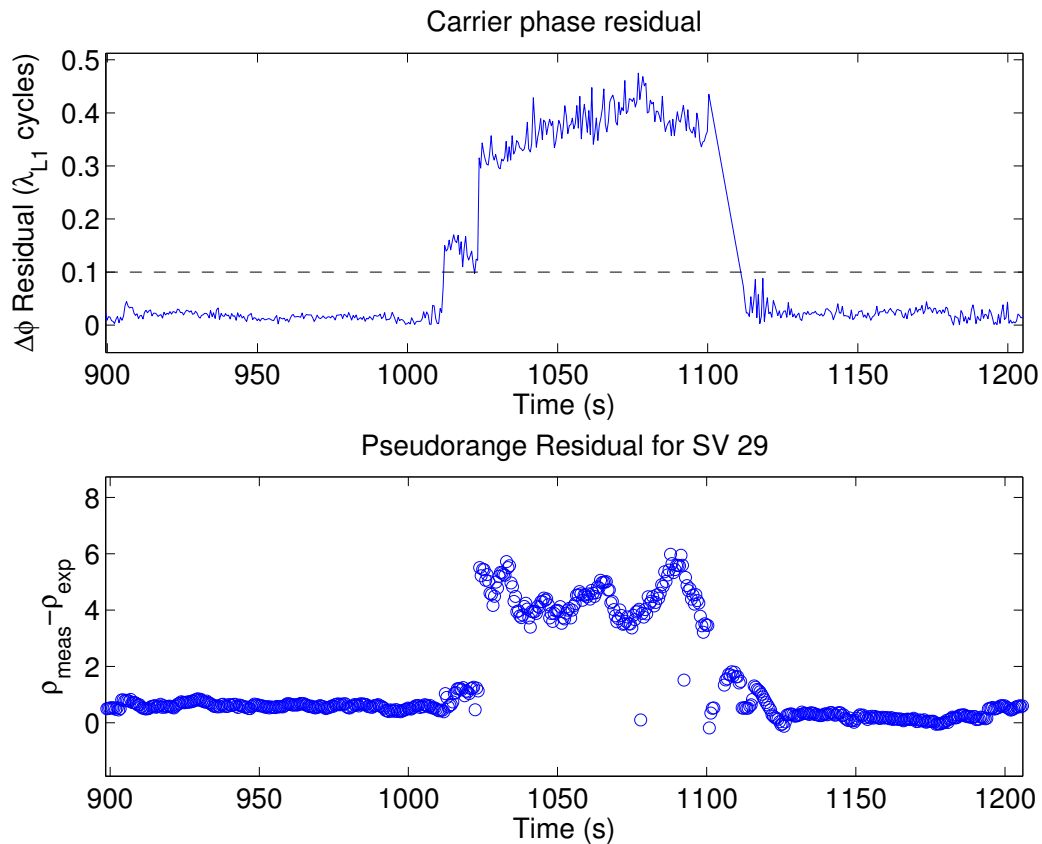


Figure 6.5: The upper sub-plot shows the carrier phase residual measurement for an antenna pair during a data run. The portion of graph above the .1-cycle threshold indicates the detection of a faulty signal. The lower sub-plot is the pseudorange residual for one of the antennas, showing the measured pseudorange differs significantly from the expected pseudorange when a faulty signal is detected.

6.5 Effects of Satellite Geometry on Positioning

In Section 6.4, it was shown that the rejection of multipath signals could yield an improved positioning solution over the original positioning solution. However, it is important to note that GPS positioning accuracy is not based entirely on signal quality. The second major component of GPS positioning accuracy is the number and geometry, or relative positions, of available satellites. The effect of the satellite geometry on positioning accuracy is summarized by dilution of position (DOP) metrics. The DOP metrics alone are unitless but may be multiplied by the standard deviation of the pseudorange measurement to yield an expected standard deviation in the positioning solution. Various DOP metrics are commonly used in the GPS community to express the expected accuracy of subsets of the positioning solutions:

- Geometrical DOP (GDOP): combined total positioning error and clock error
- Position DOP (PDOP): total positioning error
- Horizontal DOP (HDOP): total horizontal (2D) positioning error
- Vertical DOP (VDOP): vertical (1D) positioning error
- Time DOP (TDOP): clock bias error

The most important metric for this work is the HDOP, as this thesis focuses on positioning on the ground. In select cases, it was found that removing multipath signals could actually make the positioning solutions worse. One such case may be seen in Figure 6.6. The path in green depicts positioning calculations in which no satellites were rejected while the path in red depicts positioning after the rejection of signals from two satellites. In this particular case, the positioning is actually worse after the rejection of the multipath signals. The reason the positioning is worse may be seen in the HDOP, which is approximately 1.0 before the signals are rejected and 2.2 after. Given good quality GPS signals, an increase in the HDOP




Figure 6.6: Comparison of positioning with and without rejected signals. In this particular case, rejecting the signals yields a worse positioning solution.

by a factor of two generally means the standard deviations of the horizontal positioning will be expected to be twice as large. In this particular case, it appears as though the effect of the multipath error on the calculated position is less than the effect of using two fewer satellites. In most of the data, only one satellite is removed at a time, which generally changes the HDOP by less than 10 percent. This relatively small change in the HDOP, as opposed to the case for the positioning solution in Figure 6.6, explains why the positioning was improved by removing the detected multipath signals most of the time.

Chapter 7

Conclusions and Future Work

7.1 Conclusions

From the current experimental results, checking the residuals of calculated single-differenced ambiguities is the most reliable way to detect multipath; however, an improvement in the attitude filter to allow better prediction of the RPV could allow better performance in the AOA and carrier phase residual methods. The AOA method would be especially useful as it is the only method to use the single-differenced carrier phase ambiguity. The pseudorange method was judged to be unreliable with the antenna baselines on the current experimental setup due to the large error on the single-differenced pseudorange measurements, although it could potentially work better on an airplane or ship.

It is interesting to note than in Figure 6.3, the residuals generated by the predicted RPV fit the error curves much better than the residuals generated from the calculated RPV. This peculiarity is likely due to the attitude accuracy inherent in the calculated RPVs, which should result in an AOA accuracy of about 0.6 degrees for the short baseline and 0.3 degrees for the long one. The square shape of the red ambiguities, as opposed to the curve of the black ones, suggest that the residuals from the calculated RPVs are probably dominated by noise on the signal rather than attitude inaccuracy.

Since Figure 5.2 assumes that the signal itself is the only source of error (i.e. the user has “perfect” attitude estimates) in the AOA calculations, it may be used to determine the maximum performance of an AOA detection system with a given baseline. Additionally, if a user would like to know the expected performance of an AOA system with a given amount of attitude error, the curves in Figure 5.2 may be shifted up the positive y-axis by the amount

of AOA error the attitude is expected to produce. The shifted curve will represent the overall expected performance of the AOA detection system.

The data showed that removing one faulty GPS signal would usually produce an improvement in the positioning solution. However, removing two signals in certain situations gave a worse positioning solution than using the faulty signals. The likely cause of the different results is shown in the HDOP, which generally changed little with the removal of one satellite and could change by a factor of two or more with the removal of two satellites. Therefore, it may be prudent to take the HDOP into consideration when deciding whether or not to remove faulty GPS signals.

7.2 Future Work

The first improvement on this work should be to improve the attitude filter. The AOA and carrier phase residual methods have potential to match performance with the calculated RPV carrier residual method if the attitude filter is improved. Another benefit of an improved attitude filter would be the potential to reject bad satellites for a short amount of time if the GPS attitude solution is lost, potentially allowing for the recovery of the attitude solution without the use of a Kalman filter for the carrier phase ambiguities.

The anti-spoofing capabilities of these methods should also be explored. It is illegal in the United States to broadcast GPS signals outside of an anechoic chamber without permission from the federal government, so the remaining options for testing the system are to test in an anechoic chamber or with permission from the government. The indoor testing would potentially involve several spoofers representing the “real” GPS signals and another spoofer as the erroneous signal. Testing in an anechoic chamber poses a number of problems, namely:

- It may be prohibitively expensive to build an anechoic chamber large enough for the dynamic vehicle testing.

- Multiple spoofers would have to be used to imitate satellite geometry as the “real” signals in the testing.
- If the spoofers representing the GPS satellites are not sufficiently far away from the receiver antennas, relative carrier phase positioning may not work as it is assumed the signals received by each antenna are parallel to each other.

Given the difficulties of testing indoors, the most feasible test would be one performed outdoors with permission from the United States government.

Lastly, the algorithm may be modified to account for the HDOP in the decision of rejecting a faulty GPS signal. It was shown that the positioning solution could be made worse in situations where the HDOP was significantly changed by the rejection, so a method to quantify the effect of a given multipath signal on the positioning solution might be implemented so that signals may be intelligently rejected. Of course, all signals should be rejected in cases where anti-spoofing is desired.

Bibliography

- [1] *Polarx2/2e User Manual 3.2.0*. Septentrio Satellite Navigation, 2007.
- [2] Alison Brown, Dale Reynolds, D Roberts, and S Serie. Jammer and interference location system—design and initial test results. In *Proceedings of the ION GPS*, volume 99. Citeseer, 1999.
- [3] Clark Cohen. *Attitude Determination Using GPS*. PhD thesis, Stanford University, 1992.
- [4] Paul De Jonge and CCJM Tiberius. The lambda method for integer ambiguity estimation: implementation aspects. *Publications of the Delft Computing Centre, LGR-Series*, 12(12):1–47, 1996.
- [5] Warren Flenniken. Modeling inertial measurement units and analyzing the effect of their errors in navigation applications. Master’s thesis, Auburn University, 2005.
- [6] Todd E Humphreys, Brent M Ledvina, Mark L Psiaki, Brady W O Hanlon, and Paul M Kintner Jr. Assessing the spoofing threat: Development of a portable gps civilian spoofer. In *Proceedings of the ION GNSS International Technical Meeting of the Satellite Division*, 2008.
- [7] Ali Jafarnia-Jahromi, Ali Broumandan, John Nielsen, and Gérard Lachapelle. Gps vulnerability to spoofing threats and a review of anti-spoofing techniques. *International Journal of Navigation and Observation*, 2012, 2012.
- [8] Edwin L Key. Techniques to counter gps spoofing. *Internal Memorandum, MITRE Corporation*, 1995.
- [9] Scott Martin. Closely coupled gps/ins relative positioning for automated vehicle convoys. Master’s thesis, Auburn University, 2011.
- [10] Charles E McDowell. Gps multipath mitigation using a multi-element antenna array, July 16 2002. US Patent 6,421,000.
- [11] Charles E McDowell. Gps spoofer and repeater mitigation system using digital spatial nulling, July 31 2007. US Patent 7,250,903.
- [12] Paul Y Montgomery, Todd E Humphreys, and Brent M Ledvina. Receiver-autonomous spoofing detection: Experimental results of a multi-antenna receiver defense against a portable civil gps spoofer. In *Proceedings of the ION International Technical Meeting*, 2009.

- [13] PY Montgomery, Todd E Humphreys, and Brent M Ledvina. A multi-antenna defense: Receiver-autonomous gps spoofing detection. *Inside GNSS*, 4(2):40–46, 2009.
- [14] Bradford W. Parkinson and James J Spiker. *The Global Positioning System: Theory and Application*, volume 1. Aiaa, 1996.
- [15] Daniel P Shepard, Jahshan A Bhatti, and Todd E Humphreys. Drone hack. *GPS World*, 23(8):30–33, 2012.
- [16] Dan Simon. *Optimal state estimation: Kalman, H infinity, and nonlinear approaches*. Wiley-Interscience, 2006.
- [17] James Stewart. *Calculus*. Thomson Brooks/Cole, 2003.
- [18] JB Tsui. Fundamentals of global positioning system receivers: a software approach, 2005.
- [19] Hengqing Wen, Peter Yih-Ru Huang, John Dyer, Andy Archinal, and John Fagan. Countermeasures for gps signal spoofing. In *ION GNSS*, pages 13–16, 2005.

Appendix A

Determination of Position

Determination of position using GPS is most commonly achieved using pseudorange measurements. In setting up the position determination, it is helpful to rearrange Equation (2.5), giving:

$$\rho_i^k = \begin{bmatrix} (e_x)_k^i & (e_y)_k^i & (e_z)_k^i & 1 \end{bmatrix} \begin{bmatrix} x^k - x_i \\ y^k - y_i \\ z^k - z_i \\ cdt_i \end{bmatrix} + cdt^k + T_i^k + I_i^k + \epsilon_i^k \quad (\text{A.1})$$

where $\begin{bmatrix} (e_x)_k^i & (e_y)_k^i & (e_z)_k^i \end{bmatrix}$ represents the unit vector from the receiver to the satellite. Note that parameters are included in the ephemeris data for the determination of satellite clock error as a function of time, so cdt^k is a deterministic quantity. Also, T_i^k , I_i^k , and ϵ_i^k are all random variables, so they may effectively be lumped into one term. Equation (A.1) may then be rewritten:

$$\rho_i^k - cdt^k = \begin{bmatrix} (e_x)_k^i & (e_y)_k^i & (e_z)_k^i & 1 \end{bmatrix} \begin{bmatrix} x^k - x_i \\ y^k - y_i \\ z^k - z_i \\ cdt_i \end{bmatrix} + w_i^k \quad (\text{A.2})$$

Equation (A.2) is close to the format of least squares, except the unit vectors from the receiver to the satellite are unknown. Instead of attempting to estimate x_i , y_i , z_i , and cdt_i directly, it is more prudent to initialize a position estimate and compare the real

measurements to the expected measurements given the position estimate. Using an estimated position, the estimated range to a given satellite is:

$$\hat{r}_i^k = \begin{bmatrix} (\hat{e}_x)_k^i & (\hat{e}_y)_k^i & (\hat{e}_z)_k^i & 1 \end{bmatrix} \begin{bmatrix} x^k - \hat{x}_i \\ y^k - \hat{y}_i \\ z^k - \hat{z}_i \\ c\hat{d}t_i \end{bmatrix} \quad (\text{A.3})$$

If the assumption is made that $\begin{bmatrix} (e_x)_k^i & (e_y)_k^i & (e_z)_k^i & 1 \end{bmatrix} \approx \begin{bmatrix} (\hat{e}_x)_k^i & (\hat{e}_y)_k^i & (\hat{e}_z)_k^i & 1 \end{bmatrix}$, where $\begin{bmatrix} (\hat{e}_x)_k^i & (\hat{e}_y)_k^i & (\hat{e}_z)_k^i \end{bmatrix}$ is the estimated unit vector from the receiver to the satellite, Equation (A.3) may be subtracted from Equation (A.2), yielding:

$$\rho_i^k - cdt^k - \hat{r}_i^k = \begin{bmatrix} (\hat{e}_x)_k^i & (\hat{e}_y)_k^i & (\hat{e}_z)_k^i & 1 \end{bmatrix} \begin{bmatrix} \hat{x}_i - x_i \\ \hat{y}_i - y_i \\ \hat{z}_i - z_i \\ c\hat{d}t_i - cdt_i \end{bmatrix} + w_i^k \quad (\text{A.4})$$

Equation (A.4) does fit the least squares format, allowing the error in the position estimate to be solved for directly. This gives the position estimate correction as:

$$\begin{bmatrix} \delta x_i \\ \delta y_i \\ \delta z_i \\ \delta cdt_i \end{bmatrix} = (H^T H)^{-1} H^T \vec{y} \quad (\text{A.5})$$

where H , called the geometry matrix, is a matrix in which each row consists of the estimated unit vector to a satellite followed by a 1, as seen in Equation (A.4), and \vec{y} is a vector of the terms seen in the left side of Equation (A.4). It is important to note here that there is some inherent error in the geometry matrix, caused by the error in the initial guess of the user's position. However, once the position estimate correction has been applied, the new

estimate of the receiver position is closer than the previous estimate, so a new geometry matrix will be more accurate than the previous one. The new receiver position estimate and new geometry matrix may then be used to compute a new position estimate correction. By recursively carrying out the least squares calculations, the correction terms should converge towards zero quickly. It is only necessary to compute the correction terms until they become sufficiently small.

Appendix B

Kalman Filter

The Kalman filter is known as an “optimal” estimator, where optimal is defined in this sense as reducing the expected estimation error at each time step. The equations for a Kalman filter are based on a few assumptions, the first of which is that the system is linear. While very few systems are truly linear, this can be a good approximation for many systems, including nonlinear systems about an operating point. There are also modified formulations of the Kalman filter designed to be used with non-linear systems.

The second assumption is that the process noise introduced into the system and the measurement noise are Gaussian, zero-mean, uncorrelated, and white. If the noise is non-Gaussian but meets the other noise characteristics, the Kalman filter is still the best linear filter for the system. In the case that some of the other noise characteristics are also not met, the Kalman filter may be modified to solve the problem [16].

The third assumption is that the system and noise are modeled perfectly. If there is error in the system or noise models contain error, the filter is no longer an “optimal” filter.

The discrete-time Kalman filter may be divided into two parts: the time update and the measurement update.

B.1 System

For the purpose of this work, discrete-time state-space models will be presented in the following form:

$$x_{k+1} = A_k x_k + B_k u_k + w_k \tag{B.1}$$

$$y_k = C_k x_k + \nu_k \quad (\text{B.2})$$

where x_k is the column vector of system states at time instance k , A_k is the discrete-time state transition matrix, B_k is the discrete-time input matrix, u is the known input to the system, w is a column of process disturbances, y_k is the column of measurements associated with states x_k , C_k is the discrete time measurement matrix, and ν is a column of measurement noise. It is assumed that w and v are unknown and random, but their statistical properties are known.

B.2 Time Update

The time update of the Kalman filter is a “prediction” of sorts, using the system model, system states, and known inputs to the system to provide an estimate of the states of the system at the next time step. The time update of the Kalman filter is implemented through the following equations:

$$\hat{x}_k^- = A_{k-1} \hat{x}_{k-1} + B_{k-1} u_{k-1} \quad (\text{B.3})$$

$$P_k^- = A_{k-1} P_{k-1}^+ A_{k-1}^T + Q_{k-1} \quad (\text{B.4})$$

Equation (B.3) gives what is known as the *a priori* estimate, an estimate that propagates the state estimate forward using the system model. Comparing Equations (B.1) and (B.3), it is clear that they are nearly the same equation, except Equation (B.3) does not include the process disturbance term. The omission of the disturbance term makes sense intuitively, as the disturbance is random and unknown; therefore, the a priori estimate is the “best guess” of the states at the next time step, given the system model.

Equation (B.4) gives the a priori estimate covariance matrix, a measure of the uncertainty in the estimates. P_k^- will be $n \times n$ for a system with n states, with each diagonal

term representing the expected variance of each state. Each off diagonal term will be of the form $\sigma_i\sigma_j\rho$, where σ_i is the expected standard deviation of state i , σ_j is the expected standard deviation of state j , ρ is the statistical correlation of the two states, and $i \neq j$. It is important to emphasize that P_k^- represents the *expected* variances and covariances of the states. If the system model or noise characteristics are wrong, P_k^- will not represent the actual variances and covariances of the states. Q_{k-1} is the discrete time covariance matrix for the process disturbances.

B.3 Measurement Update

The measurement update of the Kalman filter serves as a “correction” of the a priori estimate using measurements of the states. The measurement update of the Kalman filter is implemented through the following equations:

$$L_k = P_k^- C_k^T (C_k P_k^- C_k^T + R_k)^{-1} \quad (\text{B.5})$$

$$\hat{x}_k^+ = \hat{x}_k^- + L_k (y_k - C_k \hat{x}_k^-) \quad (\text{B.6})$$

$$P_k^+ = (I - L_k C_k) P_k^- (I - L_k C_k)^T + L_k R_k L_k^T \quad (\text{B.7})$$

Equation (B.5) gives the Kalman gain, used in Equation (B.6). Unlike the gains used in classical estimators, which are generally chosen based on rise time, settling time, and other characteristics, the Kalman gain is recalculated with each iteration to give an optimal estimate with minimal expected error. Inspecting Equation (B.5), the Kalman gain is calculated based on the measurement noise covariance matrix (R_k), P_k^- , and C_k . It is helpful to inspect Equation (B.5) for a single state, single measurement system, in which case it becomes:

$$L_k = \frac{P_k^-}{P_k^- + R_k} \quad (\text{B.8})$$

Notice that because the system would have a single state and a single measurement, all terms in Equation (B.8) are scalars. Equation (B.8) captures the very essence of the Kalman gain: how much the state estimate is changed in Equation (B.6) is based on the relative uncertainty of the current state estimate versus the new measurements.

If the uncertainty in the state estimate is much higher than the uncertainty in the measurements ($P_k^- \gg R_k$), L_k approaches 1 and the state estimate will be corrected heavily toward the new measurement. Making the large correction makes intuitive sense, considering the measurement is a much better representation of the state than the current state estimate.

As would be expected, L_k would approach zero if the state estimate has much lower uncertainty than the measurement ($P_k^- \ll R_k$), and the measurement would mostly be ignored. Ignoring a bad measurement also makes sense intuitively if the state estimate is known to be reliable. If P_k^- is routinely much smaller than R_k , the filter is said to “go to sleep;” that is, the state estimate will change very little for each new measurement.

Equation (B.6) gives the *a posteriori* estimate, or the state estimate after the current system measurements have been taken into account. The $(y_k - C_k \hat{x}_k^-)$ term is known as the innovation or residual and is the difference between the actual measurement and the expected measurement (based on the current state estimate). How much the a priori estimate is changed is based on the product of the Kalman gain and the residual.

Equation (B.7) gives the covariance matrix for the a posteriori estimate, which is similar in concept to Equation (B.4). The difference between the two equations is that Equation (B.4) depends on the process disturbances while Equation (B.7) depends on measurement noise. In the current formulation, the new state estimate covariance may be calculated before or after the a posteriori estimate as long as the Kalman gain has been computed. Other formulations of the Kalman filter exist that calculate the a posteriori state estimate covariance before the Kalman gain, and then use the covariance in the calculation of the

Kalman gain. More detailed information about the alternate implementations of the Kalman filter, as well as details on its derivation, may be found in [16].