

Spring 1-1-2015

Improving the accuracy of in-situ lower ABL wind measurements using sUAS

Douglas E. Weibel

University of Colorado at Boulder, deweibel@gmail.com

Follow this and additional works at: https://scholar.colorado.edu/asen_gradetds

 Part of the [Aerospace Engineering Commons](#), [Environmental Sciences Commons](#), and the [Robotics Commons](#)

Recommended Citation

Weibel, Douglas E., "Improving the accuracy of in-situ lower ABL wind measurements using sUAS" (2015). *Aerospace Engineering Sciences Graduate Theses & Dissertations*. 105.

https://scholar.colorado.edu/asen_gradetds/105

This Dissertation is brought to you for free and open access by Aerospace Engineering Sciences at CU Scholar. It has been accepted for inclusion in Aerospace Engineering Sciences Graduate Theses & Dissertations by an authorized administrator of CU Scholar. For more information, please contact cuscholaradmin@colorado.edu.

**Improving the accuracy of in-situ lower ABL wind
measurements using sUAS**

by

Douglas E. Weibel

B.S., Rensselaer Polytechnic Institute, 1987

M.S., Syracuse University, 1990

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Aerospace Engineering Science
2015

This thesis entitled:
Improving the accuracy of in-situ lower ABL wind measurements using sUAS
written by Douglas E. Weibel
has been approved for the Department of Aerospace Engineering Science

Prof. Dale Lawrence

Prof. Scott Palo

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Weibel, Douglas E. (Ph.D., Aerospace Engineering Science)

Improving the accuracy of in-situ lower ABL wind measurements using sUAS

Thesis directed by Prof. Dale Lawrence

Wind measurement is a challenging problem as wind is a spatiotemporal vector field and obtaining a set of measurements with desired spatial and temporal sampling is not always possible. Small unmanned-aircraft systems (sUAS) are an emerging technology offering new capabilities for in-situ sensing in the lower atmospheric boundary layer (ABL) with cost, operational and safety advantages over larger airborne wind-measurement systems including the potential to use swarms of sUAS for denser sampling. Analysis of contemporary sUAS-based wind-measurement systems shows that the wind-measurement error is not dominated by a single source. Errors introduced by the relative-wind sensor, the airframe velocity estimate, the airframe attitude estimate, and the airframe attitude-rate estimate are all significant. This work focuses on improving sUAS velocity estimation and attitude estimation. State estimation for contemporary sUAS is based on sensors that are small, light and inexpensive but have much poorer performance than larger, navigation-grade sensors used in manned aircraft. Careful choice of a sensor fusion algorithm can allow for sUAS velocity-estimation accuracy on the order of 1 cm/s and attitude-estimation accuracy on the order of 0.1° when using inexpensive consumer grade sensors. With a high-quality relative-wind sensor these levels of state-estimation accuracy allow for wind-measurement with accuracy on the order of 1 cm/s. However, some key issues in sUAS state estimation, particularly for sUAS flying in wind, have heretofore received little attention. Wind-gust-induced motion can have a significant effect on sUAS state estimation and many sensor-fusion algorithms proposed for use with sUAS have poor performance when subjected to wind-gust-induced motion.

Specific approaches to sUAS attitude and velocity estimation are proposed based on analytic results and testing of contemporary sensors. In particular it is shown that an extended Kalman filter estimating attitude and gyroscope bias drift rate, and using time-differenced GPS veloc-

ity measurements to estimate translational acceleration, can provide the desired attitude-estimate accuracy with contemporary sUAS suitable sensors even in the presence of strong winds and turbulence. Simulation of this algorithm's performance when used on a sUAS flying in turbulent conditions show as much as an order-of-magnitude improvement in performance compared to other algorithms presented in the literature. Analysis of sUAS velocity estimation shows a strong sensitivity to the performance of the GPS receiver's velocity-measurement accuracy. Small UAS generally use inexpensive commercial-off-the-shelf (COTS) GPS receivers which often have higher error levels when experiencing accelerating motion as is typical for sUAS flying in wind. A particular COTS GPS receiver that outputs raw satellite-channel-measurement data is examined. Analysis of test data was used to develop a method of using the COTS GPS raw-measurement data to produce a velocity estimate with the desired accuracy.

Desire for an additional sensor suitable for use in either validating or improving SUAS attitude-estimation accuracy motivated the development of an optical reference-vector sensor system capable of making measurements with an accuracy better than 0.1 degrees while operating outside in full daylight at ranges in excess of 100 meters. Field tests were conducted using a typical sUAS, the optical reference-vector sensor system and contemporary sUAS-state-estimation sensors. Analysis of flight test data supports the assertion that sUAS-based wind-measurement systems with accuracy on the order of 1 cm/s are possible with contemporary sensors.

Dedication

This dissertation is dedicated to my parents. They are directly responsible for most qualities in myself that have allowed me to successfully pursue my course of studies at the University of Colorado and the support of my family, which I have enjoyed all of my life, is always a joy to me. My parents are a continual example both with regard to working earnestly towards goals and with balancing work, family and self. They set my life in motion and continue to teach and inspire me.

Acknowledgements

Professor Scott Palo and Professor Dale Lawrence receive my sincere thanks and appreciation for advising me and guiding my research. I am extremely grateful for their generosity with their time and for allowing me so much latitude in my research. I have sincerely enjoyed working with and learning from both of them.

The members of my dissertation committee, Professor Eric Frew, Professor Nissar Ahmed and Professor Lucy Pao have all been generous with their time and expertise. I thank them for their contributions and support. I further thank Professor Frew for his input and the great deal I learned about teaching from him while team-teaching Aircraft Dynamics.

I happily acknowledge as well the many friends, colleagues, students and teachers who collaborated with me and supported my research and learning. The Aerospace Engineering Science department, and the University of Colorado as a whole, has been very welcoming and supportive. I particularly wish to thank Gabe LoDolche for the many hours we spent fifteen minutes at a time collaborating over coffee.

Most importantly I wish to thank my family for the support I have received over the past years. My wife, Marjorie, my son, Drew, and my daughter, Ashley, have all been uniformly supportive and never begrudged the time I have devoted to my studies and research.

Contents

| Chapter | |
|----------|---|
| 1 | Executive summary and contributions 1 |
| 2 | Measurement of wind 8 |
| 2.1 | Contemporary wind measurement systems 12 |
| 2.1.1 | Ground-based in-situ 12 |
| 2.1.2 | Ground-based remote sensing 13 |
| 2.1.3 | Airborne in-situ 14 |
| 2.1.4 | Airborne and space-borne remote sensing 15 |
| 2.2 | Airborne wind measurement 15 |
| 3 | Measurement of wind from small unmanned aircraft 23 |
| 3.1 | Example areas of study where sUAS may have particularly value 24 |
| 3.1.1 | Study of the structure of wind turbine wakes 25 |
| 3.1.2 | Study of vertical momentum flux and coherent structures in the lower atmospheric boundary layer. 26 |

| | | |
|-------|---|----|
| 3.1.3 | Validation of the frozen turbulence hypothesis for eddy-covariance techniques for vertical-flux measurements. | 27 |
| 3.1.4 | Study of the urban boundary layer | 28 |
| 3.1.5 | Observing entrainment processes | 28 |
| 3.1.6 | Desirable system characteristics for sUAS-based wind measurement | 29 |
| 3.2 | Overview of existing sUAS-based wind measurement systems | 30 |
| 3.3 | Accuracy of wind measurement from sUAS | 32 |
| 3.3.1 | Sensitivity of the inertial-frame wind measurement to constituent measurements | 32 |
| 3.3.2 | Optimal framework for inertial-wind measurement | 36 |
| 3.4 | Research questions | 37 |
| 4 | sUAS attitude estimation | 39 |
| 4.1 | sUAS sensor characteristics | 42 |
| 4.1.1 | MEMS gyroscopes | 42 |
| 4.1.2 | MEMS accelerometers | 44 |
| 4.1.3 | Magnetoresistive magnetometers | 45 |
| 4.1.4 | Commercial off-the-shelf (COTS) GPS/GNSS receiver modules | 48 |
| 4.1.5 | Other sensors | 50 |
| 4.2 | Sensor fusion for sUAS attitude estimation | 51 |
| 4.2.1 | System model | 51 |
| 4.2.2 | Filter Structure —state estimate update variations | 55 |

| | | |
|----------|--|-----|
| 4.2.3 | Filter Types - federated versus cascaded filters | 67 |
| 4.2.4 | Specific-force correction methods | 73 |
| 4.2.5 | Expected sUAS attitude-estimation accuracy with contemporary sensors | 83 |
| 5 | sUAS velocity estimation | 85 |
| 5.1 | Cascaded versus Federated filters for velocity estimation | 87 |
| 5.2 | GPS velocity estimation for sUAS | 90 |
| 5.2.1 | GPS receiver operating principles | 90 |
| 5.2.2 | Testing of U-blox 6T GPS receiver while in accelerating motion | 94 |
| 5.3 | Expected sUAS velocity estimation accuracy with contemporary sensors | 104 |
| 6 | Optical reference vector sensor | 106 |
| 6.1 | Sensor description | 108 |
| 6.1.1 | Principle of operation | 109 |
| 6.1.2 | Link budget | 112 |
| 6.1.3 | Sensor module design | 116 |
| 6.2 | Sensor characterization | 119 |
| 6.2.1 | Prototype system parameters | 119 |
| 6.2.2 | Calibration | 120 |
| 6.2.3 | Noise performance | 121 |
| 6.3 | Potential adaptations | 124 |

| | | |
|----------|---|------------|
| 7 | Flight test analysis | 126 |
| 7.1 | Sensor calibration | 129 |
| 7.1.1 | Gyroscope calibration | 129 |
| 7.1.2 | Magnetometer calibration | 130 |
| 7.1.3 | Accelerometer calibration | 136 |
| 7.1.4 | GPS latency calibration | 137 |
| 7.2 | Effect of platform propulsion system on sensor measurements | 137 |
| 7.2.1 | Vibration and accelerometer measurements | 138 |
| 7.2.2 | Vibration and gyroscope measurements | 142 |
| 7.2.3 | Electrical fields and magnetometer measurements | 147 |
| 7.3 | Sensor fusion using actual flight data | 150 |
| 7.4 | Validation using optical reference vector sensor | 151 |
| 7.4.1 | Filter consistency | 158 |
| 7.5 | Flight test analysis summary | 161 |
| 8 | Summary and conclusions | 163 |
| 8.1 | Summary and conclusions: sUAS attitude estimation | 164 |
| 8.2 | Summary and conclusions: sUAS velocity estimation | 166 |
| 8.3 | Summary and conclusions: Optical reference-vector sensor | 167 |
| 8.4 | Summary and conclusions: field (flight) tests | 168 |
| 8.5 | General conclusions | 170 |

| | | |
|-------|---|------------|
| 8.6 | Future work | 170 |
| | Bibliography | 172 |
| | Appendices | 178 |
| 1 | Attitude representation and attitude integration using quaternions | 179 |
| 2 | Block diagonal structure of the EKF S matrix | 182 |
| 3 | The extended Kalman filter algorithm | 184 |
| 4 | Simulation of sUAS state estimation | 186 |
| 4.1 | Generation of truth data based on an aerodynamic model | 186 |
| 4.2 | Simulation of sensor measurements based on truth data | 187 |
| 4.3 | Simulation of competing sUAS state estimation algorithms | 188 |
| 5 | Optical reference vector sensor schematics, board layouts, and firmware | 190 |
| 5.1 | Firmware | 192 |
| 5.1.1 | main.c file listing | 192 |

Tables

Table

| | | |
|-----|---|----|
| 3.1 | Wind measurement parameters | 33 |
| 3.2 | Induced vertical wind measurement error - nominal case | 34 |
| 3.3 | Induced vertical wind measurement error - worst case | 35 |
| 4.1 | Attitude estimation comparison for different filter types at 3 wind/turbulence levels. | 64 |
| 4.2 | Attitude estimation comparison for different filter structures at 3 wind/turbulence levels. | 71 |
| 4.3 | Attitude estimation comparison for filters with different accelerometer measurement models at 3 wind/turbulence levels. | 80 |

Figures

Figure

| | | |
|-----|---|----|
| 2.1 | Hadley circulation model: Figure 7.5 in The Atmosphere, 8th edition, Lutgens and Tar- | |
| | buck, 8th edition, 2001. | 11 |
| 2.2 | Relative Wind Sensors. Clockwise from upper-left: hot wire anemometer, multi-hole | |
| | probe, sonic anemometer, Pitot probe with vanes | 18 |
| 2.3 | Wind Measurement Geometry. | 18 |
| 4.1 | A Maxim 21100 sensor including a 3-axis gyroscope, 3-axis accelerometer and 3- | |
| | axis magnetometer; this device represents typical physical sizing for contemporary | |
| | MEMS sensors. | 43 |
| 4.2 | Hard and Soft Iron Distortions. | 47 |
| 4.3 | Attitude Estimation Filter feedback diagram. Almost all attitude-estimation filter | |
| | structures conform to this general feedback structure. | 57 |
| 4.4 | 1 degree of freedom platform example system. | 60 |
| 4.5 | 1DOF example - Gain variation as a function of pitch angle. | 62 |
| 4.6 | Attitude Estimation comparison for different filter structures | 65 |
| 4.7 | Federated versus Cascaded filters. | 68 |

| | | |
|------|---|-----|
| 4.8 | Attitude Estimation comparison for comparable federated and cascaded filters. . . . | 72 |
| 4.9 | Error introduced by traditional specific-force correction methods (Equation (4.31)) when yaw angle is perturbed. | 78 |
| 4.10 | Attitude estimation comparison for filters with different accelerometer measurement models. | 80 |
| 4.11 | Errors in estimated translational acceleration. | 81 |
| 4.12 | Error components in estimated translational acceleration using the rotation-rate- based method. | 82 |
| 5.1 | Phase-lock tracking loop models | 94 |
| 5.2 | GPS receiver motion test set-up | 97 |
| 5.3 | Histograms of Lock type versus acceleration and carrier-to-noise ratio | 99 |
| 5.4 | Distributions of GPS velocity-estimate error for data set with 87% phase-lock mea- surements | 100 |
| 5.5 | Distributions of GPS velocity-estimate error for data set with 28% phase-lock mea- surements | 101 |
| 5.6 | Distributions of GPS velocity-estimate error for data set with 99.5% phase-lock mea- surements | 101 |
| 5.7 | GPS channels with PLL phase lock during a sUAS flight - U-blox 6T GPS receiver with an un-amplified patch antenna. | 103 |
| 6.1 | Optical reference vector sensor system concept of operation | 110 |
| 6.2 | Optical reference vector sensor block diagram | 110 |

| | | |
|------|---|-----|
| 6.3 | Front-end noise model | 114 |
| 6.4 | Sensor module front and back | 116 |
| 6.5 | Signals after analog to digital conversion | 118 |
| 6.6 | Sensor calibration setup using a 2-axis CNC positioning system to move beacon relative to the sensor | 122 |
| 6.7 | Pre-calibration error due to non-linearity of the sensor detector | 122 |
| 6.8 | One sigma measurement noise versus power density at sensor | 123 |
| 7.1 | Test flight aircraft | 128 |
| 7.2 | Magnetometer offset calibration results | 133 |
| 7.3 | 12-coefficient Magnetometer calibration results | 135 |
| 7.4 | Accelerometer measurement spectra | 140 |
| 7.5 | X-axis accelerometer measurement spectra | 141 |
| 7.6 | Gyroscope measurement spectra | 144 |
| 7.7 | X-axis Gyroscope measurement spectra | 145 |
| 7.8 | Linear interpolation resampling example frequency response | 148 |
| 7.9 | ORVS measurement versus estimated reference vector based on attitude estimate and GPS position for 5 overhead passes during the March 14, 2015 flight. | 153 |
| 7.10 | Enlarged view of Pass 4 measured and estimated angles from Figure 7.9 | 156 |
| 7.11 | Enlarged view of Pass 4 difference between measured and estimated angles from Figure 7.9 | 157 |
| 7.12 | Magnetometer measurement innovations during flight | 160 |

7.13 Frequency spectra of the x-axis magnetometer measurement innovations during flight 160

Chapter 1

Executive summary and contributions

Movement of air in the atmosphere, commonly termed wind, is a complicated process of interest to numerous fields of research. Wind is a spatial and temporal vector field with features that span length scales from hundreds of kilometers to fractions of a meter and time scales from days to fractions of a second. This poses a fundamental difficulty with measurement of wind; the number of spatial/temporal measurement points that may be desirable for a particular field of study can be very large. Some sensor systems are able to collect measurements with high temporal frequency but little or no spatial diversity. Other systems are able collect measurements with both spatial and temporal diversity, but with limited spatial and/or temporal frequency. Wind measurement from a small unmanned aircraft system (sUAS) is similarly limited, but does provide a sampling regime that is otherwise difficult to duplicate and that is desirable for many fields of study. sUAS-based wind measurement allows for high-accuracy measurements to be made in the lower atmospheric boundary layer (ABL) with high temporal frequency along a flight path. The high temporal measurement rate, combined with low sUAS flight speed, yields high spatial sampling rates along the flight path. Spatial diversity is limited to measurements along a path, but this path may be chosen based on the wind feature under study and swarms of sUAS allow simultaneous sampling along multiple paths. Wind measurement from sUASs is very similar to traditional aircraft-based wind measurement. The significant differences are that sUAS-based wind measurement systems may be significantly less expensive, may be safer allowing operations in environments where operation of

manned aircraft would be wholly impractical, may offer the potential for denser sampling using swarms of sUAS, and may offer higher measurement accuracy and higher spatial measurement frequency due to lower flight speeds.

The focus of this work is to provide insight into the question of **how best to make high-accuracy wind measurements using small, low-cost unmanned aircraft systems**. A particular metric of interest is sUAS-based wind-measurement accuracy of $O(1 \text{ cm/s})$ and this is used to judge desirability of results.

Error contributions to aircraft-based wind measurement include measurement or estimation of the aircraft's velocity, attitude, and attitude rate as well as the error introduced by the sensor measuring the air velocity relative to the aircraft. For contemporary sUAS-based wind measurement systems using low cost sensors, attitude and velocity estimation errors are significant contributors to the wind measurement error. Understanding these error sources and how to minimize their effects drive the first two research questions considered:

- Using contemporary low-cost sensors suitable for sUAS, and considering the motion environment typical of sUAS flying in wind, what sensors and sensor fusion algorithms should be used to estimate the airframe attitude to best support high-accuracy wind measurement?
- Similarly, using contemporary low-cost sensors suitable for sUAS, and considering the motion environment typical of sUAS flying in wind, what sensors and sensor fusion algorithms should be used to estimate the airframe velocity to best support high-accuracy wind measurement?

The motion environment is shown to have a significant impact on the performance of both attitude and velocity estimation in sUAS. Due to their small size and mass, sUAS experience significantly different motion than larger aircraft when flying in wind. This presents a difficulty for validating the performance of sUAS attitude-estimation systems in particular and motivates the third primary research question considered:

- Can a new sensor be constructed with low mass, volume and power consumption, and at low cost, which provides an independent observation measurement of sUAS attitude?

The primary contributions of this work are:

- **An improved understanding of wind measurement from sUAS.** Wind-measurement from an airborne platform bears complexity due to the need for knowledge of the platform state. While sUAS have cost, operational and safety advantages over larger aircraft, high-accuracy state estimation for sUAS poses a problem due to the use of less-accurate, small, low-cost sensors and to the significant wind-gust-induced motion experienced when flying in wind. Understanding the sensitivities of the wind measurement to different error sources in the context of contemporary sUAS allows for development of systems providing higher levels of wind-measurement accuracy and better knowledge of what wind-measurement accuracy is truly available from a system.
- **An improved understanding of error contributors in contemporary sUAS velocity and attitude estimation systems, particularly when operating in a turbulent wind environment.** Contemporary sUAS state-estimation systems rely on MEMS gyroscopes and accelerometers, magnetoresistive magnetometers and global positioning system (GPS) receivers. Numerous approaches have been taken to fusing the information from these sensors to produce state estimates. Analysis herein provides comparative guidance in choosing between sensor fusion approaches by understanding the various ways in which error in the sUAS state estimate arise. The effects of wind-gust-induced motion on sensors and sensor-fusion algorithms have not been previously explored, but are shown to cause significant error in sUAS state estimation when using many state-estimation algorithms.
- **A particular approach to attitude estimation for contemporary sUAS that corrects accelerometer specific force measurements using GPS velocity measurements.** This approach provides a surprising level of performance with contemporary low-

cost sensors in gusty conditions. While this approach is not novel the analysis showing why this method is significantly better than traditional alternatives is new and provides important guidance when designing sUAS-based wind measurement systems as well as sUAS-state-estimation systems for any sUAS that will be flown in wind.

- **Development and demonstration of a new sensor system providing an independent reference vector measurement for use in sUAS state estimation.** Alternatives for an additional sensor for use either to validate a sUAS attitude-estimation system, or to augment a sUAS attitude-estimation system in order to increase performance, are limited and generally unattractive. A new sensor concept has been developed and a prototype sensor system constructed which provides an independent measure of a reference vector direction. This sensor is unaffected by motion, including vibration. The airborne portion of the sensor system is small, lightweight and consumes little power making it suitable for use in sUAS. The ground-based portion of the system is easily transportable and requires no external power or special siting conditions. This sensor system provides new capability for improving accuracy of sUAS attitude-estimation systems and may be adapted to a variety of other uses.

Additional contributions include:

- Development of a simulation framework for evaluating sUAS velocity and attitude sensor-fusion estimation algorithms. This simulation framework includes simulation of wind-gust-induced sUAS motion and sensor error characteristics and allows for comparison of different sensor-fusion algorithms operating on data representative of that seen in a real sUAS flying in wind.
- Commercial off-the-shelf (COTS) GPS receivers offer a high level of performance at low cost, but generally offer few performance guarantees. A particular COTS GPS receiver is evaluated with respect to wind measurement from sUAS and the evaluation results provide

guidance on how to characterize the velocity-measurement accuracy of low-cost COTS GPS receivers operated in sUAS flown in wind.

- Magnetometer error caused by “soft-iron distortion” effects is often presumed to be negligible in sUAS. However flight test data evaluated herein shows that this is not always a valid assumption. A new magnetometer calibration technique is proposed that is suitable for sUAS and that significantly reduces magnetometer measurement error when soft-iron distortion is present.

The contributions described above are motivated by improving the accuracy of sUAS-based wind measurement and understanding the accuracy of sUAS wind measurement systems when operating in gusty conditions. sUAS-based wind measurement systems can provide measurements in sampling areas where it is difficult to place sensors by other means. The ability to make densely-targeted high-accuracy wind measurements in particular volumes will have broad impacts by providing data previously unavailable that will advance numerous areas of research. For example, improvements in general weather modeling improve forecasting future weather conditions. Better weather forecasts have numerous economic and life-safety benefits. Similarly, improvements in wind models provide a better understanding of various transport phenomena, which has numerous benefits to fields such as environmental science, agriculture, and energy production. In addition to impacts on wind measurement and the resulting benefits, the contributions listed above directly contribute to the field of UAS guidance, navigation and control. Many predict UAS to be a disruptive technology and the ability of UAS to operate well in gusty wind conditions will have broad impacts in diverse fields, some of which may not yet be apparent.

Chapters 2 and 3 examine historic and contemporary wind measurement. The relative advantages and disadvantages of sUAS-based wind measurement are considered and several areas of research are identified where sUAS-based wind measurement is particularly suitable. Airborne wind-measurement systems rely on accurate knowledge of the attitude and motion of the airborne

platform. For sUAS outfit with appropriate, contemporary sensors the methods of estimating the platform velocity and attitude are important as error in these estimates are significant error contributors in the wind measurement. Chapter 3 includes a description of the primary research questions considered in this work.

Chapter 4 addresses sUAS attitude estimation. Contemporary sUAS attitude estimation relies on sensors now available with low cost, size, weight and power requirements. While these sensors are very capable, considering their cost and size, they are not nearly as accurate as navigation grade sensors used in larger manned-aircraft systems. Material in chapter 4 examines the characteristics of these sensors and the most appropriate way of combining information from them to form an attitude estimate which supports high-accuracy wind measurement.

Chapter 5 addresses sUAS velocity estimation. The same sensors that are used for attitude estimation are also used for velocity estimation. However, wind-gust-induced motion causes difficulty for velocity estimation in a different manner from that encountered in attitude estimation. In particular sUAS velocity estimation is highly dependent on velocity measurement using small, single-frequency GPS receivers. The applicability of GPS error models to commercial off-the-shelf GPS receivers is ambiguous as manufacturers typically provide limited information about the internal operation of their receivers. Testing of a particular commercial off-the-shelf GPS receiver is performed to infer some characteristics important for velocity estimation during sUAS flight in wind.

The current set of sensors suitable for use in sUAS is sufficient for many purposes, but an additional sensor is desirable to support high-accuracy wind measurement both by allowing for higher-accuracy attitude estimation and by providing a means for validating sUAS attitude-estimation systems in sUAS-in-flight-in-wind conditions. Chapter 6 describes a new sensor system developed for this purpose. This sensor system provides a measurement of the vector direction between the sUAS and a ground-based optical beacon. Characterization of this sensor system shows that it is appropriate both for validating sUAS attitude-estimation systems or for use as an

additional sensor to increase the accuracy of sUAS attitude-estimation systems. The sensor system is also adaptable for a number of different uses.

Chapter 7 describes flight testing conducted to investigate practical aspects of the analytic conclusions reached in chapters 4 and 5, and to demonstrate the use of the sensor system described in chapter 6. Producing and operating a physical system brings to light many problems which are trivial from an analytic perspective but which require due diligence to handle correctly in a physical system. Examples of such issues considered in chapter 7 include calibration of the magnetometer, managing asynchronous discrete time sensor signals, and mitigating the effects of vibration.

Chapter 2

Measurement of wind

Wind has a significant impact on agriculture, transportation, construction, architecture, power production, and many other human endeavors, and has been of interest for centuries. Maury [47] [48] began compiling data and published the first studies of winds over large geographic extents in the mid-nineteenth century. Knowledge of expected wind speeds and directions, even on an average basis, over potential sailing routes greatly improved transit times. The impact of his work on shipping spurred further research into wind and weather patterns.

Prior to the last century all measurement of wind occurred in the lowest extent of the atmospheric boundary layer. Although observation of some atmospheric phenomena provides clues about the three-dimensional character of the wind field, little data about wind direction aloft, and even less data about wind speeds aloft, existed prior to manned flight. The three dimensional spatial structure of the wind field plays an important role and ability to forecast winds based solely on knowledge gained through measurements at the surface is limited. Wind, or more specifically the velocity of the atmospheric constituents, is a very complex phenomena. Study of wind spans length scales from the molecular scale (continuum mechanics representations) to fractions of a meter (turbulent inertial sub-range scale) to the planetary scale (cyclonic scale) and spans time scales from fractions of a second to many days. Measurement of wind from aircraft, and later from remote-sensing systems, allowed for winds to be measured throughout the spatial field, albeit with sparse sampling, and has facilitated current understanding of wind and weather processes.

Historical needs for measurement of wind were primarily driven by weather research and prediction. Contemporary needs for accurate wind measurement are still found in weather research and prediction as well as research, monitoring and prediction in myriad other fields such as agriculture [29], air quality and pollution [14], construction and civil engineering [43], ecology [76], sports [34], wildfire suppression [30], wind energy [28], and many others. These diverse applications for wind measurement have an equally diverse set of requirements. However, in many cases, measurement methodologies are driven by consideration of what instrumentation is practically available based on cost and operational considerations.

The importance of wind stems from its role in various transport phenomena. Pure diffusion occurs at time and length scales that make it too slow to be of interest in most atmospheric processes. Almost all atmospheric transport, whether of mass, momentum, heat, moisture, gaseous concentration, or particulate matter, and even processes as diverse as microbial or insect transport, are driven by advection, that is to say by the wind.

A cursory introduction to wind theory may be developed with simple models. Wind occurs due to differences in air density primarily due to changes in temperature, with resulting changes in pressure. Pressure gradients create motion. The simplest model may be constructed on a planetary scale by considering that solar radiation causes more heating at lower latitudes than at higher latitudes. High pressure, cold air near the poles moves towards the equator along the earth's surface while displaced low pressure air moves in the upper atmosphere from the lower latitudes towards the poles. This model lacks a critical effect; since the earth is rotating Coriolis force causes the north/south circulation described above to curve, when viewed in the earth-fixed frame. This would induce a consistent east/west flow at higher/lower altitudes in the northern hemisphere and the converse in the southern hemisphere. This model still has a major deficiency in that it does not explain normal trade wind patterns. Hadley proposed a new model in 1735 in the **Philosophical Transactions of the Royal Society** that was widely accepted after it was rediscovered by John Dalton and other researchers in the 19th century. Hadley's model has three separate zones, or cells,

in each hemisphere as shown in Figure 2.1. At low latitudes behavior is similar to the model above, but sinking motion in the subtropics divides the lower latitude cell from the middle latitude cell. In the middle latitudes considerable energy transport is accomplished by cyclones and anti-cyclones that move warm air polewards and cold air towards the equator in the same horizontal plane. At high latitudes the overturning mechanism again dominates. Hadley's model was improved by William Ferrel, publishing interestingly in the **Nashville Journal of Medicine and Surgery**, in the mid nineteenth century, and more recently by contemporary researchers such as Isaac Held and Arthur Hou [31]. The Held-Hou model is a conceptual model that describes atmospheric circulation in the absence of turbulence.

Large-scale models like the Held-Hou model do not recognize local forces such as friction and turbulent diffusion, surface interactions such as orographic lifting, moisture processes involving clouds and precipitation, and localized heating and convection. Numerical models can incorporate numerous processes and forces where analytical models become intractable. Models may be constructed on smaller spatial scales from fluid mechanics based dynamical equations that model events like a single thunderstorm induced downburst [53] or a wind turbine wake [56]. However, models that incorporate the fine level of detail needed for study of these small-scale phenomena cannot be practically scaled to synoptic or planetary scales while preserving small-scale details. Larger-scale models must replace processes that are too complex or small scale with parameters representing the action of these small-scale processes within a grid cell in a larger-scale model. This is not only a problem for large-scale models; even small-scale models such as those mentioned above must parameterize processes that cannot be directly resolved on the model grid scale. For example formation of cloud droplets, and their growth into falling rain, occurs at the molecular scale that is much too small to be represented directly in any model on the scale of a whole cloud [71]. The need to understand small-scale atmospheric phenomena for use in parameterizing models is a key driver for measuring wind on fine spatial and temporal scales.

The vertical pressure gradient is much larger than horizontal pressure gradients, but is bal-




Figure 2.1: Hadley circulation model: Figure 7.5 in The Atmosphere, 8th edition, Lutgens and Tarbuck, 8th edition, 2001.

anced by gravity. Gravity acts to stop, or slow, the vertical flow of air, so vertical winds are much less than horizontal winds. Most vertical winds are on the order of 1 mph, however some downdrafts and updrafts can be up to 60 mph [72]. The tendency for vertical stability, with much of the transport phenomena taking place in striated horizontal layers, makes the interface between layers a particularly interesting location for wind modeling and measurement. Parameterization of vertical motion across layer boundaries is critical to accurate modeling. Measurement of the vertical component of wind is particularly challenging, however, as the vertical component is typically of much smaller magnitude than the horizontal components. The importance of measurement of the vertical wind component for understanding many small-scale atmospheric phenomena leads to the partial focus on measurement of this particular wind component herein.

2.1 Contemporary wind measurement systems

Current systems and sensors have a variety of capabilities. Wind measurement systems may be categorized as either ground-based, or airborne or space-based. They may be further categorized based on if they use a sensor making a local or remote measurement. Each system type has advantages and disadvantages, making some systems more favorable than others for particular applications.

2.1.1 Ground-based in-situ

Numerous types of in-situ wind sensors have been developed. Simple methods such as observing flags or smoke provide a certain level of accuracy in measuring wind direction and a qualitative measure of wind speed. Mechanical sensors such as vanes, cup or turbine anemometers, Pitot tubes, venturies, etc., have long been used to measure the horizontal-wind direction and speed. The preferred instrument for 3D wind measurements from towers or other ground-based structures is the sonic anemometer. Current scientific-grade sonic anemometers [10] [35] [75] [49] are com-

mercially available, offering sampling rates up to 1Khz, wind speed component resolution down to 1 mm/s, and wind speed component accuracy down to 4 cm/s (note: these specifications cannot all be achieved by the same instrument). The desirability of the sonic anemometer stems from its ability to make accurate measurements without the need for temporal averaging.

Measurement of the wind speed at any non-trivial height above the surface requires mounting a sensor on a tower or other tall structure, which is the most significant limitation of ground-based in-situ systems. Towers are expensive to erect and only allow measurements in a single column over the ground. Diverse spatial sampling is generally impractical even at the surface and certainly at higher levels when using these systems. Ground-based in-situ systems are able to make persistent measurements. It is possible to collect a data set spanning many years once such a system is in place.

2.1.2 Ground-based remote sensing

Ground-based remote sensing systems (radar wind profilers, sodar (acoustic) wind profilers and lidar (light) wind profilers) rely on the Doppler shift in acoustic, radio or light waves reflected by turbulence induced density gradients or by aerosols. They can make measurements at greater heights than are practical with towers and the measurement point can be steered (most precisely with lidar). However, the underlying measurements are a Doppler shift along a line of sight, which equates to a velocity along the line of sight. To make a 2D or 3D measurement, the sensor must “look” in multiple directions and then, assuming homogeneity over a volume, transform the raw measurements to a vector measurement.

For measurements in the lower boundary layer, lidar is becoming the preferred remote sensing instrument with significant proliferation in commercial units targeted towards the wind energy market. While they cannot match the range of radar wind profilers, they are significantly more portable than either radar wind profilers or sodars. Lidar specifications can vary significantly

depending primarily on their maximum range. SgurrEnergy manufactures two lidar suitable for lower ABL measurements with differing ranges. Their Galion G250 [67] has a range of 250m with a spatial resolution (in the beam direction) of 24m and an accuracy of 0.1m/s. The Galion G4000 has a range of 4km with a spatial resolution (in the beam direction) of 30m and an accuracy of 0.1m/s. Mikkelsen [50] surveyed commercially available lidar wind instruments (including the Galion instruments) and states that raw measurement rates up to 500Hz are available but multiple beams are necessary for vector measurements, and averaging is often used, resulting in longer measurement periods for vector measurements.

2.1.3 Airborne in-situ

A variety of platforms have been used for airborne in-situ wind measurements. Kites allow instruments to be lifted over a fixed location and balloons allow instruments to be carried in a Lagrangian manner. Airships allow for targeted spatial measurements, but can be difficult to operate in any but light winds. Rotorcraft have been used for wind measurements, but special care must be taken so that the induced airflow does not contaminate the measurements. Airplanes have long been used but are too costly and/or have operational constraints making them unattractive for many situations. Unmanned Aircraft Systems (UAS), particularly small UAS (sUAS) are a promising platform for wind measurement as they are less costly to operate than manned aircraft, often are able to operate in airspace too difficult or hazardous for manned aircraft and their lower airspeed allows for the use of instruments with smaller dynamic ranges.

Airborne in-situ wind measurement adds complexity over ground-based in-situ measurement. A sensor is used to measure the wind relative to the airborne platform, but information about the platform attitude and motion is required to produce a wind measurement in the earth-fixed frame. Measurement or estimation of the platform attitude and motion can be a significant limitation to the accuracy of the resulting wind measurement. Airborne in-situ wind measurement may utilize a sensor similar to those used for ground-based in-situ wind measurement, but other sensors not

typically used for ground-based measurements are also frequently used for airborne measurements. Issues particular to airborne wind measurement are examined in Section 2.2.

2.1.4 Airborne and space-borne remote sensing

Airborne remote sensing systems have the ability to make measurements over large geographic areas as the remote sensor (radar, sodar, lidar) can make measurements at multiple locations relative to the platform and the platform can be flown to different locations. This capability comes with the added expense of operating the platform and with the liability of requiring information about the platform attitude and motion to translate the sensor measurements to a measurement of wind in the earth-fixed frame. These systems have the same limitations as ground-based remote sensing systems with respect to making measurements with high spatial resolution.

Space-borne remote sensing systems are similar to airborne remote sensing systems but are more extreme in most regards. They are able to cover extremely large areas but due to the length scales involved vector measurements from space-borne systems necessarily require very large measurement volumes so spatial sampling frequency is very low. Cost to field a space-borne system is extremely high and measurement locations may be constrained by orbit considerations.

2.2 Airborne wind measurement

Wind measurements, as required for atmospheric modeling, are a measure of the velocity of the air relative to the earth-fixed reference frame. This is usually a trivial consideration for ground-based wind-measurement systems that are typically in a fixed location and attitude with respect to the ground. For airborne (or space-borne) systems, however, the ability to make measurements from a moving platform adds the complexity of accounting for the platform's position, attitude and motion in making a measurement relative to the earth-fixed frame.

For coarse measurement of the wind, the wind may be estimated by observation of the effect

of the wind on the dynamical behavior of the aircraft. For example a simple and well known method for measuring the horizontal wind is to fly an airplane at a constant airspeed and bank angle. The wind speed and direction may be calculated by measuring the horizontal displacement between the beginning and end locations of a 360° rotation. This method is even more straightforward when using a balloon. However, the accuracy of this method is limited by how accurately the position of the aircraft may be measured. Also, sufficient time must be allowed for the aircraft to complete a circle, or in the case of a balloon for the horizontal displacement to be sufficiently larger than the position measurement accuracy. The result is that the measurement is an average over a time and length scale rather than a point measurement.

For point measurements a sensor measuring the relative wind (the velocity of the air relative to the airframe) is required. A variety of sensors have been used for this purpose. Cup or turbine anemometers are not often used on aircraft due to their limited dynamic range. A Pitot probe (or Pitot tube) is a one-dimensional measurement of relative wind speed based on the difference between the dynamic stagnation pressure at the tip of a probe and the static pressure (measured along the side of the probe or elsewhere). Systems using a sensor that make a scalar measurement of the relative wind (i.e. an airspeed sensor) are commonly used within aircraft avionics systems to measure wind, but a sensor that can make a vector measurement of the relative wind is required to make vector measurements at discrete points. Fixed wing aircraft require sustained airspeed for flight so the relative wind vector is usually within a small angular range. Although Pitot probes provide a scalar measurement, a combination of a Pitot probe with two directional vanes can provide this vector measurement. Another method of making a vector wind measurement is by measuring the dynamic pressure at a set of points on a blunt body. This is typically done with a multi-hole probe similar to a Pitot probe but with multiple dynamic pressure ports arranged on a blunt tip rather than just a single dynamic pressure port. An alternative to measuring the pressure at points on a probe tip is to measure pressure at points directly on the airframe itself; thin pressure transducers may be used for this purpose. Hot wire anemometers measure

wind velocity by measuring heat transfer from a low-mass heated element. By arranging multiple heating elements so that they experience differing airflow dependent on the incident wind angle a vector wind sensor may be constructed. Sonic anemometers measure the time-of-flight differences of acoustic signals; the effective speed of sound is increased for signals traveling in the direction of the wind and decreased for signals traveling against the direction of the wind. Using signals transmitted in multiple directions a vector measurement may be constructed. Figure 2.2 illustrates several of these sensor types.

In addition to the relative wind measurement, a measurement of wind in the inertial frame from a moving platform requires a measurement/estimate of the sensor's attitude and velocity. Large aircraft may have a very accurate inertial navigation system (INS) providing high accuracy attitude, attitude rate and velocity information. Information about the attitude, attitude rate and velocity of small aircraft may be less well known.

Derivation of an expression of the wind measurement dependencies begins with the geometry shown in Figure 2.3. Here an in-situ measurement from a fixed-wing aircraft is considered, but the analysis may be modified for other airborne wind measurement systems. Consider a relative wind sensor mounted on an aircraft and an a small air parcel being observed by that sensor. The air parcel is at a location \vec{r}_{air} in the inertial frame and at a location \vec{r}_{air}^{sensor} in the sensor frame. In extreme cases the difference between the earth-fixed frame, which is slowly rotating at approximately $7 * 10^{-5}$ (rad/sec), and the inertial frame must be considered, but for most wind measurement systems the earth-fixed frame may be taken as an inertial frame. The aircraft is at a location \vec{r}_{body} in the inertial frame, which is the origin of the body frame and the sensor is at a location \vec{r}_{sensor}^{body} in the body frame, which is the origin of the sensor frame. This geometry leads to the vector relation $\vec{r}_{air} = \vec{r}_{body} + \vec{r}_{sensor}^{body} + \vec{r}_{air}^{sensor}$. The quantity of interest, the wind in the inertial frame, is the time derivative of the air parcel position with respect to the inertial frame:

$${}^I\vec{v}_{air} = I\frac{d}{dt}\vec{r}_{air} = I\frac{d}{dt}\vec{r}_{body} + I\frac{d}{dt}\vec{r}_{sensor}^{body} + I\frac{d}{dt}\vec{r}_{air}^{sensor} \quad (2.1)$$

Here the leading superscripts indicate the derivatives are taken with respect to the inertial frame.



Figure 2.2: Relative Wind Sensors. Clockwise from upper-left: hot wire anemometer, multi-hole probe, sonic anemometer, Pitot probe with vanes




Figure 2.3: Wind Measurement Geometry.

The first term on the right-hand side is the vector velocity of the airframe, ${}^I\vec{v}_{body}$, and the second and third terms on the right-hand side are better understood in the body frame and sensor frame respectively. The second term is the velocity of the sensor relative to the airframe. This term accounts for movement of the sensor relative to the body frame due to flexibility in the airframe, sensor or sensor mounting. This term will be zero when averaged over time. The third term is the velocity of the air relative to the sensor itself. Expanding terms using the expression for the total derivative in a rotating frame:

$${}^I\vec{v}_{air} = {}^I\vec{v}_{body} + {}^B\frac{d}{dt}\vec{r}_{sensor}^{body} + {}^I\vec{\omega}_{body} \times \vec{r}_{sensor}^{body} + {}^S\frac{d}{dt}\vec{r}_{air}^{sensor} + {}^I\vec{\omega}_{sensor} \times \vec{r}_{air}^{sensor} \quad (2.2)$$

Requiring the sensor frame to be defined so that the origin is at the point of measurement, e.g. at the probe tip if the sensor is a multi-hole probe, causes $|\vec{r}_{air}^{sensor}|$ (but not ${}^S\frac{d}{dt}\vec{r}_{air}^{sensor}$) to equal zero allowing for some simplification.

If the potential exists for the sensor to move relative to the aircraft body frame, due to limited rigidity in the airframe, sensor and sensor mounting, the motion of the sensor relative to the airframe (the second term on the right-hand side of Equation (2.2)) must be considered. The motion of the sensor frame relative to the body frame may have random and deterministic components; the particular design and fabrication will determine the relative importance of these components. However, for many situations the dominant motion will be cyclic and result from forcing by the airframe propulsion-system vibration. If the frequency of the cyclic motion is assumed sufficiently high such that it may be effectively filtered then its effect on the wind measurement may be disregarded. If this assumption is made, then ${}^B\frac{d}{dt}\vec{r}_{sensor}^{body} \approx 0$. Parameterizing the remaining terms in (2.2), with the wind measurement parameterized in the inertial frame and for convenience parameterizing ${}^I\vec{v}_{body}$, \vec{r}_{sensor}^{body} and ${}^I\vec{\omega}_{body}$, and ${}^S\frac{d}{dt}\vec{r}_{air}^{sensor}$ in the inertial, body, and sensor frames respectively, leads to:

$$[{}^I\vec{v}_{air}]_I = [{}^I\vec{v}_{body}]_I + R_I^B([{}^I\vec{\omega}_{body}]_B \times [\vec{r}_{sensor}^{body}]_B) + R_I^B R_B^S([{}^S\vec{v}_{air}^{sensor}]_S) \quad (2.3)$$

where ${}^S\vec{v}_{air}^{sensor} = {}^S\frac{d}{dt}\vec{r}_{air}^{sensor}$ is the relative wind velocity and R_I^B and R_B^S are rotation matrices quantifying the rotation between the body and inertial, and sensor and body frames. The nominal mounting orientation of the sensor may typically be chosen so that the sensor frame is translated but not rotated relative to the body frame and again assuming the sensor motion relative to the body is limited to high frequency components that may be removed with filtering a final approximation of $R_B^S \approx I$ yields

$$[{}^I\vec{v}_{air}]_I = [{}^I\vec{v}_{body}]_I + R_I^B([{}^I\vec{\omega}_{body}]_B \times [\vec{r}_{sensor}^{body}]_B + [{}^S\vec{v}_{air}^{sensor}]_S) \quad (2.4)$$

The primary assumption made in this development was that the sensor position and orientation is fixed in the body frame. Constructing a system so that this is a valid approximation is perhaps the simplest approach. Two other approaches may also be considered.

- The sensor itself may be instrumented (rather than the airframe) such that the velocity and orientation of the sensor relative to the inertial frame is known. In this case Equation (2.4) simplifies to $[{}^I\vec{v}_{air}]_I = [{}^I\vec{v}_{sensor}]_I + R_I^S[{}^S\vec{v}_{air}^{sensor}]_S$. This approach is currently mechanically impractical for many relative wind sensors.
- The motion of the sensor may be measured relative to the body. In this case Equation 2.3 with additional terms for ${}^I\vec{\omega}_{sensor} \times \vec{r}_{air}^{sensor}$ will be valid. This approach is generally less practical than the first of these two.

Considering Equation (2.4) as a set of scalar equations, and recognizing that a rotation may be minimally represented with a set of three scalar values (Euler angles for example), the right side of the equation shows that wind measurement from an airborne platform requires the measurement or estimation of fifteen scalar values. These fifteen scalar values represent the platform orientation, the platform orientation rate vector, the sensor frame translation from the body frame, and the relative wind vector. The sensitivity of the wind measurement to each of these lower level measurements/estimates may be found by taking the partial derivative of Equation (2.4) with

respect to each of the components on the right hand side. This may produce more insight when handled as separate scalar cases for the three inertial frame components. As an example the vertical wind component is considered below.

A relative wind sensor may produce a vector measurement parameterized in a variety of ways. A useful way to consider this measurement is as a magnitude, V_a (airspeed), and α_{relw} and β_{relw} angles (traditional aerospace angle of attack and sideslip angles), where β_{relw} is the angle between the relative wind vector and the sensor XZ-plane and α_{relw} is the angle between the relative wind vector and its projection on the sensor XY-plane. The relative wind vector parameterized in the sensor frame is

$$[{}^S \vec{v}_{air}^{sensor}]_S = \begin{bmatrix} -V_a \cos \alpha_{relw} \cos \beta_{relw} \\ -V_a \sin \beta_{relw} \\ V_a \sin \alpha_{relw} \cos \beta_{relw} \end{bmatrix}. \quad (2.5)$$

The orientation of the airframe is perhaps most intuitively parameterized as roll, ψ , pitch, θ and yaw, ϕ , (Euler) angles. To reduce the number of parameter dependancies, with no significant consequence to the loss of generality, consider the case where the sensor is mounted on the body frame X-axis so that

$$\vec{r}_{sensor}^{body} = r_{sensor}^{body} {}^S \mathbf{e}_x, \quad (2.6)$$

where ${}^S \mathbf{e}_x$ is the sensor frame X-axis unit vector. Using Equations (2.4) and (2.5), the vertical component of the wind measurement may be expressed as

$$\begin{aligned} wind_z = & V_a \sin \theta \cos \beta_{relw} \cos \alpha_{relw} - V_a \sin \phi \cos \theta \sin \beta_{relw} \cos \alpha_{relw} \\ & + V_a \cos \phi \cos \theta \sin \alpha_{relw} + {}^I v_{bodyz} - r_{sensor}^{body} \sin \phi \cos \theta {}^I \omega_{bodyz} + r_{sensor}^{body} \cos \phi \cos \theta {}^I \omega_{bodyy} \end{aligned} \quad (2.7)$$

Now the sensitivity of the vertical wind measurement to the measurement/estimate of pitch angle, for example, may be determined as

$$\begin{aligned} \frac{\partial wind_z}{\partial \theta} = & V_a \cos \theta \cos \beta_{relw} \cos \alpha_{relw} + V_a \sin \phi \sin \theta \sin \beta_{relw} \cos \alpha_{relw} \\ & - V_a \cos \phi \sin \theta \sin \alpha_{relw} + r_{sensor}^{body} \sin \phi \sin \theta {}^I \omega_{bodyz} - r_{sensor}^{body} \cos \phi \sin \theta {}^I \omega_{bodyy} \end{aligned} \quad (2.8)$$

The instantaneous sensitivity of the wind measurement to the pitch angle is dependent on airspeed, relative wind angles, pitch and roll angles, pitch rate, and sensor offset. This sensitivity may be evaluated for a nominal case based on nominal values, or may be evaluated as a worst case using ranges for these values. In this particular case it can be seen that if the airframe maintains nearly level flight so that the roll and pitch angles remains near zero and the relative wind angles are small then the sensitivity is dominated by the airspeed

$$\left(\frac{\partial wind_z}{\partial \theta}\right)_{\phi, \theta \approx 0} \approx V_a - r_{sensor}^{body} \theta^I \omega_{body y} \quad (2.9)$$

The sensitivity of the vertical wind measurement to the pitch angle measurement/estimate may be now be easily calculated for the nominal case. However this simplification should not necessarily be taken to diminish the significance of the terms that have been dropped as they may be important in a worst case analysis.

This method of examining the partial derivatives of the wind measurement components with respect to the fifteen scalar terms making up the right-hand side of Equation (2.4) is used in section 3.3.1 to examine the significant error contributors to sUAS-based wind measurement.

Chapter 3

Measurement of wind from small unmanned aircraft

In-situ, remote, ground-based, airborne and space-borne wind measurement systems all have advantages and disadvantages. Airborne in-situ systems fill a unique roll in that they can make measurements with high spatial resolution and with diverse sampling in a three-dimensional space. The two most common airborne platforms have historically been balloons and manned aircraft. Balloons are relatively inexpensive to deploy and are embedded in the weather forecasting infrastructure as they are used to collect temperature and humidity profiles from numerous locations on a daily, or more frequent, basis. However balloons have a limited sampling regime of either a Lagrangian ascent or a tethered ascent/descent profile, both of which are fixed to a single point on the ground and are dependent on the wind itself. Manned aircraft provide much greater flexibility in sampling locations, but have disadvantages including:

- They are relatively expensive to acquire and operate,
- They may require expensive ferry operations to reach the desired sampling area,
- They have significant safety and operational limitations with regards to operating close to people or structures, or operating in hazardous environments such as in canyons or over open water,
- fixed-wing aircraft must maintain a minimum airspeed that may impare high-accuracy or high-spatial-frequency sampling and significant measures must be taken to keep rotary-wing

aircraft airflow characteristics from disturbing wind measurement, and

- Simultaneous sampling from multiple aircraft adds considerable expense and safety concerns.

Small unmanned aircraft systems are an emerging technology offering significant cost and operational advantages over manned aircraft and sUAS-based wind-measurement systems offer the potential for much better spatial sampling than ground-based systems while offering lower cost and advantageous operational characteristics, and potentially higher accuracy, than manned-aircraft-based systems.

3.1 Example areas of study where sUAS may have particularly value

Measurement of the wind is inherently difficult as wind is a vector quantity and making measurements with full spatial and temporal coverage over a region of interest is generally impractical. In some limited cases, such as using particle image velocimetry in a wind tunnel, a full wind field may be measured. However, there are no techniques allowing this sort of full field measurement to be made on the scales required for measurement of atmospheric dynamics. Researchers have long been accustomed to getting by with very sparse spatial coverage in wind measurement data. Low cost, quick deployability, low safety risk due to low mass, low flight speeds yielding high spatial sampling densities, and the potential for dense sampling using multiple sUAS all allow sUAS to provide more flexibility in where and how measurements are collected in numerous research situations. The following sections highlight several areas of research where sUAS may be particularly helpful.

3.1.1 Study of the structure of wind turbine wakes

Current research into wind turbine wake structure at the National Wind Technology Center in Colorado could be significantly aided by the ability to make in-situ measurements with sUAS (J. Lundquist, personal interview, February 20, 2014). Lidar has been used for spatial sampling in turbine wakes due to the impracticality of erecting additional towers, but with significant drawbacks. In particular, it is difficult to study the three dimensional structure of a turbine wake since lidar can only make a vector measurement representing an average over a large disk-shaped volume.

Study of wind turbine wakes offers an ideal example of the potential utility of sUAS-based wind measurement systems to circumvent some limitations of manned aircraft. Flying a manned aircraft in close proximity to a turbine poses an obvious hazard, and there is a limit to how close samples might be taken to the turbine itself. A sUAS, however, may be suitably designed so that an inadvertent collision between the unmanned aircraft and the turbine would have no more damage probability for the turbine than a bird strike, a naturally occurring event considered in turbine design. This is not to imply that a collision need be likely. On the contrary it is easily conceivable that a sUAS system might be suitably designed so that with little risk it could use information from the turbine on the blade phase and make sampling passes while actually flying through the blade disk itself.

The accuracy of wind measurements provided by sonic anemometers is sufficient for most turbine wake research purposes and high temporal sampling rates (100-1000Hz) are highly desirable as this allows observation of inertial-subrange turbulent kinetic-energy dissipation. For turbine-wake research, unlike some other fields of research, there is no preferential direction (u, v, w) with respect to accuracy of the wind measurement. However, the ability to take closely spaced samples, particularly in the direction across a wake and with varying horizontal positions and altitudes would be of significant benefit to turbine wake research and a sUAS-based system would clearly be suitable for making such measurements.

3.1.2 Study of vertical momentum flux and coherent structures in the lower atmospheric boundary layer.

Researchers have recently proposed that compact coherent structures play an important role in large-eddy behavior. Foster et al. [22] present evidence in support of the view that coherent structures in the form of transient alternating bands of relatively higher and lower stream-wise velocity, called streaks, facilitate the formation of the primary coherent structures of smaller scale, called ejections and sweeps, that maintain the surface stress. They use a large-eddy simulation with a 3km long, 1km wide and 750m high domain and 6.25 meter cubic grid spacing to support their assertion.

This grid spacing is smaller than the spatial resolution generally achieved with lidar, and the study of vertical momentum flux and coherent structures offers another example of research that can benefit from sUAS based wind measurement. sUAS instrumentation may be constructed to provide all information necessary to produce wind measurements at a rate on the order of 100Hz, and sUAS flight speeds are typically in the range of 10 to 30 m/s . This allows for spatial sampling at much higher frequency than needed for the grid used by Foster. In addition sUAS are able to be re-tasked in real time to follow an area of interest as it drifts with the mean wind field whereas a ground based system has very limited flexibility to move on demand.

Lidar systems, such as the high-resolution Doppler lidar (HRDL) [84], have been the primary measurement tool used by researchers investigating small scale coherent structures [16] [4]. The HRDL lidar system has a velocity precision of 0.1 m/s but current simulation results for coherent structure mean velocities have standard deviations on the order of 0.003 m/s [22]. Both the limited spatial resolution and the greater than one order-of-magnitude difference between the precision present in simulation models and the precision available for experimental observation demonstrates the need for a better wind measurement to support this area of study.

3.1.3 Validation of the frozen turbulence hypothesis for eddy-covariance techniques for vertical-flux measurements.

The eddy-covariance method allows measurement of turbulent transport (vertical flux) of atmospheric constituents by correlating concentration fluctuations with vertical wind speeds. Application of the eddy-covariance method typically depends on use of the frozen-turbulence hypothesis so that ground-based in-situ temporally-measured data can be converted to represent spatial distribution. For example, Buzorius et al. [9] studied vertical aerosol-particle fluxes measured by eddy-covariance techniques above a pine forest canopy using data from fixed towers. Sonic anemometers and particle-concentration instruments on towers were used to measure eddy covariance, but this measurement technique relies on the assumption that time-lagged measurements at a single point may be related through the mean advecting velocity to be an accurate measurement of spatially-separated points. Hill [33] notes that the frozen-turbulence hypotheses provides only a lowest-order approximation of an infinite series of even higher-order approximations, and provides corrections based on fluctuations of the advecting velocity. When an eddy-covariance study is to be performed in the field the ability to gather data and compare the characteristics of time-lagged data from a single point with spatially-separated data (at a single time) can provide confidence in the validity of using data from ground based measurements collected over long time scales.

Operation of manned aircraft at low altitude over a forest canopy is undesirable due to safety issues; however, operation of sUAS in the same environment may be conducted with far less risk and would allow spatially-distributed measurements of winds. Such a measurement set could validate the frozen-turbulence hypothesis and thereby justify the use of tower mounted sonic anemometers for a larger campaign. The Gill Instruments R3-50 Research Anemometer, marketed as being ideal for eddy-covariance studies, has a resolution of 0.01 m/s and 100 Hz sampling rate suggesting that a sUAS wind measurement system for validation of the frozen-turbulence hypothesis have similar capability.

3.1.4 Study of the urban boundary layer

A developing interest in knowledge of the wind field in and over urban environments is driven by applications including air quality, power generation and building design. Instrumented masts and roof-mounted equipment provide wind profile data at and below the rooftop level but it is challenging to obtain observations in many locations in the urban boundary layer. A sUAS could be used to collect such data and may be particularly useful in the roughness sublayer (RSL) where relatively little is known about the flow and turbulent structure over real urban or suburban surfaces. [61]. Christen et al. [12] conclude that coherent structures (sweeps and ejections) are important in momentum and sensible heat transport in the RSL. Their conclusions are based on observations made with mast-mounted 3D sonic anemometers (0.01m/s resolution) and they acknowledge that the lack of true spatial information leaves the driving processes a matter of speculation.

A sUAS-based wind measuring system with comparable accuracy to the sonic anemometers would allow focused spatial sampling and give heretofore unavailable information about coherent structures in the urban RLS. Here, again, sUAS have a significant operational advantage over manned aircraft from both cost and safety perspectives. sUAS guidance, navigation and control technology is rapidly evolving with abilities to operate safely in close proximity to buildings, etc., improving such that sampling scheme's previously impractical will soon be possible.

3.1.5 Observing entrainment processes

Martin et al. [46] assessed the feasibility of using sUAS to observe entrainment processes at the top of the convective boundary layer. They found that small-scale processes and sharp gradients were captured by sUAS measurements that were not detected by met-towers, radio soundings or tethered balloons.

According to Martin et al.:

During daytime over land, radiative heating of the surface drives intense thermal turbulence in the atmospheric boundary layer (ABL). This leads to the formation of a convectively driven mixed layer that is usually separated from the stably stratified free atmosphere by a temperature inversion (Stull 1988). This capping inversion restrains the domain of intense surface-generated turbulence and acts as a lid to the vertical exchange (mixing) of trace constituents as water vapour and pollutants (Stull 2000). The most energetic convective updrafts might be able to penetrate into this capping inversion, thereby initiating the downward transport of warmer, drier, and less polluted air from the free atmosphere into the ABL; this process is termed entrainment. Entrainment thus is a local phenomenon of discrete events that finally supports ABL growth (e.g., Bange et al. 2007).

Capturing entrainment phenomena is significantly easier with aircraft than with surface-based sensors; the aircraft may be directed to the altitude of the entrainment zone and, once there, may take spatially-diverse samples with high resolution. Traumnner et al. [78] posit that while in-situ (manned) aircraft measurements have heretofore been used to estimate entrainment flux, they do not allow fine measurement of boundary layer growth and Doppler lidar better satisfies measurement needs for studying entrainment processes. However, if sUAS-based wind measurement is improved such that measurement resolution/accuracy comparable to lidar is available, then sUAS systems would satisfy the need for high-quality data while allowing targeted measurements with greater spatial resolution and spatial diversity providing better opportunity to capture entrainment phenomena.

3.1.6 Desirable system characteristics for sUAS-based wind measurement

The use cases examined above provide some information on what should be considered desirable characteristics for a contemporary sUAS-based wind-measurement system. There is general agreement that for ground based in-situ measurements the sonic anemometer is the sensor of choice for scientific measurement. Sonic anemometers are available with absolute accuracy on the order of 0.1 m/s and resolution on the order of 0.01 m/s. Even higher levels of accuracy may be desirable particularly for the study of small scale phenomena as simulation results have variance levels below 0.01 m/s. Based on these considerations a target level for sUAS-based wind measurement of 0.01

m/s for each axis is used herein.

Discussions with several atmospheric researchers (Lundquist - National Wind Technology Center, Casanno - Cooperative Institute for Research in Environmental Sciences, Balsley - Cooperative Institute for Research in Environmental Sciences) on the topic of desired measurement rate indicate that spatial and temporal sampling rates that approach the lower limit of the turbulent inertial subrange are often desirable. The inertial subrange in the atmosphere is usually expected to extend down to the centimeter scale. For a sUAS flying at 10 m/s sampling at 1 cm intervals requires a measurement rate of 1KHz, which may be considered a target for sUAS-based wind measurement frequency.

3.2 Overview of existing sUAS-based wind measurement systems

While the idea of using fixed-wing (or rotary-wing) unmanned aircraft for wind measurement is not new, previous implementations have been limited. Recent advances in sensors and embedded processors are making sophisticated sUAS-based wind-measurement systems possible, but this is a relatively recent development. Rowland [62] published one of the first uses of a small fixed-wing unmanned aircraft for wind measurement in 1972. Rowland used a sUAS for making repeated measurements in a vertical column using a 5.2 kg, 2.5 meter wingspan airframe and a heated thermistor sensor to measure airspeed. Rough estimates of the vertical-wind-component were made based on airspeed and descent rate while the airframe was flown as a glider with the engine off. Data was telemetered to the ground and recorded on magnetic tape in analog form. In just over forty years sensors and embedded processors available off-the-shelf at low cost make such a system look very antiquated.

Elston et al. [17] provide a survey of sUAS-based wind-measurement systems. This survey provides an overview of different wind measurement methods, sensors, and airframes used for a variety of applications. The literature indicates that some UAS-based wind-measurement systems

have had accuracy on the order of 0.01 m/s [40] but it is unclear that any UAS small enough to be considered a sUAS have met this level of accuracy yet.

Airframe production may initially appear to be the the most significant hurdle to constructing a sUAS-based wind-measurement system. However, small airframe design has benefitted from a long history of model-aircraft construction and an active hobbyist community driving development of a wide variety of applicable off-the-shelf components. Still, there are a variety of criteria that must be considered when choosing an airframe such as the required flight time, altitude gain and flight speed, use of a pusher versus tractor configuration and accompanying sensor mounting issues, size, mass, propulsion type (electric versus combustion), launch type (hand launch, bungee, catapult, rolling takeoff) etc. There is no single airframe design that is most suitable for wind measurement. Much depends on the particular application. In general, though, it appears that researchers have not had particular difficulty in acquiring or producing suitable airframes for wind measurement.

Existing systems have used a variety of wind estimation methods, with some methods providing only average background wind estimates, some providing vector wind measurements with high spatial and temporal resolution and others with intermediate capabilities. Only systems producing vector wind measurements with high spatial and temporal resolution will be considered herein as these systems are most suitable for scientific data collection and such systems may now be produced with small volume, mass, and power requirements and at low cost.

As discussed in Section 2.2 a single point vector wind measurement from an airborne platform requires a sensor providing a vector measurement of the relative wind. Frequent use has been made of multi-hole probes (e.g., [39], [81]). At present multi-hole probes appear to be the best relative-wind sensor available for sUAS. Flush pressure sensors, which measure pressure at different points on an airframe and thereby make a relative-wind measurement in the same manner as a multi-hole probe have recently been used by Quindlen and Langelaan [57]. This type of sensor removes issues associated with sensor motion relative to the airframe, but is difficult to calibrate. Vane type systems are rarely used on sUAS due to their fragile nature when fabricated in small sizes.

Multi-wire hot-wire anemometers are difficult to fabricate and calibrate, and are quite fragile. Sonic anemometers have been used for wind measurement from unmanned aircraft, but have not been small or light enough for use in sUAS.

Single point vector wind measurement from an airborne platform also requires measurement or estimation of the airframe velocity, attitude and attitude-rate, which is referred to as state estimation. With a few exceptions such as the DataHawk [42] (which uses thermopiles, a magnetometer and GPS receiver), almost all sUAS use a sensor fusion scheme to fuse information from an inertial measurement unit, a magnetometer, a GPS receiver, and potentially other sensors. Additionally, the majority of systems either used the sUAS autopilot's state estimate or use an off-the-shelf system such as a GPS-INS system. Of those researchers who disclose their sensor fusion scheme the three types noted by Elston et al. include complimentary filters and extended Kalman filters. Desirable methods of performing this sensor fusion are examined in Chapters 4 and 5.

3.3 Accuracy of wind measurement from sUAS

3.3.1 Sensitivity of the inertial-frame wind measurement to constituent measurements

To understand the sensitivity of the wind measurement to the relative wind-measurement and to the different components of the airframe state estimation, nominal and worst case sensitivity based on the partial derivatives of Equation (2.4) are analyzed. As measurement of the vertical wind is a particularly interesting case the analysis will examine that component of the wind measurement using the partial derivatives of Equation (2.7). To begin, expected nominal values and ranges of various parameters must be established. A nominal airspeed of 12 m/s is chosen, which is typical for sUAS in the two to three kilogram range, along with a sensor offset of 1 meter which is as large as would be encounter in sUAS. Wind measurement will be expected to occur during level flight, with pitch and roll angle ranges established to allow for turbulence-induced motion. sUAS airframes

have relatively low inertia and are expected to weathervane quickly around the longitudinal and yaw axes, so modest ranges are established for the relative wind angles, along with ranges for pitch rate and yaw rate. These ranges, centered on the nominal values, appear in the “Expected range” column of Table 3.1. Given the nominal and extreme values of these parameters the sensitivity of the wind measurement to each parameter may be calculated. The error in the wind measurement (to first order) is the cumulative result of the sensitivity of the measurement to each parameter times the error in each parameter value. Given an error in the measurement or estimation of a particular parameter the resulting error in the wind measurement may be calculated based on nominal or extreme values of the other parameters. The column of Table 3.1 labeled “Measurement accuracy” represents the best expected accuracy of that parameter’s measurement or estimation in the context of contemporary sUAS systems.

Table 3.1: Wind measurement parameters

| Parameter name | Symbol | Expected range | Measurement accuracy |
|--------------------------------|---------------------|-------------------|----------------------|
| Relative wind magnitude | V_a | 12 m/s | 0.1 m/s |
| Relative wind vertical angle | α_{relw} | $\pm 5^\circ$ | $\pm 0.1^\circ$ |
| Relative wind horizontal angle | β_{relw} | $\pm 5^\circ$ | $\pm 0.1^\circ$ |
| Airframe pitch angle | θ | $\pm 10^\circ$ | $\pm 0.1^\circ$ |
| Airframe roll angle | ϕ | $\pm 20^\circ$ | $\pm 0.1^\circ$ |
| Airframe vertical velocity | v_{bodyz} | n/a | 0.01 m/s |
| Sensor offset | r_{sensor}^{body} | 1 m | 0.001 m |
| Pitch Rate | $I\omega_{bodyy}$ | $\pm 250^\circ/s$ | $0.1^\circ/s$ |
| Yaw Rate | $I\omega_{bodyz}$ | $\pm 250^\circ/s$ | $0.1^\circ/s$ |

Table 3.2 contains the results of a nominal case sensitivity analysis. Using the nominal values from Table 3.1 the partial derivatives of Equation (2.7) yield the sensitivities shown in the Table 3.2 column labeled “Sensitivity”. Multiplying this sensitivity by the measurement accuracy from

Table 3.1 produces the error induced in the wind measurement by the measurement error for each individual parameter when the other parameters are nominal. For the nominal values chosen it can be seen that the vertical-wind measurement is sensitive to measurement of the relative-wind vertical angle, the airframe pitch angle and the airframe vertical velocity.

Table 3.2: Induced vertical wind measurement error - nominal case

| Parameter name | Symbol | Sensitivity | Induced error |
|--------------------------------|---------------------|---------------|---------------|
| Relative wind magnitude | V_a | 0 m/s / m/s | 0 m/s |
| Relative wind vertical angle | α_{relw} | 12 m/s / rad | 0.021 m/s |
| Relative wind horizontal angle | β_{relw} | 0 m/s / rad | 0 m/s |
| Airframe pitch angle | θ | 12 m/s / rad | 0.021 m/s |
| Airframe roll angle | ϕ | 0 m/s / rad | 0 m/s |
| Airframe vertical velocity | v_{bodyz} | 1 m/s / m/s | 0.01 m/s |
| Sensor offset | r_{sensor}^{body} | 0 m/s / m | 0 m/s |
| Pitch Rate | $I\omega_{bodyy}$ | 1 m/s / rad/s | .002 m/s |
| Yaw Rate | $I\omega_{bodyz}$ | 0 m/s / rad/s | 0 m/s |

Similarly, Table 3.3 contains the results of a worst case case sensitivity analysis. The partial derivatives of Equation (2.7) are evaluated using the worst case set of values from Table 3.1 to yield the sensitivities shown in the Table 3.3 column labeled “Sensitivity”. Multiplying this sensitivity by the measurement accuracy from Table 3.1 produces the error induced in the wind measurement by the measurement error for each individual parameter when the other parameters are at their value that maximizes the sensitivity. For the chosen range of parameter values it can be seen that the vertical-wind measurement is again sensitive to measurement of the relative-wind vertical angle, the airframe pitch angle and the airframe vertical velocity. However whereas the measurement accuracy of the relative-wind magnitude (airspeed) had no impact on the vertical wind measurement in the nominal case, it can be seen that airspeed measurement accuracy is also important.

Table 3.3: Induced vertical wind measurement error - worst case

| Parameter name | Symbol | Sensitivity | Induced error |
|--------------------------------|---------------------|-------------------|---------------|
| Relative wind magnitude | V_a | 0.14 m/s / m/s | 0.014 m/s |
| Relative wind vertical angle | α_{relw} | 11.1 m/s / rad | 0.019 m/s |
| Relative wind horizontal angle | β_{relw} | -4.2 m/s / rad | 0.007 m/s |
| Airframe pitch angle | θ | 12.0 m/s / rad | 0.021 m/s |
| Airframe roll angle | ϕ | 4.2 m/s / rad | 0.007 m/s |
| Airframe vertical velocity | v_{body_z} | 1 m/s / m/s | 0.01 m/s |
| Sensor offset | r_{sensor}^{body} | -2.6 m/s / m | 0.003 m/s |
| Pitch Rate | $I\omega_{body_y}$ | -0.94 m/s / rad/s | 0.002 m/s |
| Yaw Rate | $I\omega_{body_z}$ | 0.34 m/s / rad/s | 6 e-4 m/s |

The results in Tables 3.2 and 3.3 indicate that vertical wind measurement accuracy on the order of 1 cm/sec requires a relative wind sensor capable of measuring airspeed with an accuracy on the order of 0.1 m/s and wind angles with an accuracy on the order of 0.1° , requires a pitch attitude estimate with an accuracy on the order of 0.1° , and requires an airframe vertical-velocity measurement with an accuracy on the order of 0.01 m/s. There is no single underlying measurement that causes the dominant error in the wind measurement.

The underlying measurements may be divided into two categories: those related to the relative wind sensor and those related to estimation of the airframe state. Estimation of the airframe state is the focus of the research described herein. This group contains nine of the fifteen (minimum representation) parameters in Equation (2.4): the velocity, attitude, and attitude-rate of the airframe with respect to the inertial frame. In contemporary sUAS state-estimation systems measurement/estimation of attitude and attitude-rate are closely related, with estimation of attitude being the more complex problem. Desire to understand the problem of sUAS state estimation with respect to wind measurement leads to the questions presented in Section 3.4

3.3.2 Optimal framework for inertial-wind measurement

Measurement of wind in the inertial frame is sensitive to error in any of the terms in Equation (2.4), so a natural question is how to produce these constituent measurements in a manner that will minimize the resulting error. Obviously if the error in each term in Equation (2.4) may be minimized independent of an effect on the other terms this will minimize the resultant error. If the measurements are not independent, then they should be produced in a manner that preferentially reduces the errors to which the inertial wind measurement is most sensitive. Still, if a measurement is independent of other measurements it should be produced in the manner that minimizes its error without consideration of the other measurements.

A relative-vector-wind sensor may combine pressures or Doppler measurements in some manner to produce a vector measurement. A relevant question is “How should information from other sensors be combined with these underlying measurements to arrive at an estimate the inertial wind-measurement with the minimum error?” The crux of this question is whether the relative-wind sensor measurements from the current time (only) should be combined to arrive at a relative-wind measurement that is then combined with a current attitude and velocity estimate to form an inertial-wind measurement or whether all sensor measurements from the current and preceding times should be fused directly in some fashion. The inertial wind, and consequently the relative wind, is a stochastic process and individual time samples must be measured independently. To combine relative wind sensor measurements from multiple times in some way to produce an estimate would require presupposition of some form of structure of the wind. For purposes of measurement of the wind this would be highly undesirable. This differs from the platform state, which is highly dependent on the platform state at previous times and where measurements from previous times are used in forming an estimate of the platform state at the current time. For these reasons it is undesirable, and in fact would be suboptimal because of the stochastic nature of the wind, to combine the underlying measurements of the relative-wind measurement in any manner other than that which uses the current measurement only and which minimizes the error in the relative-wind

measurement. This relative-wind measurement is then combined with the current attitude and velocity estimate geometrically to produce an inertial-wind estimate.

Although measurement of the sensor offset, $[\vec{r}_{sensor}^{body}]_B$, is typically a static, external measurement a similar argument applies.

The question of optimal estimation of the platform velocity, platform attitude, and platform attitude-rate is considered in section 4.2.3.

3.4 Research questions

The research presented herein considers three questions that are central to understanding issues accompanying wind measurement from sUAS and are focused on estimating the state of the airframe.

First, wind measurement in the inertial frame is sensitive to the estimation of the airframe attitude. sUAS have size, weight, power and cost limitation for avionics and payload systems, so navigation-grade attitude and heading reference systems, inertial navigation systems, or other systems that might be used in manned aircraft to provide a high-quality attitude estimate cannot be used in sUAS. This leads to the first question;

- An inexpensive sensor set consisting of a 3 axis MEMs gyroscope, a 3 axis MEMs accelerometer, a 3 axis magnetoresistive magnetometer and a commercial off-the-shelf single frequency GPS receiver module has become the typical basis for sUAS attitude estimation. Based on analysis of contemporary sensors composing this sensor set, analysis of sensor fusion algorithms, and consideration of typical sUAS flight characteristics and contemporary microprocessor capabilities, what approaches to sensor fusion yields good results for sUAS attitude estimation with respect to flight in wind? How well do they work and what is the impact of their performance on sUAS-based wind measurement?

The wind measurement is also sensitive to the estimation of the airframe's inertial velocity, but the same sUAS limitations that affect attitude estimation also affect velocity estimation. This leads to the second question;

- The sensor set above is also the typical basis for sUAS velocity estimation. Based on analysis of contemporary sensors composing this sensor set, analysis of sensor fusion algorithms, and consideration of typical sUAS flight characteristics and contemporary microprocessor capabilities, what approaches to sensor fusion yields good results for sUAS velocity estimation with respect to flight in wind? How well do they work and what is the impact of their performance on sUAS-based wind measurement?

Other sensors are sometimes used in sUAS attitude estimation but suitable, available sensors either do not have direct observability of a reference vector or have poor accuracy. This leads to the third research question:

- An additional reference-vector sensor is useful for either validating or augmenting a sUAS attitude-estimation system, but available sensors suitable for sUAS use have significant limitations. Optical systems have been suggested but have generally either involved lasers and precision tracking of the airframe for outdoor use or have used pixel-based imaging of large fields of view yielding poor accuracy. What approach to an optical reference-vector sensor is suitable for use in sUAS and yields good results when used at useful (flight) ranges in an outdoor, daylight environment? How well will such a system perform and how can it be used to validate or augment a sUAS attitude estimation system?

Chapter 4

sUAS attitude estimation

The attitude of an aircraft, which is characterized by the rotation of the aircraft body frame relative to an inertial frame, is relatively difficult to determine. This rotation may be described by an axis of rotation in the inertial frame and an amount (angle) of rotation around that axis. Numerous alternate methods of parameterizing this rotation have been used in aerospace systems including Euler angles, rotation matrices, quaternions, Rodriguez parameters, etc. Quaternions are the primary attitude representation used in this chapter, with rotation matrices and Euler angles also appearing. Readers unfamiliar with quaternion attitude representations or quaternion algebra are referred to Appendix 1.

Two basic methods are available for determining attitude. One method is to directly measure the rotation of one frame/body with respect to another. This is the basis of most early attitude determination systems in manned aircraft. A spinning gyroscope will resist a change to its axis of rotation, so if fitted into a gimbal such that it is free to rotate it will maintain its rotational axis while the platform to which it is attached rotates. By fitting two gyroscopes with non-colinear axes into a gimbal, instrumenting the gimbal to measure the rotation of the gyroscopes relative to the platform, and starting the system in a known orientation it is possible to measure the orientation of the platform relative to its initial orientation in the inertial frame. Essentially the gyroscopes maintain a fixed orientation in the inertial frame and the gimbal measures the rotation between the platform and the gyroscopes. The contemporary analogue to this method is examined below.

The second method relies on the observation of two reference vectors in both the inertial frame and the airframe body frame to determine the airframe orientation (Wahba's problem) [82]. If a reference vector is fixed in one frame, then no measurement of that vector in that frame is required. For example the gravity vector is always in the local downward direction and for sUAS systems which travel limited distances and do not have very-high-accuracy attitude estimates ($O(1$ arc-minute) or better), the gravity vector may be considered fixed in the downward direction in either a North/East/Down (NED) or East/North/Up (ENU) inertial frame. So, classically, one reference vector is the local vertical. This may also be measured with electrostatic-potential sensors [32] or thermopiles [42] which attempt to identify a local horizontal plane based on the earth's electrostatic field or the local infra-red optical horizon, respectively. However, both these sensors are difficult to calibrate and produce relatively inaccurate results when the vehicle deviates significantly from a level orientation, or when atmospheric conditions change. Accelerometers are frequently used in static platforms to measure the local vertical based on the gravitation vector. However, accelerometers measure specific force, which is the difference between gravitational acceleration and translational acceleration of the sensor, and removal of the translational acceleration portion of the measurement for an accelerating vehicle can be difficult and will introduce errors.

The second classical reference vector is the local Earth's-magnetic-field vector, which is also mostly immune to weather conditions. However, properly calibrating a magnetometer to remove hard-iron-distortion and soft-iron-distortion effects, and having accurate knowledge of the local ambient magnetic field can both be difficult, leading to errors in the attitude estimate. Additional (or alternative) reference vectors can be obtained from specialized multi-antenna global navigation satellite systems (GNSS, e.g., GPS) receivers that determine the differential range to the spacecraft from each antenna. However, this approach often requires a high level of computational power to resolve the integer ambiguity problem and the short baselines (distance between antennas) on a sUAS do not allow for high-accuracy vector measurements. Another approach is to use on-board imaging to resolve the bearing to known ground landmarks [18] [83]. Although camera and

imaging processing capabilities are steadily improving, this method does not work in clouds or heavy precipitation and often lacks accuracy.

An alternative to direct attitude measurement is the integration of inertial rotation rates from a known starting attitude to produce attitude estimates thereafter. This method is associated with the term ‘strapdown navigation system’ as the sensors are fixed, or strapped down, to the platform rather than being gimbaled like a gyroscopic system. While any noise or bias in the inertial rate measurement is integrated, causing a growing variance with time, this method is still viable and electronic systems of this type have supplanted systems relying on a rotating wheel and mechanical components. Ring-laser-gyroscope and fiber-optic-gyroscope-based systems are now the standard for use in new manned-aircraft and spacecraft systems. Although there is no spinning wheel or other moving parts in these sensors (which measure rotation rate versus direct measurement of rotation) they share the term ‘gyroscope’ (or ‘gyro’) in their name with older, mechanical systems.

The advent of low cost MEMS gyroscopes makes this approach attractive for sUAS. Micro-electromechanical-systems (MEMS) gyroscopes are constructed with nanoscale electromechanical systems and supporting electronics packaged in a microchip. MEMS gyroscopes are typically based on a vibrating structure mechanism, such as a tuning fork gyroscope, or a hemispherical or cylindrical resonator gyroscope, and often have multiple sensors in a single package to provide a measure of rotation rate in 3 dimensions. Further references to a gyroscope (singular) assume a three-axis device and this convention will be used for accelerometers and magnetometers as well. The difficulty with using MEMS gyroscopes is that the zero point drift rates on this class of sensor do not enable accurate dead reckoning over even a few minutes. Higher quality sensors, such as fiber-optic gyros or ring-laser gyros are well beyond the cost, size, and power constraints of most sUAS vehicles. Correcting the attitude estimate integrated from MEMS gyroscope measurements with information from a magnetoresistive magnetometer, a MEMS accelerometer, and a GPS receiver (for use in estimating translational acceleration) allows for attitude estimation with an accuracy far better than that possible with either the direct measurement (accelerometer and magnetometer) or

integrated estimate (gyroscope) method alone. A sensor set commensurate with sUAS size, mass, and cost constraints, that will work in all weather conditions, includes a MEMS gyroscope and accelerometer, magnetoresistive magnetometer, and low cost GPS receiver. This set of sensors, possibly augmented with other sensors, has become the most frequently used sensor set for sUAS attitude estimation. Numerous approaches to the specific way that the sensor information is fused exist. The best method of fusing information from sensors depends on the particular sensors used as well as the operating environment (the wind environment in particular with respect to sUAS), and the precision required for the application of interest.

4.1 sUAS sensor characteristics

Understanding the general operating and error characteristics of sensors suitable for sUAS is necessary for analyzing sensor fusion schemes.

4.1.1 MEMS gyroscopes

MEMS gyroscopes have become the foundation of most sUAS attitude estimation systems. As these sensors measure rotation rate, which must be integrated to yield attitude information, and as the integration process causes the gyro error to have a cumulative effect on the integrated attitude, the gyro error characteristics are very important. Even high-end MEMS gyros have enough zero-point drift to cause integrated attitude error to accumulate rapidly. Dead-reckoned attitude using only MEMS gyros is only good for minutes (or seconds).

A measurement model adapted from Titterton [77] for a single (x) axis MEMS gyroscope is

$$\hat{\omega}_x = (1 + S_x)\omega_x + M_{xy}\omega_y + M_{xz}\omega_z + B_{fx} + \eta_x \quad (4.1)$$

where ω_n is the rotation rate about the N-axis, S_x is a scale factor constant, M_{xy} and M_{xz} are alignment constants, B_{fx} is the zero-point bias and η_x is a noise term. This model does not consider the




Figure 4.1: A Maxim 21100 sensor including a 3-axis gyroscope, 3-axis accelerometer and 3-axis magnetometer; this device represents typical physical sizing for contemporary MEMS sensors.

frequency response of the sensor itself, but does consider the frequency content of the measurement error. The terms S_x , M_{xy} , M_{xz} and B_{fx} all predominantly arise from deterministic sources such as misalignment, mis-calibration, temperature sensitivity, supply-voltage sensitivity, and acceleration sensitivity. The variability in these terms due to some of these sources can be removed based on a deterministic model. However there will still be residual variability, due to inadequacy of the model, and the error resulting from this residual variability is treated as stochastic, low-frequency error. In particular the zero-point bias is generally modeled as a (relatively) slow first-order auto-regressive process. As such integration of this error, which will have a non-zero expectation for significant, discrete time periods, results in a generally linear divergence with time of the integrated attitude. The term η_x represents high-frequency noise, which may be considered as white for discrete time samples. Integration of white noise results in a random walk of the integrated attitude with a variance that increases linearly with time.

A sensor fusion algorithm should work to both remove the angle-random-walk error produced by high-frequency gyro noise and to estimate the current zero-point bias so that it may be removed before integration and prevent the linear divergence of the integrated attitude.

4.1.2 MEMS accelerometers

MEMS accelerometers measure the specific force on a proof mass. Several definitions may be found for ‘specific force’, for example inertial acceleration times mass minus gravitational force. The accelerometer measurement is in fact a measure of the force required to hold the proof mass stationary relative to the accelerometer case. For a single axis accelerometer at rest this force is zero when the axis is orthogonal to the gravity direction and plus or minus 1g (at the Earth’s surface) when the axis is aligned with the gravity direction, with the sign dependent on the orientation of the accelerometer.

The fact that accelerometers directly measure neither acceleration nor gravitational force

complicates their use in state estimation. Ideally a sensor would be available to measure the gravity direction for use in estimating attitude, and a sensor would be available to measure inertial acceleration for use in estimating velocity.

Beyond this functional complication, MEMS accelerometers also have error characteristics. A measurement model adapted from Titterton [77] for a single (x) axis MEMS accelerometer is

$$\frac{\hat{s}f_x}{mass} = (1 + S_x)a_x - \frac{g_x}{mass} + M_y a_y + M_z a_z + B_{fx} + \eta_x \quad (4.2)$$

where *mass* is the mass of the accelerometer proof mass, S_x is a scale factor constant, M_y and M_z are alignment constants, B_{fx} and η_x are low-frequency and high-frequency error terms, a_n is the inertial acceleration in the N-axis direction and g_x is the gravitational force in the x direction. These error terms are similar to those for MEMS gyroscopes. For use in velocity estimation the gravity component is removed and the resulting acceleration measurement is integrated. This integration makes zero-point bias drift the most important error component. For use in attitude estimation the acceleration must be removed to arrive at a measurement of the gravity force, or in three dimensions the gravity vector. While not an internal error source, the external estimation of translational acceleration used to remove translational acceleration from the accelerometer measurement can be the most significant error with respect to sUAS attitude estimation.

4.1.3 Magneto-resistive magnetometers

Several types of electronic magnetic-field sensors are available including fluxgate, magneto-resistive, magnetoinductive, and others. Fluxgate based systems are often used in manned aircraft but tend to be too large and heavy for sUAS, and often have slow response times on the order of 2 to 3 seconds. Magneto-resistive magnetometers are sometimes incorrectly referred to as MEMS magnetometers as they are packaged in microchips, and sometimes are packaged with MEMS sensors in a single package. Currently available magneto-resistive sensors have sensitivities on the order of 0.1 milligauss and response times on the order of one microsecond.

While magnetoresistive sensors have internal errors, the errors associated with a magnetoresistive magnetometer also include distortions to the Earth's magnetic field due to magnetic sources and ferrous metals near the sensor. A measurement model for a three axis magnetometer mounted in an airframe adapted from Renaudin, Afzal, and Lachapelle [60] is

$$\hat{h} = \mathbf{S} \mathbf{M}(\mathbf{A}_{si}[\vec{h}_{earth}]_B + [\vec{b}_{hi}]_B) + b_{fx} + \epsilon = \mathbf{A}_e[\vec{h}_{earth}]_B + b_e + \epsilon. \quad (4.3)$$

where \vec{h}_{earth} is the undistorted ambient (Earth) magnetic field, \mathbf{S} is a scale-factor-constant diagonal matrix, \mathbf{M} is a misalignment-constant matrix, \mathbf{A}_{si} is a soft-iron-distortion matrix, \vec{b}_{hi} is a hard-iron-distortion vector, b_{fx} is a zero-point-bias column vector, and ϵ is a stochastic-error column vector term. The hard and soft-iron terms account for the fact that the magnetic field at the sensor differs from the undistorted local Earth's magnetic field. The deterministic error terms (\mathbf{S} , \mathbf{M} , \mathbf{A}_{si} , $[\vec{b}_{hi}]_B$, and b_{fx}) may be collected into a single square matrix, \mathbf{A}_e , and column vector, b_e .

The hard-iron-distortion vector accounts for other magnetic sources influencing the measurement. Consider a magnetic sensor operating in close proximity to a bar magnet, as shown in the left-hand panel of Figure 4.2. The magnetic field being measured by the sensor, for example as indicated by the heavy red arrow, is an additive combination of the Earth's magnetic field (light red lines) and the bar magnet's magnetic field (black lines). Permanent magnets, magnetized ferrous-metal components and magnetic fields produced by electric currents may all be present in an airframe. Proper calibration of a magnetometer requires characterizing the magnetic field produced by these sources so that the measurement of the Earth's magnetic field may be isolated. Magnetic fields produced by electric currents may be particularly troublesome in sUAS as electric propulsion systems may produce significant magnetic fields that may also be time varying due to changing thrust demands. The additive magnetic field resulting from a (constant) magnetic source in the airframe appears in the sensor measurement as a fixed bias. The combined error term b_e includes both the hard-iron distortion and zero-point bias errors

The soft-iron-distortion matrix accounts for the effect of ferrous metals on the Earth's mag-




Figure 4.2: Hard and Soft Iron Distortions.

netic field at the sensor as shown in the right-hand panel of Figure 4.2. Ferrous metals will distort, or bend, the Earth's field, but unlike magnetic sources in the airframe this bending effect at the magnetometer is dependent on the orientation of the airframe with respect to the Earth's magnetic field. For example if the magnetometer were in line with the ferrous material with respect to the Earth's magnetic field, as depicted by the arrow labeled '1' in the figure, then there would be little or no distortion. In some other orientation, for example as depicted by the arrow labelled '2', where the ferrous material is not in line with the magnetometer with respect to the Earth's magnetic field, then there would be a distortion. The distortion is dependent on the relative positions of the ferrous material and magnetometer with respect to the Earth's magnetic field, which is dependent on the airframe's orientation. This distortion changes both the strength and direction of the local magnetic field. The resulting effect is that the magnitude of the magnetic-field measurement in the body frame taken over all orientations will form an ellipsoid rather than a sphere; a fact that is useful for constructing a calibration algorithm as described in section 7.1.2.

The deterministic error terms in Equation (4.3), collected into the matrix \mathbf{A}_e and vector b_e , include a total of 12 parameters. Using estimates of these 12 parameters a calibrated measurement may be formed which removes most of the deterministic error. Residual errors must be treated as stochastic. Section 7.1.2 describes procedures for determining calibration parameters for a

magnetometer installed in a sUAS.

4.1.4 Commercial off-the-shelf (COTS) GPS/GNSS receiver modules

Global positioning system (GPS) receiver modules are the most common type of global navigation satellite system (GNSS) receiver module. GNSS includes the U.S. GPS system, the European Galileo positioning system, the Russian GLONASS system and the Chinese Compass navigation system. The latest generation of GNSS receiver modules are able to use signals from more than one of these systems simultaneously, improving performance. GNSS receiver modules have two primary uses in sUAS state estimation: as one of the primary sensors for position and velocity estimation and as a sensor involved in the attitude estimation sensor fusion. For the latter use the GNSS velocity estimate is used to estimate inertial acceleration for use in conjunction with the accelerometer measurement.

GPS measurements contain error due to a number of sources including errors in the satellite orbit and clock models, the effect of the ionosphere and troposphere on propagation of the satellite signals, multi-path reflections of the signals, receiver clock instability, and receiver noise. Some of these factors, such as ionosphere and troposphere effects have a much greater effect on position measurement than on velocity measurement. Commercial off-the-shelf (COTS) GPS modules' velocity measurement error generally is modeled as simple additive noise.

The basis for GNSS receiver velocity measurement/estimation is measurement of the Doppler frequency of the signals received from multiple satellites. When signals from four or more satellites are being tracked, and a Doppler measurement is available for each satellite signal, a least-squares optimization is used to estimate the receiver velocity in the inertial frame and the receiver clock drift rate. The GNSS satellites all have atomic-clock timebases, and the GNSS systems' ground segments monitor and model the drift and drift rate of the satellite clocks and provide correction information in the satellite ephemeris data. The receiver clock, however, will be either a crystal

oscillator or a temperature-controlled crystal oscillator, and the receiver's clock drift rate must be included as a parameter in the least squares optimization. For this reason a minimum of four satellites are required to produce a 3D velocity estimate.

GNSS receivers utilize a frequency-lock loop to track the Doppler frequency of each satellite signal. When the frequency-lock loop has a good lock the signal is handed over to a phase-lock loop. The phase-lock loop will typically use a Costas discriminator, and as long as lock is maintained the total phase change between measurement epoch's can be measured to a fraction of a cycle. For measurement epochs spaced at typical COTS GNSS receiver module output rates (10Hz or less) this results in quite accurate estimates of the radial range rate between the satellite and receiver. When a receiver is able to maintain phase-lock on a good number of satellites (6+) then velocity measurement error can be on the order of 1 cm/sec, even for inexpensive COTS GNSS receiver modules. The conditions under which these results may be expected in examined further in Chapter 5.

Specialized GNSS receivers with more than one antenna, or systems using multiple GNSS receivers, may also be used to directly measure attitude information. Such a receiver/system may directly measure the orientation of the baseline between two antennas in the inertial frame. There are two issues with using this technique for sUAS. First, the difference in phase length between a satellite and each antenna must be determined. This involves resolving the integer number of cycles in the difference. Various techniques are available to solve this problem, but all generally require significant computational resources (or data storage), making this technique unattractive for sUAS. More significantly, the accuracy of the measurement of the baseline attitude is limited by the length of the baseline. For sUAS baseline lengths are typically limited to less than two meters. The GPS L1 frequency, used by most COTS receiver modules, has a wavelength of 19.0 cm. If a receiver is able to measure the relative phase to within 1/10 cycle, then for a 1 meter baseline the resulting attitude measurement will have an error on the order of one degree. This is adequate for some purposes, but inadequate for high-accuracy sUAS-based wind measurement.

4.1.5 Other sensors

Other sensors are sometimes used in sUAS state estimation. Although the set of four sensors considered above (gyroscope, accelerometer, magnetometer, GNSS) are the sensors used in this research, and are the core sensor set for most contemporary UAS state estimation systems, a brief explanation of some other sensors found in the literature is provided.

Pitot/static pressure sensors are common on aircraft and are integral to most relative-wind sensors. Information from these sensors may be included in state-estimation sensor fusion. A specific case is further examined later in this chapter.

Readily available thermopile sensors provide a measure of thermal infrared radiation, typically in the band between 7.5 and 13.5 microns. A set of three or more thermopiles may be employed to estimate a vector perpendicular to the local horizon plane, which will generally coincide with the local vertical. Infrared horizon sensing has been used in satellite stabilization since the 1950s and prior to recent advances in MEMS sensors was often employed in sUAS state-estimation or control systems. While calculation of a vector orthogonal to the local horizon from the response of multiple thermopile sensors with different fields of view is relatively straightforward, several problems arise. First, the local horizon may vary significantly with azimuth due to terrain, and even if the local horizon is flat it may not coincide with the local horizontal plane. In clear sky conditions the sky/ground temperature difference is generally tens of degrees Celcius, but clouds, drizzle or fog can reduce this difference significantly, reducing the accuracy of the horizon measurement. Further, sky conditions that vary with azimuth can introduce a bias. Finally contamination on the sensor aperture, such as a precipitation droplet can significantly affect performance.

Electrostatic potential sensors have also been used [32] and can produce a measurement of the local vertical. Accuracy of this method is limited in sUAS by the limited baseline lengths and anthropogenic noise is an increasing issue. There have been no recent examples of use of this sensor type in UAS in the literature.

Optical systems have been employed based on angular measurements to known landmarks, simultaneous location and mapping (SLAM) techniques, or optical flow techniques. Computational loads have prevented widespread use of these systems, but continued improvement in lightweight, low power computer hardware may make these systems more suitable for sAUS use in the future. Other optical systems such as laser rangefinders have been used in sUAS for limited purposes in state estimation such as estimation of height over terrain. Chapter 6 introduces a new optical reference vector sensor which is suitable for use in sUAS state estimation.

4.2 Sensor fusion for sUAS attitude estimation

4.2.1 System model

Attitude, which is part of an aircraft 'state', is governed by a dynamic process involving forces and moments on the airframe. These forces and moments will be partially stochastic due to the effect of the (unknown) inertial wind. Numerous filters, or estimators, have been developed which are directly applicable to this problem [15]. Development of a filter to estimate attitude invariably begins with producing a model of the system. Initially it would appear that the system model should be a model incorporating the aerodynamic forces and moments on the airframe and the airframe mass and moment-of-inertial properties. System inputs would include control surface deflections and propulsion system thrust. The difference between the instantaneous wind and the mean wind would be treated as a disturbance. In practice, however, this sort of model is almost never used for UAS state estimation for two reasons. First, producing an accurate parametric model of this type is difficult, and sUAS are often subject to variation from unit to unit which would require characterizing each airframe, and re-characterizing each airframe after repair or modification. Also, and more significantly for the problem at hand, the disturbances caused by wind can be large and basing the filter on a model with large disturbances results in a low-accuracy state estimate.

Rather than a force-and-moments model, a simple rigid-body kinematic-motion model is

used. This model uses the rotation rates and acceleration, as measured by the gyroscope and accelerometer, as inputs. The ‘state’ of this model, which includes attitude, velocity and position states, is

$$\mathbf{x} = \begin{bmatrix} \mathbf{q}_I^B{}^T, & [{}^I\vec{v}_B]_I{}^T, & [{}^I\vec{p}_B]_I{}^T \end{bmatrix}^T = \begin{bmatrix} \mathbf{q} \\ \mathbf{v} \\ \mathbf{p} \end{bmatrix}, \quad (4.4)$$

where $\mathbf{q}_I^B = \mathbf{q}$ is the quaternion representation of the rotation of the body frame relative to the inertial frame, $[{}^I\vec{v}_B]_I$ is the velocity of the body frame relative to the inertial frame in inertial frame coordinates, and $[{}^I\vec{p}_B]_I$ is the position (translation) of the body frame relative to the inertial frame in inertial coordinates. The kinematic model relating state derivatives and motions is

$$\dot{\mathbf{x}} = \begin{bmatrix} \dot{\mathbf{q}} \\ \dot{\mathbf{v}} \\ \dot{\mathbf{p}} \end{bmatrix} = f(\mathbf{x}, \mathbf{u}) = \begin{bmatrix} \frac{1}{2}\mathbf{q} \otimes [{}^I\vec{\omega}_B]_B \\ \mathbf{q}^* \otimes [{}^I\vec{a}_B]_B \otimes \mathbf{q} \\ \mathbf{v} \end{bmatrix} \quad (4.5)$$

where $\mathbf{u} = [[{}^I\vec{\omega}_B]_B{}^T, [{}^I\vec{a}_B]_B{}^T]^T$ and $[{}^I\vec{\omega}_B]_B$ and $[{}^I\vec{a}_B]_B$ are the rotation rate and acceleration of the body frame relative to the inertial frame in body frame coordinates. Appendix 1 describes quaternion attitude representation and operations as used in Equations (4.4) through (4.8). The state of this model is related to the accelerometer, magnetometer, and GNSS (velocity and position) measurements by

$$\mathbf{y} = \begin{bmatrix} \mathbf{sf}_{meas} \\ \mathbf{b}_{meas} \\ \mathbf{v}_{meas} \\ \mathbf{p}_{meas} \end{bmatrix} = h(\mathbf{x}, \mathbf{u}) + \boldsymbol{\eta} = \begin{bmatrix} [{}^I\vec{a}_B]_B - \mathbf{q} \otimes [{}^I\vec{g}]_I \otimes \mathbf{q}^* + \eta_{accel} \\ \mathbf{q} \otimes [{}^I\vec{b}]_I \otimes \mathbf{q}^* + \eta_{mag} \\ \mathbf{v} + \eta_{veloc} \\ \mathbf{p} + \eta_{pos} \end{bmatrix}, \quad (4.6)$$

where \vec{b} is the local earth’s magnetic-field vector and \vec{g} is the local gravity vector. The dependence of the accelerometer and magnetometer measurements on the attitude state comes through the body-frame measurement of the gravity and magnetic-field vector. Here, the η terms encompass all (post-calibration) measurement error sources.

This model can be used to incrementally estimate attitude, velocity and position by using a gyroscope measurement, $\boldsymbol{\omega}_{meas}$, and accelerometer measurement, \mathbf{sf}_{meas} , instead of the true rotation rate and acceleration:

$$\dot{\tilde{\mathbf{x}}} = f(\tilde{\mathbf{x}}, \tilde{\mathbf{u}}) = \begin{bmatrix} \frac{1}{2}\tilde{\mathbf{q}} \otimes \boldsymbol{\omega}_{meas} \\ \tilde{\mathbf{q}}^* \otimes \mathbf{sf}_{meas} \otimes \tilde{\mathbf{q}} - [\tilde{\mathbf{g}}]_I \\ \tilde{\mathbf{v}} \end{bmatrix} \quad (4.7)$$

where the $\tilde{\cdot}$ symbol denotes a measured or estimated parameter and $\tilde{\mathbf{u}} = [\boldsymbol{\omega}_{meas}^T, \mathbf{sf}_{meas}^T]^T$. The incremental estimate error will drift over time, but the incremental estimates can be compared with absolute measurements to correct this drift. Using the observations model (4.6) to produce a sensor measurement prediction yields

$$\tilde{\mathbf{y}} = \begin{bmatrix} \mathbf{sf}_{pred} \\ \mathbf{b}_{pred} \\ \mathbf{v}_{pred} \\ \mathbf{p}_{pred} \end{bmatrix} = h(\tilde{\mathbf{x}}, \tilde{\mathbf{u}}) = h(\tilde{\mathbf{x}}) = \begin{bmatrix} [{}^I\tilde{\mathbf{a}}_B]_B - \tilde{\mathbf{q}} \otimes [\tilde{\mathbf{g}}]_I \otimes \tilde{\mathbf{q}}^* \\ \tilde{\mathbf{q}} \otimes [\tilde{\mathbf{b}}]_I \otimes \tilde{\mathbf{q}}^* \\ \tilde{\mathbf{v}} \\ \tilde{\mathbf{p}} \end{bmatrix}. \quad (4.8)$$

Note that a value for $[{}^I\tilde{\mathbf{a}}_B]_B$ is required to evaluate this expression; estimating $[{}^I\tilde{\mathbf{a}}_B]_B$ is discussed in section 4.2.4. Now the prediction error $(\mathbf{y} - \tilde{\mathbf{y}})$ and an appropriate gain matrix K may be used to correct the estimates in a state observer structure as follows;

$$\dot{\tilde{\mathbf{x}}} = f(\tilde{\mathbf{x}}, \tilde{\mathbf{u}}) + K(\mathbf{y} - \tilde{\mathbf{y}}). \quad (4.9)$$

To place this in a discrete-time framework for ease of implementation and comparison, use a simple first-order Euler approximation to obtain

$$\tilde{\mathbf{x}}_k = \tilde{\mathbf{x}}_{k-1} + f(\tilde{\mathbf{x}}_{k-1}, \tilde{\mathbf{u}}_k)\delta t + K(\mathbf{y}_k - h(\tilde{\mathbf{x}}_{k-1})) = \tilde{\mathbf{x}}_{k-1} + f_k\delta t + K\mathbf{z}_k. \quad (4.10)$$

Note that \mathbf{z}_k is often termed the ‘innovation’ in the ‘observation’ sensor measurement at time-step k . The difference between algorithm **types** and **structures** will be seen in the form of the gain matrix, K , as described below.

Before leaving the topic of system models it is noted that since the kinematic model uses rotation rate and acceleration as inputs, and since (calibrated) gyroscope and accelerometer measurements are equal to the true rotation rate and acceleration plus an additive error, and since a portion of this additive error may be treated as deterministic or slowly varying, it is often useful to augment the system model with terms relating to the gyroscope and accelerometer error. In particular, filter algorithms allowing estimation of gyroscope biases will have significantly better attitude-estimation performance because filters without a method of accounting for gyro-bias drift will have a steady-state attitude error whenever there is a non-zero zero-point bias. Similarly, filters without a method of accounting for accelerometer bias drift will have a steady-state velocity error whenever there is a non-zero zero-point bias in the accelerometer measurements. Some filter **structures**, such as the extended Kalman filter (EKF), account for bias using additional elements in the state that are estimated. Other filters, for example many complimentary filter implementations [19] [44], do not explicitly estimate the bias as part of the state, but use proportional-integral (PI) feedback where the integrator states track the bias. For high-accuracy attitude estimation using contemporary MEMS gyroscopes, use of a filter **structure** which accounts for gyroscope biases is essential, and all attitude filter simulation results herein either include gyroscope bias in the estimated state (EKF, unscented Kalman filter (UKF)) or incorporate PI feedback where the integrator states track gyroscope bias. However bias states are omitted in analytical analyses of sections 4.2.2 and 4.2.3 where particular filter tradeoffs are considered since the inclusion of bias states or integrator states is not germane to the particular tradeoff analyses. This is not intended to imply that gyroscope bias tracking is not important and can be omitted in final filter design. Similarly estimation of accelerometer biases may be necessary to obtain desired performance. However for many attitude-estimation cases, calibration of the accelerometer over temperature may be sufficient as residual bias may be relatively small and unlike the gyroscope measurement integration of the accelerometer measurement is not needed for attitude estimation.

4.2.2 Filter Structure—state estimate update variations

All contemporary sUAS attitude estimation systems and state estimation systems are discrete time systems implemented with digital computers. Many discrete time algorithms exist for fusing sensor data. The term filter **structure** will be used to differentiate between different algorithms which operate on a particular set of sensor inputs to produce a particular set of estimated quantities. Examples of different filter **structures** include fixed-gain observers, complimentary filters, extended Kalman filters, particle filters, etc.

The question of what filter **structure** is most appropriate for sUAS attitude estimation is considered here. Aside from the filter **structure** employed, sUAS attitude may be estimated with a filter estimating only attitude, commonly termed an attitude and heading reference system (AHRS), or may be estimated with a filter estimating attitude, velocity and position, commonly termed an inertial navigation system (INS). The question of which of these two approaches should be used is deferred to section 4.2.3. AHRS will be used for comparison of filter **structures** for attitude estimation.

While the extended Kalman filter (EKF) is a prevalent attitude-estimation filter, and has become the de facto standard for low-cost AHRS systems, there are still a wide range of filter **structures** in use. This includes a host of non-linear observers and complimentary filters [19] [44] [85] [13] [38] as well as Kalman filter derivatives (EKF, UKF, etc.) [15] [20] [37] [45]. Interestingly many of these can be shown to be related through underlying cost functions.

All the attitude filters considered herein, based on MEMS gyroscopes and accelerometers, magnetometers and GNSS receivers, can be viewed in a general framework as a feedback system as shown in Figure 4.3. In the context of a complimentary filter the gyroscope is viewed as an accurate sensor at high frequencies, and the accelerometer and magnetometer are used only for low-frequency (drift) correction. While it would be possible to construct a complimentary filter in an explicit fashion by high-pass filtering an integrated gyroscope signal and low-pass filtering an

attitude estimate from a Wahba problem solution [82] based on accelerometer and magnetometer measurements, and then combining these attitude estimates, a filter design employing feedback, as in Figure 4.3, is almost always used in practice. In the context of a Kalman filter the plant model is the kinematic rotation of the airframe with the plant model input being the rotation rates as measured by the gyroscope (Equation (4.7)), and the observation sensors are the accelerometer and magnetometer. Some formulations treat the accelerometer and magnetometer as independent observation sensors while other formulations combine the accelerometer and magnetometer information through a Wahba problem solution and use the resulting attitude estimate as an observation.

As described by Figure 4.3, attitude estimation filters operate by integrating the gyroscope measurement, corrected by a filtered error signal, to form an attitude estimate. The error signal is formed by comparing the accelerometer and magnetometer measurements with their a-priori estimates calculated with the current attitude estimate using mappings \mathbb{F} and \mathbb{H} . Other sensor information, for example GNSS velocity measurements, or other internal information, for example filter integrator values when using PI feedback in the mapping \mathbb{G} [44], may also be used in the a-priori estimate calculations. The error signal may undergo filtering before use as the correction signal; for example PI feedback is commonly used in complimentary filter attitude estimators [19].

The relationship between different filter **structures** may be seen by considering a generalized least squares cost function leading to different choices for the mapping function \mathbb{H} . Consider the generalized cost function for a four state quaternion based attitude filter

$$J_k = \frac{1}{2} \sum_{i=1}^k \mathbf{\Lambda}(k-i) [\mathbf{z}_i - \mathbf{H}_i \delta \tilde{\mathbf{q}}_k]^T \mathbf{W}_i [\mathbf{z}_i - \mathbf{H}_i \delta \tilde{\mathbf{q}}_k]. \quad (4.11)$$

where $\mathbf{\Lambda}$ is a scalar windowing or ‘forgetting’ function, $\mathbf{z}_i \left(= \begin{bmatrix} \mathbf{sf}_{meas} - \mathbf{sf}_{pred} \\ \mathbf{b}_{meas} - \mathbf{b}_{pred} \end{bmatrix} \right)$ is the innovation of the i ’th observation, \mathbf{W}_i is a directional weighing function, \mathbf{H}_i is the Jacobian of the mapping function from attitude to observation sensor measurement, and $\delta \tilde{\mathbf{q}}_k$ is the error in the state estimate at time k . A large variety of attitude-estimation filters use a feedback signal equivalent to the state estimate error, $\delta \tilde{\mathbf{q}}$, that minimizes this cost function, with differences between

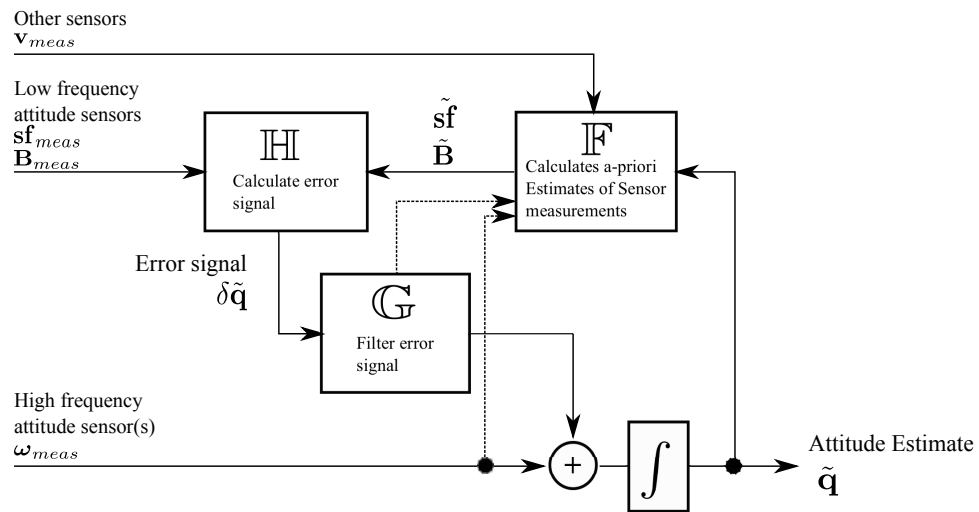


Figure 4.3: Attitude Estimation Filter feedback diagram. Almost all attitude-estimation filter structures conform to this general feedback structure.

filters resulting from different choices of $\mathbf{\Lambda}$ and \mathbf{W} . Several complimentary filters result from using uniform weighting ($\mathbf{W}_i = \mathbf{I}$) and a windowing function which considers only the current observation sensor measurement, \mathbf{z}_k , meaning $\mathbf{\Lambda}(\cdot) = \delta(\cdot)$ (dirac delta function), which results in $J_k = \frac{1}{2} [\mathbf{z}_k - \mathbf{H}_k \delta\tilde{\mathbf{q}}_k]^T [\mathbf{z}_k - \mathbf{H}_k \delta\tilde{\mathbf{q}}_k]$. The error quaternion minimizing this cost is

$$\delta\tilde{\mathbf{q}}_k = [\mathbf{H}_k^T \mathbf{H}_k]^{-1} \mathbf{H}_k^T \mathbf{z}_k. \quad (4.12)$$

Using this solution for the error quaternion update is problematic as the inverse term can be poorly behaved particularly at high/low latitudes where the gravity and magnetic-field vectors are nearly co-linear, causing linear dependance between rows of \mathbf{H} , or when processing accelerometer and magnetometer measurements separately, which causes \mathbf{H} to be rank 3. However, several filters use modified forms of this error quaternion calculation:

- Fixed-step-size gradient-descent filter [44]

$$\tilde{\mathbf{q}}_k = \tilde{\mathbf{q}}_{k-1} + f_k \delta t + \alpha \frac{\mathbf{H}_k^T \mathbf{z}_k}{|\mathbf{H}_k^T \mathbf{z}_k|} \quad (4.13)$$

- Gradient-descent filter [55]

$$\tilde{\mathbf{q}}_k = \tilde{\mathbf{q}}_{k-1} + f_k \delta t + \alpha \mathbf{H}_k^T \mathbf{z}_k \quad (4.14)$$

- Modified-gradient-descent filter [21]

$$\tilde{\mathbf{q}}_k = \tilde{\mathbf{q}}_{k-1} + f_k \delta t + \alpha [\mathbf{H}_k^T \mathbf{H}_k + \mathbf{Q}_0]^{-1} \mathbf{H}_k^T \mathbf{z}_k \quad (4.15)$$

In the first two cases above α is a gain term treated as a tuning parameter. For the modified-gradient-descent filter the \mathbf{Q}_0 matrix is added in the inverse term as a heuristic measure to counteract ill-conditioning and is treated as a tuning parameter. The resulting form bears interesting similarity to the EKF state-mean-estimate predict and update equation pair

$$\tilde{\mathbf{q}}_k^- = \tilde{\mathbf{q}}_{k-1}^+ + f_k \delta t \quad (4.16)$$

$$\tilde{\mathbf{q}}_k^+ = \tilde{\mathbf{q}}_k^- + \mathbf{P}_k^- \mathbf{H}_k^T [\mathbf{H}_k \mathbf{P}_k^- \mathbf{H}_k^T + \mathbf{R}_k]^{-1} \mathbf{z}_k. \quad (4.17)$$

See Appendix 3 for a summary of the EKF algorithm. The EKF update equation can be derived using a uniform windowing function ($\mathbf{\Lambda}(\cdot) = 1$) in Equation (4.11), and then following a derivation approach as presented by Tapley, Schutz and Born [73]:

- adding a term to the cost function for an initial epoch a-priori state-estimate error,
- following a minimum variance estimate analysis,
- following the typical derivation of the sequential update filter algorithm from the batch update algorithm filter, and
- recognizing that the initial model derivation used herein prefers re-linearization of the problem at each update time epoch.

The UKF differs from the EKF in the propagation of the mean-state-estimate and state-error-covariance using sigma points versus a linear approximation to the non-linear kinematic (plant) model function, but the structure uses the same underlying cost function with additional weighting across sigma points.

Examining Equations (4.14) and (4.17) illustrates the difference between most fixed-gain-observer attitude filters and EKF attitude filters. Both filters use an update that applies a linear correction step based on the measurement innovation. The gradient-descent filter takes a correction step in the direction of the gradient of the prediction error and with a magnitude dependent on the gradient magnitude, the innovation, and the tuning parameter, α . The EKF takes a correction step in the direction of the gradient modified by the state-error covariance and a term incorporating the gradient, state-error covariance, and the observation-sensor-error covariance, using a step size that similarly incorporates all this information. The implication is that the gradient-descent filter update incorporates information about the observability of the state from the observation, while the EKF incorporates both information about the observability of the state from the observation and information about the state-error covariance.




Figure 4.4: 1 degree of freedom platform example system.

4.2.2.1 Example analysis for a 1-degree-of-freedom (1DOF) platform

To gain some understanding of the frequency-response characteristics of the gradient-descent filter and the EKF, the filters are examined in the context of a simplified system. Consider a platform of unit length which is free to rotate in pitch (only), and for which we wish to estimate the pitch angle, as shown in Figure 4.4. Sensors used are a gyroscope measuring rotation rate about the platform pitch axis and a sensor measuring the height of the platform's raised edge (i.e., y).

Both filters follow the generalized observer framework of Equation (4.10), with the f term calculated from the gyroscope measurement and the \mathbf{z} term calculated from the height sensor measurement. The transfer function from the height sensor to the estimated pitch angle is

$$H_{y \text{ to } \hat{\theta}}(s) = \frac{K}{s + K}. \quad (4.18)$$

For the gradient-descent filter the gain term is $K = \alpha \mathbf{H} = \alpha \cos \tilde{\theta}$. This results in a somewhat intuitive behavior of the gain term with the gain decreasing as the pitch angle increases; as the pitch angle increases the observation measurements have less information about the pitch angle. Note that for a fixed pitch angle the behavior is essentially a first-order low-pass filter but the pole is not strictly fixed due to its dependency on the estimated pitch angle.

For the EKF the Kalman gain, $K = \mathbf{P}\mathbf{H}^T(\mathbf{H}\mathbf{P}\mathbf{H}^T + \mathbf{R})^{-1}$, has two dependencies. First, as is the case for the gradient-descent filter the feedback gain depends on the pitch angle, with the same functional dependence through \mathbf{H} . It also depends on the state-error covariance estimate, \mathbf{P} , which affects the (scalar in this example) terms \mathbf{P} and $(\mathbf{H}\mathbf{P}\mathbf{H}^T + \mathbf{R})^{-1}$ (versus α in the gradient-descent

filter). Again, for a fixed pitch angle the behavior is close to a first-order low-pass filter but the pole will move somewhat due to the gain term's dependence on the estimated pitch angle and the estimated variance of the pitch estimate. **It is particularly interesting** to note that for a fixed pitch value the state-error covariance estimate and the gain term will converge, generally quite quickly, and that the rate of convergence is somewhat insensitive to the values of R and Q . The convergence evolves following the non-linear equation

$$P_{k+1} = P_k \frac{R}{P_k \cos^2 \bar{\theta} + R} + Q. \quad (4.19)$$

Figure 4.5 shows the (converged) gain behavior for both the gradient-descent filter and the EKF for this 1DOF platform example. The gain behavior for a gradient-inverse filter (Equation (4.12)) is also shown, illustrating the problematic gain behavior when the observability becomes small. The behavior of the EKF gain is governed both by the information content (observability) of the observation measurement and the state-error variance. The observability drops off as cosine of the pitch angle. However, the state-error variance increases with pitch angle due to the decreasing information content in the observation measurements. The net effect is a modest increase in the gain value with increasing pitch angle up to a relatively high angle where the observability becomes very poor and the gain drops rapidly. This result is perhaps more desirable than the gradient-descent filter behavior; the EKF has little change in feedback gain, and therefore the projection sensor to pitch estimate frequency response is nearly constant until the observability becomes very poor (at which point the pitch sensor is disregarded). This is the desired behavior because the frequency dependent qualities of the gyroscope and projection sensor do not vary with pitch angle.

This analysis of the frequency response of attitude-estimation filters provides an important insight. While these filters do not have a fixed corner frequency, the corner frequency will be constant, and calculable, when the observability remains constant for more than a few discrete-time updates of the filter. The corner frequency is a function of observability in complimentary filters and similar filters, while it is a function of both observability and state-estimate covariance in EKF and similar filters. The state-estimate covariance dependance on observability provides a

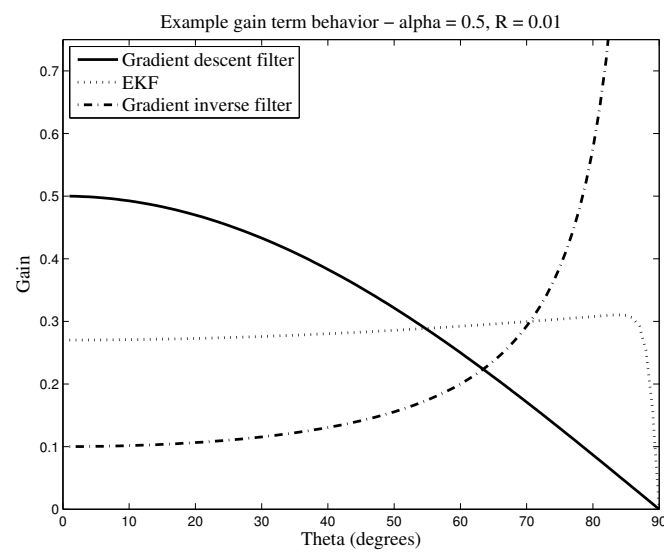


Figure 4.5: 1DOF example - Gain variation as a function of pitch angle.

stabilizing effect on the corner frequency (as a function of observability) in EKF. This feature may make EKF type filters desirable for wind-measurement systems when the wind measurements have interesting frequency content at or above the motion frequencies of the aircraft.

4.2.2.2 Performance comparisons in wind

The EKF offers the potential for better performance than the gradient-descent filter, or other fixed-gain filters, as the update corrections incorporate state-error covariance information in addition to observability information. And the UKF offers the potential for better performance than the EKF as the state mean estimate and error covariance estimates may be more accurately propagated in non-linear systems using sigma points. However, sensor measurement errors may dominate the state estimation errors, rather than sub-optimality of the estimate updates. For example, Fiorenzani et al. [20] found that for a small VTOL UAS performance between an EKF and UKF were very similar except when large initial errors were introduced. This leads to the question of how sophisticated a filter is required in the context of attitude estimation for sUAS flown in wind using contemporary sensors. At some level it is expected that errors introduced by the sensors will be far larger than errors introduced by one filter **structure** versus another.

Figure 4.6 and Table 4.2 show simulation results comparing the performance of 4-state AHRS gradient-descent filter with PI feedback, and 7-state EKF and UKF AHRS filters, using the standard simulation sensor noise levels given in Appendix 4 as well as with the simulated sensor noise increased and decreased by a factor of four. Two points are made by Figure 4.6 and Table 4.2. First, with the sensor noise levels considered herein the EKF performs a bit better than the gradient-descent filter, but the performance of the EKF and UKF are very nearly equal. Second, as sensor noise levels increase there is more difference between the performance of different filter **structures**; this implies the choice of sensor fusion algorithm is important when using lower performance sensors to meet cost/power/weight constraints.

Table 4.1: Attitude estimation comparison for different filter types at 3 wind/turbulence levels.

| | 1/4 * Standard sensor noise and drift | Standard sensor noise and drift | 4 * Standard sensor noise and drift |
|--|---------------------------------------|---------------------------------|-------------------------------------|
| Roll estimate error (degrees-r.m.s.) | | | |
| Gradient descent filter | 0.058 | 0.142 | 0.467 |
| EKF (or UKF) | 0.041 | 0.078 | 0.265 |
| Pitch estimate error (degrees-r.m.s.) | | | |
| Gradient descent filter | 0.052 | 0.113 | 0.374 |
| EKF (or UKF) | 0.047 | 0.063 | 0.200 |

Using sensor characteristics representative of those currently available and applicable to sUAS attitude estimation (the standard sensor noise case), there is a definite benefit to using an EKF filter **structure** versus a gradient-descent filter **structure**. For this simulation the gradient-descent filter was hand tuned for optimal performance - a fairly simple exercise as this filter has only a single P gain and a single I gain. Although higher than the error produced by the EKF the error levels are still comparable making the gradient-descent filter an appropriate choice in some situations, as for example when the attitude estimation must be accomplished with a low power micro-controller unable to support the computational load of an EKF.

The potential advantage of a UKF, however, is not realized as the error due to propagation method of mean and covariance estimates is negligible in comparison to error due to the imperfect sensors. At the scale shown in the figure the EKF and UKF estimates are indistinguishable.

In the simulation results presented in Figure 4.6 there is error in the attitude solution due to error in the gyroscope, accelerometer, and magnetometer sensors and from the GNSS-based translational acceleration estimate. The error introduced by just the gyroscope sensor alone, however, is still much greater than that introduced by the propagation of mean and covariance estimates using either the EKF's linearization or the UKF's sigma point process; in a separate simulation (results not shown) performed with the platform at rest and perfect magnetometer and GNSS sensors the difference between the EKF and UKF attitude estimates is still indistinguishable at the scale of

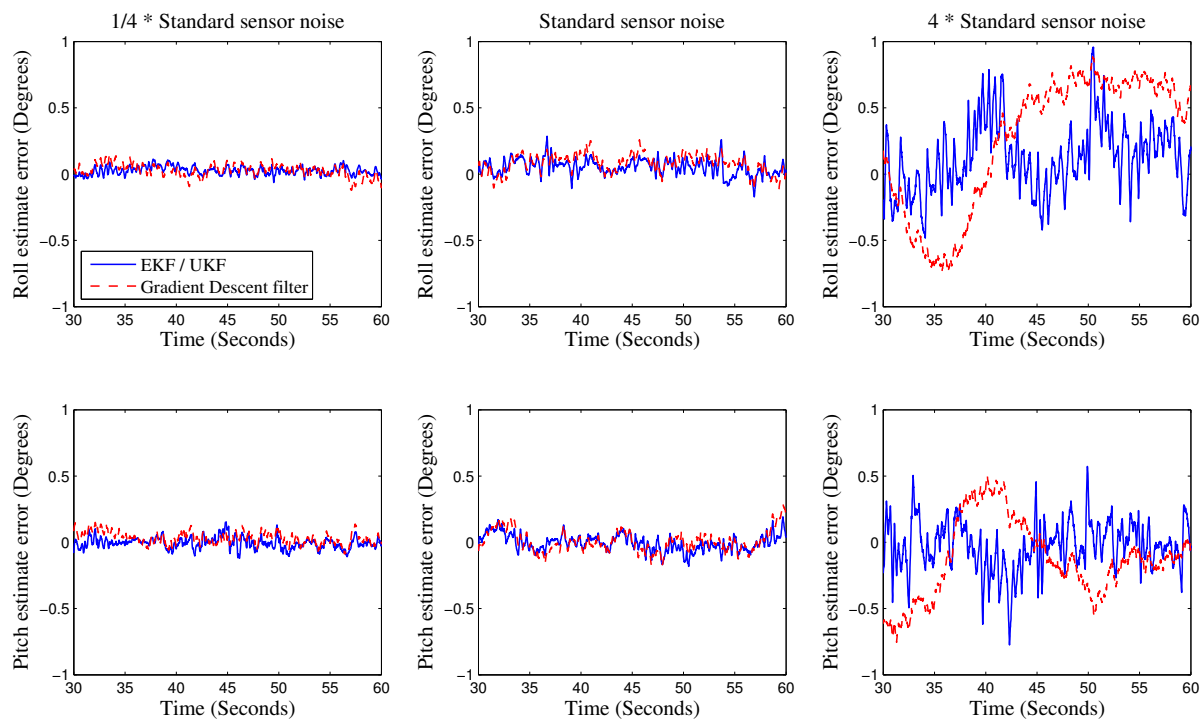


Figure 4.6: Attitude Estimation comparison for different filter **structures**.

Figure 4.6

While the integration of rotation rate to attitude is an inherently non-linear process the performance equivalence of the EKF and UKF for this application is not surprising given that the amount of rotation in any time step is small with high-rate gyro sampling and integration. Linearizing the rotation about the current attitude causes little error when propagating the covariance estimate. Similarly there are a variety of more accurate methods that may be used with the EKF that have little or no performance advantage in this application. These include:

- Use of higher order discrete time approximations to the rotation integration. The rotation integration described in Equation (4.7) is usually converted to discrete time with the first-order approximation $\tilde{\mathbf{q}}_{k+1} = \tilde{\mathbf{q}}_k + \frac{1}{2}(\tilde{\mathbf{q}}_k \otimes \boldsymbol{\omega}_{meas})\delta t$. Second and higher order approximations may be used [65], but provide little benefit for this application when the filter integration occurs at rates at or above 100Hz, as commonly supported by MEMS gyroscopes, and rotation rates are modest.
- Use of Van Loan method or other matrix exponential approximations. The Van Loan method is used in some EKF implementations to numerically convert the continuous-time linearized dynamics model and process noise matrices into the discrete-time state transition and discrete-time process noise matrices [52]. This provides little benefit over using a first-order approximation for the matrix exponential ($e^A \approx I + A$) for this application as the error introduced using the first-order approximation is insignificant when compared to the error introduced by the sensors.
- Use of a Wahba solution for observation measurements. Several EKF implementations pass the magnetometer and accelerometer measurements through a Wahba problem solution and use the resulting attitude estimate as an observation [45] [64]. For this application the observation innovations generally represent a small rotation and consequently this technique provides little or no improvement. Also, this technique requires processing accelerometer

and magnetometer measurements at the same rate whereas actual sensors may impose differing rates.

4.2.2.3 Effect of smoothing algorithms

The EKF is a recursive filter which incorporates all preceding observation measurements into the current estimate. For wind measurement, as for some other applications, the end measurement is not always needed in real time, and consequently sensor data may be post-processed to yield an attitude estimate. This allows the alternative of using a batch filter, or using first an EKF and then using a ‘smoother’ so that the estimate at each time step incorporates all observation measurements, both preceding and succeeding observations.

The Rauch-Striebel-Tung (RTS) [59] smoother, which uses ‘future’ samples (in post-processing) to improve the estimate at each time step. was examined in conjunction with an EKF. Simulation results showed that use of the RTS smoother did smooth the time series of attitude estimates, but had little effect on the average attitude estimate error. Even so, use of a smoother may be appropriate, for example in the case of wind measurement, where the time or spatial frequency characteristics of the attitude estimate are important as the smoothing operation removes some high frequency content from the attitude estimate error.

4.2.3 Filter Types - federated versus cascaded filters

In the case of in-situ wind measurement from a sUAS, as in many systems, a position and velocity estimate is needed in addition to an attitude estimate. This leads to two **types** of state estimation systems as shown in Figure 4.7. The first **type**, designated as federated filters and specifically for this application inertial navigation systems (INS), use a single estimator stage to estimate position, velocity, and attitude. The second **type**, designated as cascaded filters, use two estimator stages in series with the first stage, designated an attitude-heading reference system

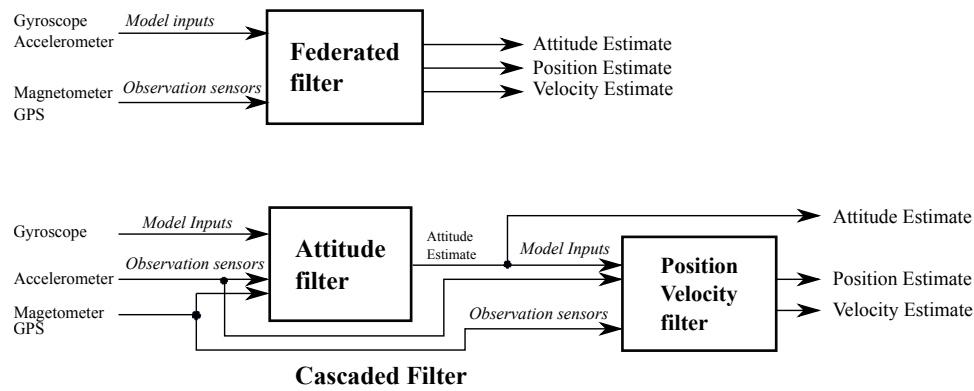


Figure 4.7: Federated versus Cascaded filters.

(AHRS), producing an attitude estimate and the second stage producing a position and velocity estimate. The term **Type** is used herein to differentiate between AHRS and INS filters used to produce attitude estimates. In both cases mean wind speed is often estimated (not explicitly shown in figure), as well as gyro and accelerometer biases and possibly other parameters; these states are omitted in the following analysis for clarity of comparison.

The federated-filter approach is often preferred as it has the advantage of (potentially) offering a near-optimal state estimate, and avoiding problems with needing specific-force correction of the accelerometer measurement (examined in section 4.2.4). The cascaded-filter approach is generally considered less accurate, while being less computationally costly. However, while the cascaded approach results in a sub-optimal estimate of the full state, the estimation of the attitude states is near optimal and is more accurate than the estimate of attitude states produced by a federated filter. Further, it is shown in section 5.1 that the velocity estimate produced by the second stage of a cascade filter should be as accurate, or more accurate, as that produced by a federated filter. The reason can be seen by considering the effect of the generalized K matrices (gain matrix of Equation (4.10)) below.

Different filter **structures** will calculate the K matrix in different manners, but the block

structure will be consistent. The observer gain matrix K may be written in block form as

$$K = \begin{bmatrix} K_{q,sf_{meas}} & K_{q,b_{meas}} & K_{q,v_{meas}} & K_{q,p_{meas}} \\ K_{v,sf_{meas}} & K_{v,b_{meas}} & K_{v,v_{meas}} & K_{v,p_{meas}} \\ K_{p,sf_{meas}} & K_{p,b_{meas}} & K_{p,v_{meas}} & K_{p,p_{meas}} \end{bmatrix}. \quad (4.20)$$

Here the $K_{q,sf_{meas}}$ block is the gain applied the accelerometer (specific force) measurement innovation to correct the attitude (quaternion) portion of the state, etc. In the federated-filter case the accelerometer is used as a kinematic-model input and is not used for observer feedback and this restriction requires the first column of the generalized K matrix to be zero so that

$$K_{federated} = \begin{bmatrix} 0 & K_{q,b_{meas}} & K_{q,v_{meas}} & K_{q,p_{meas}} \\ 0 & K_{v,b_{meas}} & K_{v,v_{meas}} & K_{v,p_{meas}} \\ 0 & K_{p,b_{meas}} & K_{p,v_{meas}} & K_{p,p_{meas}} \end{bmatrix}. \quad (4.21)$$

Attempting to use the accelerometer measurement as both a kinematic-model input and an observation measurement causes problems due to the introduction of correlated process noise and observation-measurement noise. This violates the assumptions used in the derivation of the Kalman filter. In practice the observation vector for federated filters does not include the accelerometer measurement and the K matrix does not include the left column of the generalized K matrix above.

For the cascaded filter, the cascade approach allows for the accelerometer to be used as an observation sensor for correcting/updating the attitude states and only the lower two blocks in the left row are zero,

$$K_{cascade} = \begin{bmatrix} K_{q,sf_{meas}} & K_{q,b_{meas}} & K_{q,v_{meas}} & K_{q,p_{meas}} \\ 0 & K_{v,b_{meas}} & K_{v,v_{meas}} & K_{v,p_{meas}} \\ 0 & K_{p,b_{meas}} & K_{p,v_{meas}} & K_{p,p_{meas}} \end{bmatrix}. \quad (4.22)$$

The effect of the zero blocks is easily seen in the covariance update of the EKF. Consider the EKF covariance predict/update equations [36],

$$P_{k|k-1} = F_{k-1}P_{k-1}F_{k-1}^T + Q_{k-1} \quad (4.23)$$

$$S_k = H_k P_{k|k-1} H_k^T + R_k \quad (4.24)$$

$$K_k = P_{k|k-1} H_k^T S_k^{-1} \quad (4.25)$$

$$P_{k|k} = P_{k|k-1} - K_k S_k K_k^T, \quad (4.26)$$

where P matrices represent state-estimate-error covariance, Q and R represent process and sensor-noise covariance, and F and H represent process and measurement-model Jacobians. The S matrix represents the estimated state-estimate covariance in the sensor space plus the sensor-error covariance. This matrix is often block diagonal, or sufficiently near block diagonal (see Appendix 4). A block diagonal form for S is denoted as

$$S \approx \begin{bmatrix} S_{sf} & 0 & 0 & 0 \\ 0 & S_b & 0 & 0 \\ 0 & 0 & S_v & 0 \\ 0 & 0 & 0 & S_p \end{bmatrix}, \quad (4.27)$$

where the sf , b , v and p subscripts denote the blocks related to the accelerometer measurement, magnetometer measurement, velocity measurement and position measurement, respectively. The upper-left-hand block of the rightmost term in Equation (4.26), which is the reduction in the attitude-state-error covariance produced by applying a set of observation-sensor measurements, is

$$[KSK^T](1,1)_{federated} = K_{q,bmeas} S_B K_{q,bmeas}^T + K_{q,vmeas} S_v K_{q,vmeas}^T + K_{q,pmeas} S_p K_{q,pmeas}^T \quad (4.28)$$

for the federated case, while for the cascaded case it is

$$[KSK^T](1,1)_{cascaded} = K_{q,sfmeas} S_{sa} K_{q,sfmeas}^T + K_{q,bmeas} S_B K_{q,bmeas}^T + K_{q,vmeas} S_v K_{q,vmeas}^T + K_{q,pmeas} S_p K_{q,pmeas}^T. \quad (4.29)$$

The inclusion of the accelerometer-related term, $K_{q,sfmeas} S_{sa} K_{q,sfmeas}^T$ in Equation (4.29), that does not appear in Equation (4.28), shows that the variance of the estimated attitude states is lower in the cascaded filter than in the federated filter due to the subtraction of the $K_k S_k K_k^T$ term in Equation (4.26). In addition, the accelerometer and magnetometer-related terms in Equations (4.28) and (4.29) will generally be of similar magnitude much larger than the GNSS-related terms

Table 4.2: Attitude estimation comparison for different filter structures at 3 wind/turbulence levels.

| | 1/4 * Standard sensor noise and drift | Standard sensor noise and drift | 4 * Standard sensor noise and drift |
|--|---------------------------------------|---------------------------------|-------------------------------------|
| Roll estimate error (degrees-r.m.s.) | | | |
| 1st stage of cascaded filter | 0.041 | 0.079 | 0.252 |
| Federated filter | 0.701 | 0.704 | 1.53 |
| Pitch estimate error (degrees-r.m.s.) | | | |
| 1st stage of cascaded filter | 0.047 | 0.068 | 0.220 |
| Federated filter | 0.750 | 0.879 | 1.95 |

so that the difference in attitude-state-error variance between a federated filter and a cascaded filter will be significant.

Figure 4.8 and Table 4.2 show simulation results for a single 16-state federated EKF and a 7-state (first stage of cascaded filter) EKF operating on simulated sUAS sensor data for flight in wind - see Appendix 4 for simulation parameters. Three cases of sensor errors (noise and drift rate) are shown with the center column being the standard error levels presented in the Appendix 4 and the left and right columns having the error levels decreased and increased by a factor of four, respectively. The two filters use the same covariance values for the gyroscope, accelerometer and magnetometer. While the second stage of the cascaded filter has slightly more error in the position and velocity estimate, Figure 4.8 shows that the first stage cascaded filter provides a significantly better attitude estimate than the federated filter.

If the relative error in the attitude estimate is of greater importance than the error in the position and velocity estimates, then a cascaded filter should be used. If the error in the attitude estimate and the velocity estimate are both of importance, as is the case for in-situ wind sensing from sUAS, then a combined approach may be used to minimize the error; the attitude estimate may be produced with an AHRS filter and the velocity estimate may be produced with a federated (INS) filter. However, as described in chapter 5 a cascade approach provides nearly as accurate a velocity estimate as an INS approach and may be sufficient.

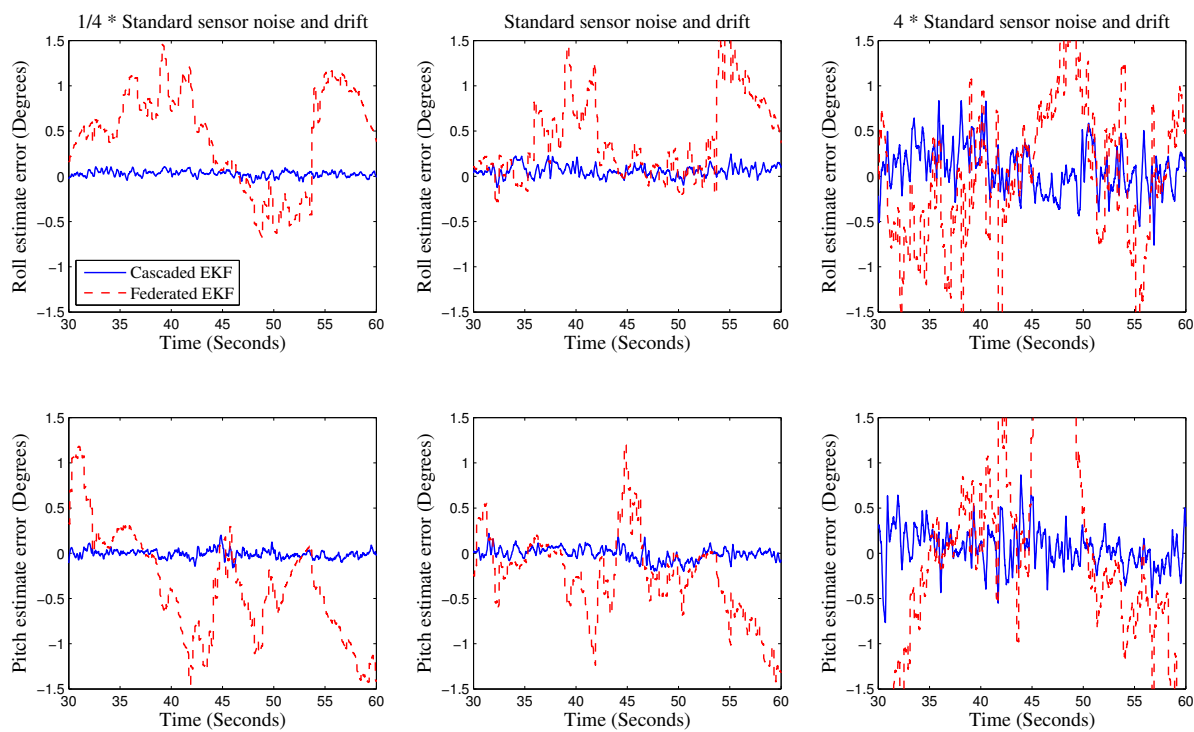


Figure 4.8: Attitude Estimation comparison for comparable federated and cascaded filters.

In addition to the fundamental difference in the attitude information content of observation-sensor measurements between federated and cascaded filters, several pragmatic considerations also make the cascaded approach more appropriate when a high-accuracy attitude estimate is needed.

- COTS GNSS receivers typically have measurement rates much lower than the other filter sensors (1-10 Hz versus 100-1000 Hz). This limits GNSS receivers effectiveness in contributing to a high-accuracy / high-update-rate attitude solution. I.E. $K_{q,v_{meas}}$ and $K_{q,p_{meas}}$ will necessarily be small.
- Federated filters will have a larger state count than either stage of a cascade filter. This can make it more difficult to construct and tune the filter so that it maintains stability while providing a high-accuracy estimate. The decoupling in a cascade filter significantly simplifies tuning.

Given this background the remainder of this paper considers only cascade **type** filters.

4.2.4 Specific-force correction methods

In an (AHRS) attitude filter the accelerometer and magnetometer measurements are used to provide a measure of two reference vectors — the Earth’s-magnetic-field vector and the gravity vector. Ideally, the accelerometer measures the gravity vector, which, along with the magnetic-field vector measurement, provides an absolute attitude reference to correct gyro drift. In reality, the accelerometer measures the specific force on a so-called ‘proof mass’. Translational acceleration, resulting from maneuvering flight or wind gusts, corrupts measurement of the gravity vector and must be removed. Failing to account for turning flight paths, which produce translational accelerations that are not transient, degrades the gyro-drift correction at low frequencies. To wit, consider that the specific force measured by an accelerometer in an aircraft in a coordinated 20° bank turn is 20° misaligned with the gravity vector. This misalignment is clearly a problem if 0.1° attitude estimation accuracy is desired. The question is how to separate gravity from acceleration in the

accelerometer measurement or, stated in an EKF context, how to estimate a priori the accelerometer measurement. For sUAS-based wind measurement this must be accomplished using low-cost sensors in the presence of significant gust accelerations.

The various filter **Structures** introduced in section 4.2.2 and considered herein use high-frequency gyroscope information and low-frequency accelerometer/magnetometer information to produce an attitude estimate. Both frequency-domain filters (e.g., complimentary filters) and statistical filters (e.g., EKF) effectively high-pass filter integrated gyroscope measurements and low-pass filter an attitude estimate based on accelerometer and magnetometer measurements. The high-pass and low-pass corner frequencies are either explicitly or effectively equal. If it were possible to set the low-pass corner frequency low enough that frequencies attributable to platform dynamics were above the corner, then translational accelerations that corrupt the gravity vector would be filtered out of the accelerometer measurement. However, unlike typical flight for manned aircraft, sampling regimes used by sUAS often result in a large percentage of their total flight time spent in turning flight — sUAS sampling flight often includes constant-rate turns that last for minutes. As a result, filtering out the effects of translational accelerations would require the corner frequency to be set very low (e.g., $O(\frac{1}{1000})$ Hz). Unfortunately, MEMS gyroscopes have bias drift and noise angular random walk that requires their high-pass corner to be set relatively high (e.g., $O(\frac{1}{10})$ Hz). (Current high-end MEMS gyroscopes drift on the order of 10° /hour [25], whereas lower-cost consumer and automotive-grade devices drift on the order of 1° /second.)

Since setting the low-pass corner frequency to remove the effect of translational accelerations is not possible, some other means of specific-force correction is required. In Figure 4.3, this functionality is incorporated in block \mathbb{F} . In the context of a frequency-domain filter, specific-force correction means taking into account the translational acceleration present in the accelerometer measurement when calculating the accelerometer-dependent error signal. In the context of an extended Kalman filter, specific-force correction means taking into account the translational acceleration present in the accelerometer measurement when calculating the a priori estimate of the accelerometer mea-

surement. The method of calculating an a-priori estimate of platform translational acceleration is an important (although perhaps the most overlooked) consideration in attitude estimator design for sUAS, and some methods frequently used provide poor performance in typical sUAS operating environments with non-trivial wind gusts.

4.2.4.1 Traditional specific-force correction methods

Barton [6] describes a specific-force-correction method commonly used in sUAS attitude estimators. Ignoring sensor error, the specific force measured by an accelerometer may be expressed as [6]

$$[s\vec{f}]_B = [{}^I\vec{a}_B]_B - [\vec{g}]_B = [{}^I\dot{\vec{v}}_B]_B - [\vec{g}]_B, \quad (4.30)$$

where $[{}^I\dot{\vec{v}}_B]_B$ denotes the derivative of the velocity vector taken with respect to the inertial frame. The inertial acceleration term in Equation (4.30) may be written in terms of the body-frame derivative, resulting in

$$[s\vec{f}]_B = [{}^B\dot{\vec{v}}_B]_B + [{}^I\vec{\omega}_B]_B \times [{}^I\vec{v}_B]_B - [\vec{g}]_B. \quad (4.31)$$

The first two terms on the right side of Equation (4.31) must be subtracted from the accelerometer measurement to determine the gravity vector. The second of these terms is commonly calculated based on ground speed, typically from GPS, and rotation rate from the gyros. Other implementations substitute airspeed for GPS ground speed based on a steady-wind assumption. These specific-force correction methods are hereafter referred to as rotation-rate based. Handling the first term on the right side of Equation (4.31) is where the problem arises. Barton [6] suggests the first term be ignored. An alternative is to consider velocity changes along the body x-axis (longitudinal direction) only and estimate this acceleration based on flight-condition changes using some combination of GPS velocity change, airspeed, thrust and flight-path angle. Euston et al. [19], for example, use a model based on angle-of-attack dynamics to estimate the body x-component, assuming that acceleration in the direction of flight may occur from time to time but accelerations

in other directions are minimal.

To illustrate the error introduced by these specific-force-correction methods, the first term on the right side of Equation (4.31) is expanded into three terms as

$$[{}^B\dot{\vec{v}}_B]_B = \left[\frac{{}^B d^2}{dt^2} \vec{r} \right]_B + \left[\frac{{}^B d^I}{dt} \vec{\omega}_B \right]_B \times [\vec{r}]_B + [{}^I\vec{\omega}_B]_B \times \left[\frac{{}^B d}{dt} \vec{r} \right]_B \quad (4.32)$$

and the second term on the right side of Equation (4.31) is expanded into two terms as

$$[{}^I\vec{\omega}_B] \times [{}^I\vec{v}_B] = [{}^I\vec{\omega}_B] \times \left[\frac{{}^B d}{dt} \vec{r} \right]_B + [{}^I\vec{\omega}_B] \times ([{}^I\vec{\omega}_B] \times [\vec{r}]_B), \quad (4.33)$$

where $[\vec{r}]_B$ represents the inertial position of the platform as components in the body frame. The three terms in Equation (4.32) are the translational, Coriolis, and Euler accelerations in the body frame and the two terms in Equation (4.33) are the Euler and centripetal accelerations [74].

Estimating the body x-component of the first term on the right side of Equation (4.32) based on changes in flight condition leads to errors caused by several types of gust-induced motion. The two most significant error types are those caused by translational accelerations and those caused by rotation rates not induced by turning flight. Errors introduced by wind-gust-induced translational acceleration are analyzed using Eqs. (4.31) through (4.33). Consider a wind gust that induces a 0.1g lateral acceleration in the body Y-direction without any rotation. Both terms on the right side of Equation (4.33) and the second and third terms on the right side of Equation (4.32) are zero; the first term on the right side of Equation (4.31) (or (4.32)) is the translational acceleration. However, rotation-rate-based methods fail to account for this translational acceleration. (It is not in the body x-direction and not associated with a change in flight condition.) The resulting specific-force correction error is equal to the wind-gust-induced translational acceleration, with a resultant error in the measured gravity vector of 5.7°.

Wind-gust-induced rotation rate and rotation acceleration also produce errors through the remaining terms in Equation (4.32) and (4.33). Consider a platform that travels in a straight horizontal line at a constant speed of 15 m/s (e.g., no translational acceleration) with no roll or

pitch; the yaw angle has a small gust-induced perturbation as shown in Figure 4.9 (left). In this situation, Equation (4.30) reduces to $[s\vec{f}]_B = -[\vec{g}]_B$, which implies $[{}^B\dot{\vec{v}}_B]_B = -[{}^I\vec{\omega}_B]_B \times [{}^I\vec{v}_B]_B$. A rotation-rate-based specific-force correction approximates the first term on the right side of Equation (4.31) as zero since the platform is at constant speed and pitch. The second term is non-zero due to the non-zero rotation rate (the yaw angle disturbance), is oriented in roughly the body Y-direction, and is not properly cancelled by $[{}^B\dot{\vec{v}}_B]_B$. The resulting measurement error in the gravity-vector direction is shown in Figure 4.9 (right).

In sUAS flight similar magnitude errors are typical due to gust-induced motion; yaw disturbances typically cause larger error than translational accelerations as seen in the simulation results of Section 4.2.4.3 below. The yaw disturbances used in this example are modest by comparison.

4.2.4.2 Alternate approach to specific-force correction

Rather than correcting the specific-force measurement based on approximating Equation (4.31), the inertial acceleration $[{}^I\ddot{\vec{v}}_B]_B$ in Equation (4.30) may be estimated directly using time-differenced GPS velocity measurements as

$$[{}^I\dot{\vec{v}}_B]_B = \mathbf{q} \otimes [{}^I\vec{a}_B]_I \otimes \mathbf{q}^* = \mathbf{R} [{}^I\vec{a}_B]_I, \quad (4.34)$$

where

$$[{}^I\vec{a}_B]_I \approx ([{}^I\vec{v}_B]_{I,t_k} - [{}^I\vec{v}_B]_{I,t_{k-1}})/(t_k - t_{k-1}), \quad (4.35)$$

\otimes is the quaternion product operator, $*$ is the quaternion conjugate operator, \mathbf{q} is the quaternion attitude estimate, and \mathbf{R} is the equivalent rotation-matrix parameterization of the attitude estimate.

It is frequently assumed that COTS GPS receivers do not have sufficient accuracy or measurement frequency to be used to estimate platform acceleration. However, comparison of the velocity measurement provided by Ublox-7 series GPS receivers [80] with the velocity measurement from a software-defined GPS receiver using a rubidium (atomic) clock as a timing reference suggests that the COTS receivers' velocity measurement error is merely a few centimeters per second. Also,

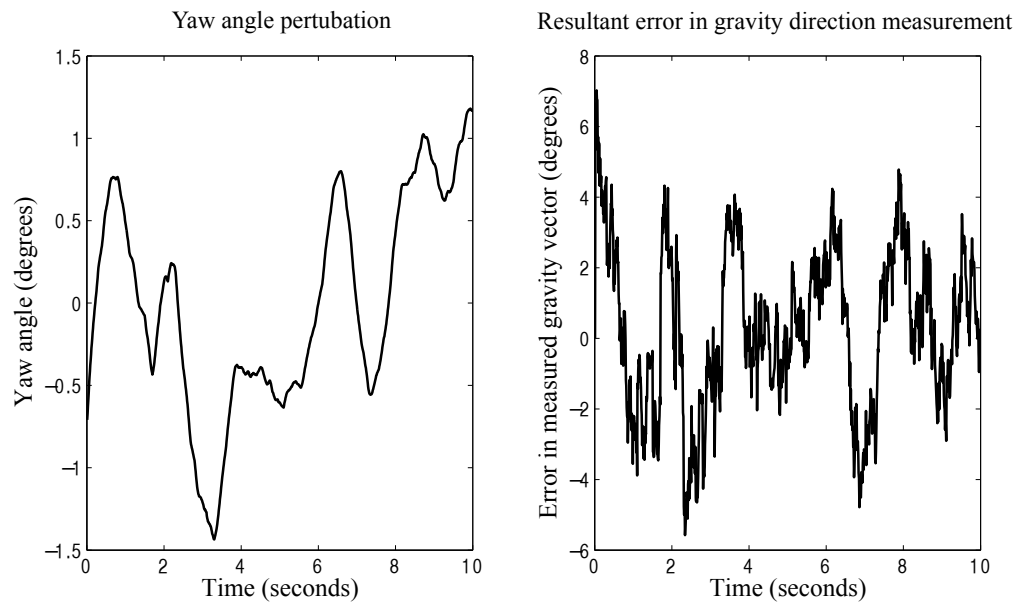


Figure 4.9: Error introduced by traditional specific-force correction methods (Equation (4.31)) when yaw angle is perturbed.

since the specific-force correction is applied to the accelerometer measurement to recover gravity for low-frequency gyro-drift correction, the GPS measurement rate is in fact relatively fast.

4.2.4.3 Comparison of approaches

Simulations of attitude estimators comparing the two approaches to specific-force correction show that estimating translational acceleration with time-differenced GPS velocity measurements provides significantly better performance than rotation-rate-based (traditional) methods when flying in wind. GPS measurements were simulated at a 5 Hz update rate and with a velocity error standard deviation of 5 cm/s for the horizontal axes and 10 cm/s for the vertical axis. A description of the simulation engine and parameters is provided in Appendix 4. Even at a 5 Hz GPS output rate, this approach reduces pitch and roll estimate error by nearly an order of magnitude versus Barton's approach [6] for sUAS flight in wind. Figure 4.10 and Table 4.3 show simulation results for two 7-state AHRS EKF filters that are identical except for the accelerometer measurement model; one models translational acceleration using time-differenced GPS velocity measurements and the other uses the rotation-rate-based calculation. The simulation is performed with a nominal sUAS airspeed of 12 m/s, three mean wind speeds (3 m/s, 6 m/s and 12 m/s), and wind turbulence modeled using a discrete Dryden wind gust model (MIL-F-8785C) [70]. Figure 4.10 shows results for a 30-second segment of the simulation and Table 4.3 shows results averaged over a 2000-second timespan. These simulation results indicate that the desired level of attitude-estimation accuracy for wind measurement, $O(0.1^\circ)$, is realistic with proper sensor fusion and contemporary sUAS-suitable sensors.

Note that the simulations in sections 4.2.3 and 4.2.2, and the results shown in Figures 4.6 and 4.8 and Tables 4.1 and 4.2, were based on a perfect estimate of translational acceleration. Only in section 4.2.4 has the effect of error in the translational acceleration estimate been included in the simulations.




Figure 4.10: Attitude estimation comparison for filters with different accelerometer measurement models.

Table 4.3: Attitude estimation comparison for filters with different accelerometer measurement models at 3 wind/turbulence levels.

| | 3 m/s wind | 6 m/s wind | 12 m/s wind |
|--|------------|------------|-------------|
| Rotation-rate-based correction | | | |
| Roll estimate error (degrees-r.m.s.) | 0.78 | 1.28 | 2.59 |
| Pitch estimate error (degrees-r.m.s.) | 0.64 | 1.39 | 3.31 |
| Time-differenced GPS correction | | | |
| Roll estimate error (degrees-r.m.s.) | 0.14 | 0.14 | 0.16 |
| Pitch estimate error (degrees-r.m.s.) | 0.14 | 0.14 | 0.17 |




Figure 4.11: Errors in estimated translational acceleration.

The reason that rotation-rate-based specific-force correction performs poorly for sUAS flown in wind can be seen in the error in the specific-force correction. Figure 4.11 shows the error in the estimated translational acceleration that is introduced into the attitude solution for the three wind conditions considered above. This figure shows the error relative to the true translational acceleration in the body Y-direction for each case. Higher winds (and accompanying turbulence levels) introduce larger errors when using a rotation-rate-based calculation but even modest amounts of wind produce enough perturbation motion to cause a significant difference between the results. Using time-differenced GPS velocity measurements results in lower attitude-estimation error; the error is relatively insensitive to the turbulence level.

Figure 4.12 shows the rotation-rate-based specific-force correction error for a 5-second segment of the 12 m/w wind case. The error component introduced by wind-gust-induced translational acceleration and the error component introduced by wind-gust-induced rotation rate are plotted as separate lines. The main contribution to the error when using the rotation-rate-based correction is misinterpreting wind-gust-induced rotation rate as rotation due to turning flight.

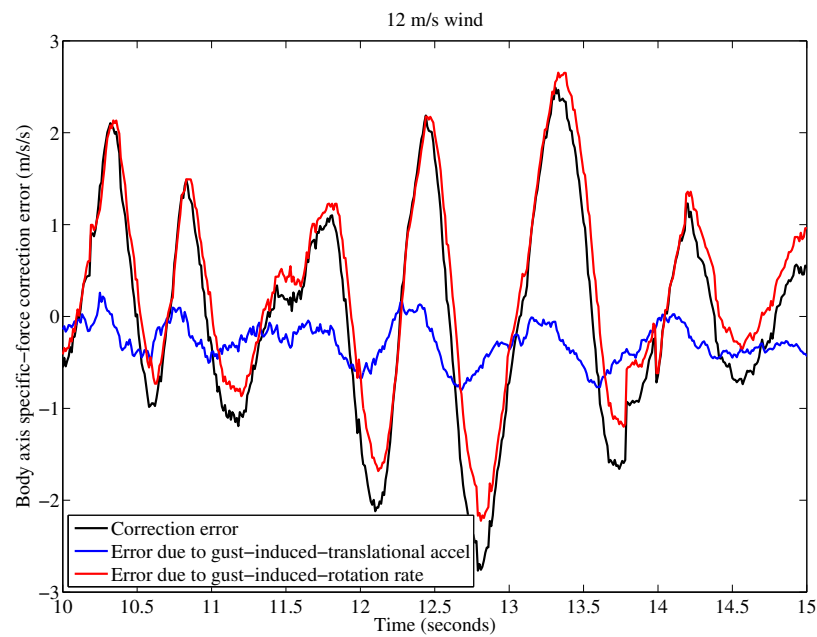


Figure 4.12: Error components in estimated translational acceleration using the rotation-rate-based method.

4.2.5 Expected sUAS attitude-estimation accuracy with contemporary sensors

Recent improvements on MEMS sensors has significantly reduced the size and power consumption of these sensors, expanded interface choices and improved on-board calibration for temperature effects, etc., but noise performance improvements have been modest. Regardless, the best MEMS sensors are quite good considering their size and power consumption. The standard sensor error levels used for the simulation results presented above were roughly based on the Analog Devices ADIS16448 Ten Degrees of Freedom Inertial Sensor. This sensor, while relatively expensive in comparison to other sensors used in contemporary sUAS systems, is still not the best (or most expensive) sensor available in this class. For example the Analog Devices ADIS16488A offers significantly better performance, but at near triple the cost. Some newer, low-cost MEMS sensors, such as the Maxim Integrated MAX21100 6+3 DoF IMU, are approaching the Analog Devices ADIS16448 in performance, making the Analog Devices ADIS16448 a reasonable choice for a current benchmark. Some of the key sensor parameters used for simulation were:

- Gyroscope angular random walk - $1.0^\circ / \sqrt{\text{hour}}$
- Accelerometer noise density - $0.25 \text{ mg} / \sqrt{\text{Hz}}$
- Magnetometer noise density - $0.4 \text{ mgauss} / \sqrt{\text{Hz}}$
- GPS velocity error - $0.05/\sqrt{3} \text{ m/s r.m.s. north \& east, } 0.10/\sqrt{3} \text{ m/s r.m.s. vertical}$

Sections 4.2.2, 4.2.3 and 4.2.4 produced specific recommendations for constructing a sUAS attitude filter. The simulation results in section 4.2.4 use an attitude filter incorporating all these recommendations. It uses a 7 state extended Kalman filter, with the states being the quaternion elements and the gyroscope bias, and employs time-differenced GPS measurements for translational-acceleration estimation. Simulation of this filter performing attitude-estimation for a sUAS flying in wind and using sensor errors representative of contemporary sensors as described above show that roll and pitch estimate error can be less than 0.2° r.m.s. even when flying in high winds.

Section 3.3.1 included analysis of the sensitivity of the vertical-wind measurement to the attitude estimate. It was shown that the vertical-wind measurement is sensitive to the pitch estimate and using representative values for the relative-wind speed and direction and the airframe attitude, attitude rate and inertial velocity this sensitivity is 12.0 m/s / rad. The error in the vertical-wind measurement induced by a 0.2° pitch estimate error is then 0.042 m/s. At moderate wind speeds (e.g., 6 m/s) the pitch estimate error for the simulated case is expected to be 0.014° and the induced error in the vertical-wind measurement is then 0.029 m/s.

Chapter 5

sUAS velocity estimation

As described in section 3.3.1, wind measurement from sUAS has a 1-to-1 sensitivity to error in the estimation or measurement of the sUAS velocity in the inertial (earth-fixed) frame. Wind measurement accuracy of $O(0.01m/s)$ requires sUAS platform velocity estimation accuracy of $O(0.01m/s)$. This chapter considers the question of what sensors and sensor fusion algorithms should be used for sUAS velocity estimation to support sUAS-based wind measurement.

A number of sensors described in previous chapters may be used in sUAS velocity estimation. GPS receivers provide a direct estimate of the platform velocity in the inertial frame, but GPS receivers suitable for sUAS have relatively slow output rates (typically < 10 Hz). Accelerometer measurements may be integrated to yield velocity, but the body-frame accelerometer measurement must be rotated to the inertial frame and integrated noise and bias create problematic error levels within a period of minutes or seconds. Other sensors such as optical systems may also be employed, but often have poor accuracy, particularly in the vertical direction.

In both federated filters and the second stage of cascaded filters, described in section 4.2.3, accelerometer measurements, rotated with the current attitude estimate, are integrated to yield velocity of the platform in the inertial frame. This integrated velocity is subject to error that increases with time and with contemporary MEMS accelerometers the integrated velocity estimate will drift outside of the desired accuracy in a short time. Consequently, other sensors are needed to provide a feedback correction; the situation is directly analogous to attitude-estimation sensor

fusion described in section 4.2.2. GPS receivers are the primary sensor used for this purpose.

Equation (4.20) introduced a general block structure for the gain matrix, K , for state-estimation systems using gyroscopes, accelerometers, magnetometers and GPS receivers. Other sensors may be included in the sensor fusion, with an additional column added to the block structure of K for each additional sensor. The structure of the K matrix determines the relevance of a particular sensor to velocity estimation. Gain terms in the second row of Equation (4.20) provide a direct feedback path for sensor measurements to affect velocity states. However, Equation (4.5) shows that the attitude estimate is used in propagation of the velocity using accelerometer measurements and, consequently, gain terms in the first row of Equation (4.20) provide an indirect feedback path for sensor measurements to affect velocity states. This indirect feedback improves the attitude estimate, which then reduces error introduced through integration of the rotated accelerometer measurements at future time-steps. By contrast direct feedback improves the velocity estimate at the current time-step based on the current sensor measurements.

Use of different filter **structures** determine whether a particular sensor has a direct feedback path to affect the velocity estimate. To illustrate this point consider the gain block $K_{v,bmeas}$, which provides a direct feedback path for magnetometer measurements to affect the velocity estimate. The velocity states have no observability from the magnetometer measurements; the velocity states do not appear in the second row of Equation (4.8), the mapping function of state to predicted sensor measurements. For filter **structures** like the complimentary filter or gradient descent filter the block $K_{v,bmeas}$ will be zero due to the lack of observability of velocity states from magnetometer measurements. For filter **structures** like the EKF, however, the block $K_{v,bmeas}$ may theoretically be non-zero due to non-zero covariance between the attitude states and the velocity states. In practice the relative gain from blocks like $K_{v,bmeas}$, that relate to sensors without observability of the velocity states, will produce much smaller feedback signals than those from sensors, like GPS, with direct observability of the velocity states. For the standard gyro/accelerometer/magnetometer/GPS sensor set considered herein, the GPS receiver is primarily responsible for the feedback correction

of the velocity estimate and the accuracy of the GPS receiver velocity estimate is a limiting factor for accuracy of the resulting sensor-fusion velocity estimate.

5.1 Cascaded versus Federated filters for velocity estimation

If the attitude and velocity estimation problem were linear with respect to the dynamic model and observation sensor mappings, and the sensor errors were Gaussian, then a Kalman filter would provide an optimal estimate of the aircraft state. Some authors (for example [11]) infer that a federated EKF should provide a near-optimal state estimate, with the linearization of the nonlinear dynamics and observation sensor models introducing the sub-optimality, and therefore a federated (INS) filter is a better choice for attitude/velocity estimation. For the problem of sUAS state estimation, neither the dynamic model nor the observation sensor mappings are linear, the sensor errors may not be Gaussian and, further, the sensors often operate at different measurement rates. In particular, the GPS measurements are typically processed separately from accelerometer and magnetometer measurements due to the relatively low measurement rates available from contemporary COTS GPS) operate at different measurement rates. It was shown in section 4.2.3 that the first stage of a cascaded filter produces a more accurate attitude estimate than a federated filter. It is further shown below that the velocity estimate produced by the second stage of a cascade filter should be more accurate than that produced by a federated filter. Accordingly a cascaded filter is recommended for sUAS-based wind measurement systems.

The EKF framework is used again to examine the difference between the velocity estimate produced by a federated filter and the second stage of a cascaded filter. Equations 4.23 through 4.26 are repeated here for convenience.

$$P_{k|k-1} = F_{k-1}P_{k-1}F_{k-1}^T + Q_{k-1} \quad (4.23)$$

$$S_k = H_k P_{k|k-1} H_k^T + R_k \quad (4.24)$$

$$K_k = P_{k|k-1} H_k^T S_k^{-1} \quad (4.25)$$

$$P_{k|k} = P_{k|k-1} - K_k S_k K_k^T, \quad (4.26)$$

Now working in a similar fashion to section 4.2.3 and using block representations of the various matrices, and collecting position and velocity terms into single blocks, represent the matrices of Equations 4.23 through 4.26 for the federated filter as:

$$P_{fed} = \begin{bmatrix} P_q & P_{q,vp} \\ P_{q,vp} & P_{vp} \end{bmatrix}, \quad (5.1)$$

$$Q_{fed} = \begin{bmatrix} Q_q & 0 \\ 0 & Q_{vp} \end{bmatrix}, \quad (5.2)$$

$$F_{fed} = \begin{bmatrix} F_{q,q} & 0 \\ F_{vp,q} & F_{vp,vp} \end{bmatrix}, \quad (5.3)$$

$$H_{fed} = \begin{bmatrix} H_{b,q} & 0 \\ 0 & H_{GPS,vp} \end{bmatrix}. \quad (5.4)$$

The second stage cascaded filter has only velocity and position as state and uses the matrices P_{vp} , Q_{vp} , $F_{vp,vp}$, and $H_{GPS,vp}$. At this point it should be noted that while the federated filter covariance propagation of Equation 4.23 will propagate the full state-error covariance matrix including both attitude and velocity states, Equations 4.24 through 4.26 will be processed separately for magnetometer and GPS measurements (using sub-matrices of the full, generalized P , H , R , S , and K matrices) due to these sensors having different measurement rates. For this reason the off-diagonal blocks of

$$P_{fed:k|k-1} = \begin{bmatrix} F_{q,q}P_qF_{q,q} + Q_q & F_{q,q}P_qF_{vp,q} + F_{q,q}P_{q,vp}F_{vp,vp} \\ F_{q,q}P_{q,vp}F_{vp,vp} + F_{q,vp}P_{vp}F_{vp,vp} & F_{vp,q}P_qF_{vp,q} + 2F_{vp,vp}P_{q,vp}F_{vp,q} + F_{vp,vp}P_{vp}F_{vp,vp} + Q_{vp} \end{bmatrix} \quad (5.5)$$

are not used in computation of the Kalman gain or covariance update. To reiterate, Equations 4.23 and 5.5 reflect the increase in covariance due to error introduced in the prediction step. In the case of velocity estimation this is due to error in the accelerometer measurements (associated with the term Q_{vp}) and due to error in the current attitude estimate (associated with the terms $F_{vp,q}P_qF_{vp,q}$

and $2F_{vp,vp}P_{q,vp}F_{vp,q}$). The term $F_{vp,vp}P_{vp}F_{vp,vp}$ represents the rotation of the covariance due to the incremental rotation in attitude.

By contrast, the covariance propagation for the second stage cascaded filter is simply:

$$P_{cas:k|k-1} = F_{vp,vp}P_{vp}F_{vp,vp} + Q'_{vp} , \quad (5.6)$$

where the Q matrix Q'_{vp} is notated with a ' to indicate that it differs from Q_{vp} (of the federated filter); a slightly different process noise model is required for the second stage cascaded filter. The calculated state-error covariance update for the cascaded filter does not have the explicit additive elements increasing the velocity-state-error covariance due uncertainty in the attitude estimate noted above for the federated filter. Rather, in the second stage cascaded filter the attitude estimate is treated as an input parameter and the process noise term Q'_{vp} must account for this source of estimation error. However, as the first stage cascaded filter produces a more accurate attitude estimate than the federated filter, the increase in Q'_{vp} in the second stage cascaded filter versus Q_{vp} in the federated filter should be less than $F_{vp,q}P_qF_{vp,q} + 2F_{vp,vp}P_{q,vp}F_{vp,q}$ in Equation 5.5. Finally, noting that H_{fed} is diagonal the covariance reduction for the velocity states through Equation 4.26 can be seen to be equal for the second stage cascaded filter and the federated filter. Hence, the cascaded filter produces a velocity estimate with lower state-error covariance.

This result is not unexpected if considering that both filters propagate the velocity estimate using the accelerometer measurement rotated by the current attitude estimate. The cascaded filter has a more accurate attitude estimate and therefore incurs less error in the prediction step of the filter cycle. When the observation sensor measurements are processed separately, as is realistic for this problem due to low measurement rates of GPS, the filter does not take full advantage of potential cross-correlations and the correction applied to the velocity estimate in the correction step of the filter cycle is the same for both filter types.

5.2 GPS velocity estimation for sUAS

Chapter 4 described how sUAS attitude estimation is sensitive to the sUAS motion environment. GPS velocity estimation is similarly sensitive to the sUAS motion environment. Documentation of commercial off-the-shelf GPS receivers is limited and analytical analysis of their velocity estimation performance with respect to different motion environments is not possible without proprietary information. However, reasons for COTS GPS receiver velocity-estimate-error sensitivity to motion may be inferred from general GPS system operating principles and performance of a particular COTS GPS receiver may be quantified through testing.

5.2.1 GPS receiver operating principles

The GPS system is the most used global navigation satellite system. Competing systems (GLONASS, Compass, Galileo) operate on similar principles with respect to velocity estimation and some newer GPS receivers are able to concurrently use satellite signals from the GPS system satellites as well as satellite signals from one or more other GNSS system. This description of GPS velocity estimation applies to low-cost COTS GNSS receivers regardless of what satellite signals are used. The GPS system includes a constellation of satellites, a ground control segment managing the satellites, and user receivers. The satellites transmit signals with encoded information about the satellite orbits. Receivers are able to decode this information and use it to compute the satellite locations in the earth-centered-earth-fixed (ECEF) reference frame. The signals also include encoded information about the time of the signals' transmission. GPS receivers track the signals received from multiple satellites and measure the relative arrival time of these signals. When four or more satellite signals are tracked the receiver performs a least-squares optimization of the tri-lateration problem to determine the receiver position in the ECEF frame as well as the receiver's clock offset.

In addition to measuring the relative arrival time of the satellite signals a receiver also makes

a measurement of the carrier frequency of each satellite signal. This occurs in one of two tracking loops for each satellite channel. When a signal is first acquired the signal is tracked in a frequency-lock loop (FLL). This loop provides a measure of the signal's carrier frequency. The difference between the known frequency at which the satellite transmitted the signal and the frequency of the received signal is the Doppler frequency due to the relative velocity between the satellite and receiver. The Doppler frequency is proportional to the range rate - the component of the velocity of the satellite relative to the receiver along the line-of-sight vector between the satellite and receiver. This may be modeled as ([51] p. 218)

$$v^k = (\vec{v}^k - \vec{v}) \cdot \vec{l}^k + \dot{b} + \epsilon^k, \quad (5.7)$$

where v^k is the range rate between the k'th satellite and the receiver, \vec{v}^k is the ECEF velocity of the k'th satellite, \vec{v} is the ECEF velocity of the receiver, \vec{l}^k is the line-of-sight (unit) vector between the receiver and k'th satellite, \dot{b} is the receiver-clock drift rate and ϵ^k is the combined error due to Doppler measurement error of the k'th satellite's signal as well as error due to ionosphere, troposphere and multi-path induced changes to the satellite signal during the measurement interval. Error due to satellite-clock drift rate is negligible. Error in the Doppler measurement due to ionospheric and tropospheric delays and multi-path are generally small [51]. Doppler measurement error is the primary limitation to accuracy of the velocity estimate. Using a minimum of four satellites' Doppler measurements a least-squares optimization yields an estimate of the three components of the receiver ECEF velocity and the receiver-clock drift rate. Using the receiver's position estimate the velocity estimate may be rotated to the local north-east-down (NED) frame. The accuracy of the three components of the receiver NED velocity may be calculated using the accuracy of the Doppler measurements and the satellite/receiver geometry.

After a satellite signal has been acquired and tracked in a FLL, it may be also be tracked with a phase-lock loop (PLL). While a FLL tracks the carrier frequency, a PLL tracks the phase of the carrier signal and the phase change over a measurement interval provides a more accurate measurement of the Doppler frequency than is provided by a FLL. GPS signals include a binary

ranging code, used as a time-mark to determine relative arrival time, and binary navigation data, used to communicate satellite ephemeris information. These two binary signals are combined into a single binary signal and modulate the carrier signal using binary phase-shift keying. While the ranging code is known a priori, the navigation data is not. The result is that the phase of a received satellite signal may change phase by 180° during a measurement interval. For this reason a coherent phase detector cannot be used in the phase-tracking loop; GPS receivers typically employ a Costas discriminator in the PLL [51]. A PLL with Costas discriminator will lose lock if the phase measurement deviates more than $\pi/2$ radians from the true phase. The maximum error in a Doppler frequency measurement made with a Costas-discriminator-based PLL occurs if the phase error is $\pi/2$ at the beginning of the measurement period and $-\pi/2$ at the end of the measurement period, or visa versa. At a 5 Hz measurement rate (0.2 second measurement period), the corresponding Doppler frequency measurement error for a GPS L1 frequency (1575.42 Mhz) signal corresponds to a range-rate error of 0.48 *m/s*. Typically when a PLL has lock the phase error will be much smaller than this maximum error. Misra and Enge [51] state that PLLs are likely to lose lock if the 1- σ phase error is $\pi/12$. If the PLL is consistently maintaining lock then range-rate accuracy of $O(1 \text{ cm/s})$ or better is expected at a 5Hz measurement rate. The key to GPS velocity-estimation accuracy of $O(1 \text{ cm/s})$ is for the receiver to have phase-lock on signals from a sufficient number of satellites. The minimum number depends on the current satellite geometry and is always greater than or equal to four.

5.2.1.1 First and second-order tracking loops

Both frequency-lock and phase-lock tracking loops use a feedback signal to control a numerically-controlled oscillator which generates a replica carrier signal to demodulate the received signal. The discriminator producing the feedback signal in a frequency lock loop produces an appropriate feedback signal over a much wider range of frequency error than the Costas discriminator in a PLL, which must maintain the phase error between $\pm\pi/2$ radians maximum ($\pm\pi/12$ radians typical).

A phase-lock tracking loop can be modeled as an linear analog control system - the Costas discriminator output is near linear for small phase errors. Figure 5.1 shows first and second-order phase-lock tracking loop models. The phase change in a signal over a period of time is equal to the time integration of the signal frequency; the operation of the numerically-controlled oscillator and discriminator over a discrete time step is modeled as an integration block. Simple proportional feedback to the numerically-controlled oscillator yields a first-order closed-loop system. Using this feedback signal the tracking loop will have a steady-state error in the replica-carrier-signal phase change for a signal with a constant (non-zero) Doppler frequency. A second-order tracking loop is able to correct for a non-zero Doppler frequency and will have no steady-state error for a signal with a constant non-zero Doppler frequency. Behavior of second-order linear systems is well known and the feedback gain values k_1 and k_2 determine the natural frequency and damping ratio of the response. Higher gain values allow for faster response that allow better tracking of changing signals, but increase the noise bandwidth of the system resulting in more noise in the Doppler measurement. Accelerating motion produces a changing Doppler frequency. If the receiver is under constant acceleration then a steady-state error results and is given by

$$\epsilon_{\Delta\phi} = \frac{a\pi}{2\lambda B^2} \text{ (radians)}, \quad (5.8)$$

where a is the acceleration, λ is the carrier wavelength, and B is the tracking loop bandwidth. At the GPS L1 frequency of 1.57542 GHz the wavelength is 19.0 cm. If the maximum phase error that can be tolerated without losing lock is $\pi/12$ radians, then tracking a signal corresponding to 1 g acceleration requires a tracking loop bandwidth of 17 Hz while tracking a signal corresponding to 1/4 g acceleration requires a tracking loop bandwidth of 8.8 Hz. The GPS designer must make a tradeoff between tracking performance (higher bandwidth) and noise performance (lower bandwidth).

Just as a second-order tracking loop can track a signal with constant Doppler frequency (corresponding to a constant velocity) without steady-state error, a third-order tracking loop can track a signal with linearly varying Doppler frequency (corresponding to a constant acceleration)

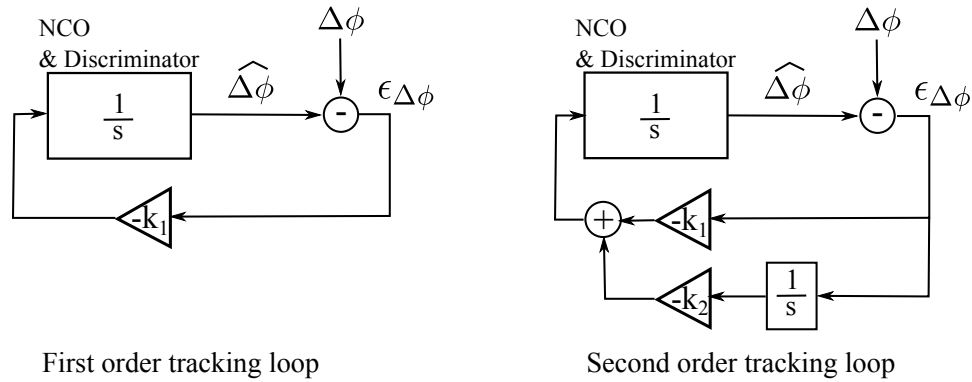


Figure 5.1: Phase-lock tracking loop models

without steady-state error. In the case of a third-order tracking loop a calculation similar to the one above may be performed to calculate the maximum amount of jerk (rate-of-change of acceleration). However third-order tracking loops are rarely employed due to stability concerns. sUAS flying in wind are subject to random wind-gust-induced motion with non-zero acceleration and jerk. Analysis of the resulting errors and loss-of-lock conditions for COTS GPS receivers is not possible as details of their tracking loop characteristics are not known. However, as it is expected that a COTS GPS receiver's phase-lock tracking loop is no more sophisticated than a second-order loop as described above, accelerations above some threshold are expected to cause loss-of-lock.

5.2.2 Testing of U-blox 6T GPS receiver while in accelerating motion

COTS GPS receivers are for the most part 'black box' instruments, with little information available about the specific methods and algorithms used by a particular receiver. User parameters give some limited insight into a receiver's design, but generally leave much unknown. Experimental evidence is examined to determine the velocity-estimate accuracy provided by a COTS GPS receiver experiencing different levels of accelerating motion. A U-blox 6T GPS receiver module [79] was selected for testing for three reasons:

- (1) This receiver is from a series of receivers often used in sUAS. Other research projects at

the University of Colorado use receivers from this series. Characterization of this receiver provides information about other receivers in the U-blox 6 series with value beyond the work herein.

- (2) This receiver will output raw pseudo-range measurements, Doppler (FLL) measurements and carrier-phase-change (PLL) measurements. Additionally the receiver output includes a quality metric for each channel indicating if the tracking loop is maintaining frequency-lock and/or phase-lock, and a carrier-to-noise measurement for each channel. This information may provide insight into the receiver carrier-phase-tracking-loop characteristics.
- (3) An external velocity estimate formed with the carrier-phase-change measurements may be compared to the GPS-receiver's velocity estimate (which is filtered). Nearly all commercial receivers apply filtering to the position and velocity solutions using a Kalman filter and a (usually undocumented) dynamics model. The dynamics model is chosen to represent the motion expected for the platform on which the GPS receiver will be used, which often is an automobile or a pedestrian. Some receivers (including the U-blox 6T) allow selection from a list of dynamic models and these lists often include aircraft, but the dynamic model itself is still unknown and the use of the filter may be inappropriate.

A swing-arm test set-up, shown in Figure 5.2, provides a known and repeatable motion environment. The desired rotation rate of the swing-arm may be chosen and the resulting rotation rate is generally constant with some variation due to the wind-gust forces on the arm. This circular motion with near-constant angular velocity results in the receiver experiencing near-constant-magnitude acceleration but changing acceleration direction. The effect on a particular satellite's Doppler signature is a sinusoidal modulation of the Doppler frequency with the amplitude of the frequency modulation dependent on the satellite's elevation and the phase of the modulation dependent on the satellite's azimuth. The swing-arm set-up allows the GPS output to be recorded and matched to the actual GPS-receiver velocity in post-processing. Data was collected with the arm stationary and rotating at various rates up to approximately 1 rotation per second resulting in the receiver

experiencing no acceleration and acceleration levels up to approximately 40 m/s^2 .

For the test case with no acceleration an overwhelming majority (99.5%) of the individual measurements (for a particular satellite channel and particular measurement epoch) have a quality indicator which indicates that the tracking loop has strong carrier-phase-lock throughout the measurement period. However for the data sets with accelerating motion there are varying lock-quality indications for individual measurements; some data sets have a significant percentage of measurements for which the tracking loop did not have phase-lock. Direct correlation of acceleration or acceleration rate to lock type for individual measurement epochs failed to show a significant correlation. For modest levels of acceleration the carrier-to-noise ratio is a better predictor of lock quality than acceleration. The carrier-to-noise ratio is not considered in the analysis of section 5.2.1.1, but noise in the input signal acts as a disturbance input to the tracking loop and a low carrier-to-noise ratio results in both measurement noise and, as the noise level increases relative to the signal level, a loss of tracking-loop lock. The combination of Doppler-frequency dynamics and noise act together to cause loss-of-lock in the phase-lock tracking loop.

Figure 5.3 shows statistics for 0.39 Hz and 0.81 Hz arm-rotation-rate data-sets. The histograms on the left show lock type versus acceleration. The horizontal axis represents histogram bins at different acceleration levels where the acceleration level represents the projection of the receiver's instantaneous acceleration in the receiver to satellite baseline direction during a particular epoch for a particular satellite channel. The four colored bars represent the number of measurements falling in each acceleration bin for each of four reported lock types - (dark blue) frequency lock only, (light blue) phase lock with slip during the measurement epoch, (yellow) phase lock, or (brown) strong phase lock. The vertical axis is the number of counts scale for the colored bars. It can be seen that phase lock is only achieved at lower acceleration levels; in the lower left histogram there are zero measurement counts for measurements with phase lock or strong phase lock at acceleration levels above 15 m/s^2 . However there are many measurement samples where phase lock is lost even at low acceleration levels. The histograms on the right show lock type versus




Figure 5.2: GPS receiver motion test set-up

carrier-to-noise ratio. Here it can be seen that high carrier-to-noise ratio is strongly correlated with phase lock and low carrier-to-noise ratio is strongly correlated with loss of phase lock. So, while accelerating motion has a significant effect on tracking loop performance, knowledge of the acceleration or acceleration rate alone (at modest acceleration levels) is not a good predictor of tracking loop performance. Satellite channel carrier-to-noise levels have a significant effect and are a better predictor of tracking loop performance than acceleration except at high acceleration levels (e.g., above 5 m/s²).

Figure 5.4 shows histograms of the velocity-estimate error for a dataset where 87% of the tracking loop measurements came from the (locked) PLL and 13% came from the FLL. The left-side histogram shows the velocity-estimate error for the GPS-receiver-produced velocity estimate. The center histogram shows the velocity-estimate error for a velocity estimate produced external to the GPS receiver using the receiver's raw carrier-phase-change and Doppler measurements. This externally produced velocity estimate has less than half the error seen in the GPS-receiver's velocity estimate. The expected cause is filtering applied to the velocity estimate by the GPS receiver. The GPS-receiver's configuration includes a dynamic model parameter; the 'airborne <4g' dynamic model is used for this test, which is the most dynamic choice available and which should apply the least filtering. Testing with other dynamic model selections showed higher error levels for this test.

The velocity estimate reflected in the center histogram of Figure 5.4 uses all the available carrier-phase-change and Doppler measurements. Since the carrier-phase-change measurements are higher quality than the Doppler measurements it is of interest to see if the Doppler measurements have a positive or negative contribution to the velocity-estimate error. The right-side histogram shows the velocity-estimate error for a velocity estimate produced external to the GPS receiver using the receiver's raw carrier-phase-change measurement but not the Doppler measurements. This results in a smaller number of measurements being used at each measurement epoch. However, as the FLL Doppler measurements are of relatively poor quality, dropping the FLL Doppler measurements and using only the PLL carrier-phase-change measurements results in a slightly better

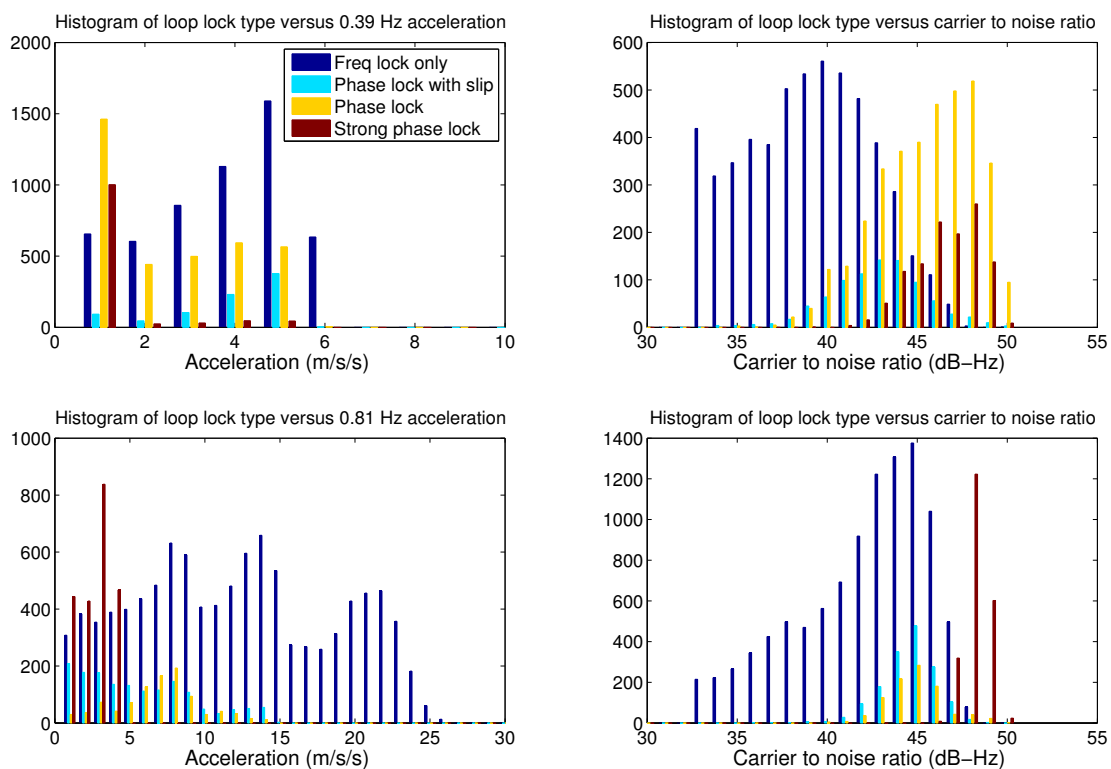


Figure 5.3: Histograms of Lock type versus acceleration and carrier-to-noise ratio




Figure 5.4: Distributions of GPS velocity-estimate error for data set with 87% phase-lock measurements

velocity estimate. Note that the relatively small difference between these two methods may be attributed to using a weighted estimate based on channel carrier-to-noise ratio; the FLL Doppler measurements typically arise in channels with low carrier-to-noise ratio so their impact on the estimate is small.

Figure 5.5 shows histograms of the velocity-estimate error for a dataset where only 28% of the tracking loop measurements came from the (locked) PLL and 72% came from the FLL. The left-side histogram shows the velocity-estimate error for the GPS receiver produced velocity estimate and the right-side histogram shows the velocity-estimate error for a velocity estimate produced external to the GPS receiver using the receiver's raw carrier-phase-change and Doppler measurements. In this case no velocity estimate was produced using PLL carrier-phase-change measurements only as there were not 4 PLL channel measurements available at most measurement epochs. It is interesting to note that in the data sets used for Figures 5.4 and 5.5 the carrier-to-noise ratio statistics are very similar. The poorer tracking loop performance for the second case is due to the more dynamic motion (higher swing-arm rotation rate) used; the arm rotation rate is 74% faster resulting in three times greater (10.9 m/s^2) acceleration.

Figure 5.6 shows histograms similar to those of Figure 5.5. In this case the GPS receiver was at rest (zero acceleration).

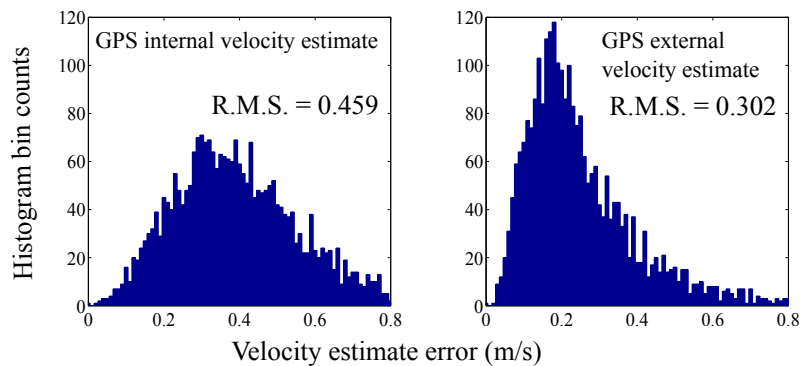


Figure 5.5: Distributions of GPS velocity-estimate error for data set with 28% phase-lock measurements




Figure 5.6: Distributions of GPS velocity-estimate error for data set with 99.5% phase-lock measurements

Wind-measurement accuracy of $O(0.01m/s)$ requires airframe velocity-estimation accuracy of $O(0.01m/s)$. However, when fusing information from the accelerometer and GPS some level of averaging may be applied to the GPS velocity estimate. The amount of averaging is dependent on the quality of the accelerometer - in particular the noise and bias drift characteristics. The 16-state (federated) filter simulation used in section 4.2.3 was used to investigate the necessary GPS velocity-estimate accuracy needed to produce a platform-velocity-estimate using contemporary accelerometers and gyroscopes with $O(0.01m/s)$ accuracy. The simulation suggests that GPS velocity-estimate accuracy of $0.1 m/s$ or better is adequate.

The error levels seen in Figures 5.5 and 5.6 show that GPS velocity-estimation accuracy of $O(0.1m/s)$ requires that a GPS receiver be able to maintain PLL phase lock in most satellite channels most of the time - 95% of measurements made with phase lock appears a conservative predictor. This is not, however, an unreasonable expectation. Figure 5.7 shows the number of satellite channels with PLL phase lock at each measurement epoch during a flight test described in Chapter 7 and shows that an adequate number of channels have phase lock during a high percentage of the flight time.

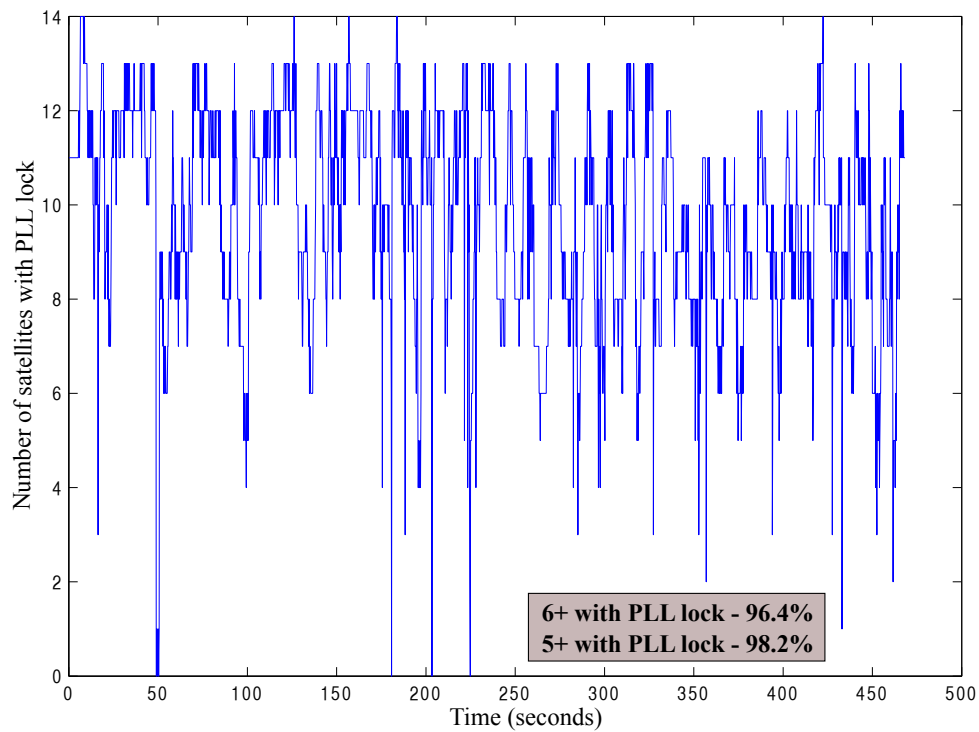


Figure 5.7: GPS channels with PLL phase lock during a sUAS flight - U-blox 6T GPS receiver with an un-amplified patch antenna.

5.3 Expected sUAS velocity estimation accuracy with contemporary sensors

The lack of visibility into the inner working of COTS GPS receivers makes broad statements about what results may be expected from sUAS velocity-estimation systems difficult, but two important results have been shown above. First, filtering applied by the GPS receiver may introduce significant error into the GPS velocity estimate. It may be that the filtering applied is suitable for sUAS velocity estimation, depending on how the sUAS velocity-estimation filter is constructed, but details of the filtering applied by the GPS receiver are generally unknown. There are two ways to work around this issue. Some GPS receivers may allow for all filtering to be turned off. Other receivers provide raw FLL and PLL measurement data and an external solution may be computed for the velocity measurement. The second result is that the GPS velocity-estimate accuracy is strongly dependent on having PLL phase lock in a sufficient number of satellite channels. When the estimate is formed using a weighted least squares solution then channels with FLL Doppler measurements do not significantly degrade the estimate so long as there are sufficient channels (5+) with PLL phase lock. A GPS receiver reporting which channels are using PLL versus FLL measurements allows for monitoring the quality of the velocity estimate and using it accordingly. For example, the weighting applied to the GPS velocity measurement in the velocity-estimate sensor fusion may be varied - E.G., the EKF R matrix may be computed based on a velocity estimate quality metric.

The previous section showed that at least one COTS GPS receiver, the U-blox 6T, is capable of the desired velocity estimation accuracy of $O(0.1m/s)$ while aboard a sUAS in flight. Due to the fast pace of change in the consumer electronics market that receiver is now at the end of its life cycle and will be replaced with the U-blox M8T, which will offer enhanced sensitivity and concurrent dual-constellation multi-GNSS raw measurements. It is reasonable to expect that other GPS receiver choices are available that offer similar or better performance to the U-blox 6T, and more will become available in the future. With the U-blox 6T the onboard filtering of the velocity

estimate limited the accuracy which can be achieved in a high-dynamics motion environment, but this issue may be circumvented by using raw measurement data supplied by the receiver to produce an external least-squares solution for the velocity estimate.

The absolute accuracy of a sUAS velocity-estimation system is dependent on the accuracy of the GPS velocity estimate. However, based on the attitude-estimate accuracy described in section 4.2.5, and contemporary MEMS accelerometer performance (e.g., the Analog Devices ADIS16488 noise density of $0.25 \text{ mg}/\sqrt{\text{Hz}}$), the accuracy of a sUAS velocity-estimation system may be nearly an order of magnitude better than the accuracy of the GPS velocity estimate. With the ability to produce velocity estimates with accuracy of $O(0.1\text{m/s})$ using sUAS-appropriate COTS GPS receivers, airframe velocity estimates with an accuracy of $O(0.01\text{m/s})$ may be produced, supporting high-accuracy sUAS-based wind measurement.

Chapter 6

Optical reference vector sensor

No existing sensor suitable for use in sUAS for observation of a reference vector is particularly attractive. Accelerometers require specific-force correction. Magnetometer measurements are corrupted by hard-iron and soft-iron effects. Thermopile and electrostatic-potential sensors are difficult to calibrate and subject to weather disturbances. GPS baseline systems are noisy when using sUAS scale baseline lengths. Optical systems generally have limited accuracy, may be reliant on particular lighting conditions and cannot operate in some weather conditions. This chapter describes a new sensor suitable for sUAS and capable of making a reference-vector-direction measurement.

An additional sensor providing a reference-vector measurement is useful for several purposes. First, the sensor may be used to validate sUAS attitude-estimation systems. Chapter 4 illustrated the fact that attitude-estimation-system accuracy is dependent on the motion environment. The ability to validate an sUAS attitude-estimation system while flying in a sUAS that is subject to wind-gust-induced motion is not possible with larger/heavier sensors. sUAS-based wind measurement of the accuracy considered here requires an attitude-estimation system accurate to the order of 0.1° . The usual approach of validating a sUAS attitude-estimation system against a second higher-quality attitude-estimation system may be impractical in a sUAS due to the size, weight and power requirements of this more accurate system. A higher-quality attitude-estimation system flown in a larger aircraft, with different levels and frequencies of platform motion than a sUAS, may produce

misleading results. The sensor described herein is small, light, and has low power requirements making it suitable for use in sUAS and therefore useful for validating sUAS attitude-estimation systems.

A second use for the sensor is to include its measurement in the sUAS attitude-estimation-system sensor fusion. This is advantageous when a high-accuracy attitude estimate is desired. Observation of the gravity vector with accelerometers requires the platform inertial acceleration be removed from the measurement, but this acceleration cannot be directly measured. It must be estimated based on other sensors introducing error. The magnetic field measured by a magnetometer is distorted by magnetic sources and ferrous metals in the platform and calibration to remove these effects is difficult, particularly in sUAS using electric propulsion systems. The optical sensor introduced here, on the other hand, is not influenced by accelerations, rotation rates, magnetic field effects, or any other effect associated with the platform so long as the aperture has a clear view of the ground-based beacon. Including this sensor in the attitude sensor fusion may significantly reduce the attitude-estimation error.

Finally, this sensor may be used for sUAS state estimation in special situations such as when operating in GPS-denied environments. This concept has not been developed, but is described further in section 6.3.

Use of optical beacons or beams for attitude measurement has been previously explored. Sepp [66] and Aruga and Igarashi [3] used a laser beacon from the ground to a satellite as a directional reference for stationkeeping. Use of optical position sensing detectors (PSDs) for attitude measurement has also been developed. Salcudean [63] used a PSD-based attitude measurement system for machine control. To the author's knowledge these concepts have not been combined in a system for use in aircraft. Alkeryd [2] suggested the use of optical PSDs for limited UAS position measurement, but did not extend to attitude estimation or produce a system suitable for use in UAS. The system described herein uses a monochromatic, time-modulated, ground-based optical beacon, single-element small-diameter refractive optics, a two-dimensional position-

sensing photodiode detector, signal conditioning and a microprocessor to make a vector-direction measurement of the vector between the beacon and sensor aperture. Measurement of the distance between the beacon and sensor may also be possible, based on signal strength or time of flight, but this capability has not been developed. The system requires line-of-sight between the beacon and sensor, but operates outside in full daylight conditions. The optical beacon consumes only a few watts of power, does not employ a laser (eliminating aviation regulatory concerns), and is eye-safe. The entire system may be constructed at low cost commensurate with use in sUAS.

6.1 Sensor description

The prototype sensor requirements relate to validating sUAS attitude-estimation systems for sUAS-based wind measurement. Many system parameters may be varied to tailor the performance for other uses. The primary requirements were set as:

- A sensor suitable for use in sUAS; the on-board portion of the system must be small ($< 70 \text{ cm}^3$), lightweight ($< 100 \text{ g}$), have low power consumption ($< 1 \text{ W}$), and be inexpensive ($< \$500$ in small quantities).
- The ground-based beacon must be eye-safe, have power consumption low enough to use batteries under 2 kg for a 6 hour run time, and not require an active pointing mechanism.
- Accuracy on the order of 0.1° must be possible over a range of 30° at ranges up to 100 meters.
- The system must be able to operate in full daylight conditions.
- Continual measurement of the reference vector is not required. Measurements may be made during occasional periods when the sUAS flies through the beacon beam and the reference vector is in the sensor's field of view.

6.1.1 Principle of operation

The sensor is fitted in a downward looking manner in the sUAS airframe as shown in Figure 6.1. The sensor uses simple optics to focus optical power from the beacon into a spot on a two-dimensional position-sensing photodiode detector. Background optical radiation will also reach the detector and will obscure the desired signal. To deal with this background radiation both optical and signal-frequency separation filtering are used. The beacon uses a narrowband emitter (LED with 20 nm maximum optical bandwidth) and an optical band-pass filter (with a 10 nm bandwidth) in the receiver optics that removes a large portion of the background optical radiation while allowing the beacon signal to pass. Also, the beacon output is modulated using a free-running 1 kHz square wave. The background optical radiation has a signal spectrum that is concentrated at low frequency compared to the 1 kHz beacon frequency. The background is not stationary due to a non-homogeneous scene in the sensor's field of view and changing platform location and attitude, but has little frequency content above ~ 10 Hz. Depending on intended operations and the link budget (discussed below) a wide, fixed (non-tracking) beacon beam may be used. Otherwise a narrower beam may be used and gimbaled to point towards the sUAS. In either case measurements can only be made when the airborne sensor is in the beacon's beam, and when the beacon is in the airborne sensor's angular field of view. The detector provides continuous analog outputs related to the position of the centroid of the light spot on the detector. Anode and cathode currents from the detector pass through analog amplification and filtering channels and are then digitized. The digital signals are processed in a microcontroller to produce the vector-direction measurement. Figure 6.2 shows a block diagram for the sensor system.

A commercial two-dimensional position-sensing photodiode detector (PSD) is utilized. PSDs are fabricated with either a duo-lateral structure or a tetra-lateral structure. Duo-lateral detectors have a pair of anodes along one pair of parallel edges of a photo-resistive structure and a pair of cathodes along the other edges. A bias is applied between anodes and cathodes and the amount of current flowing in/out of a particular cathode/anode is dependent on the amount and location




Figure 6.1: Optical reference vector sensor system concept of operation

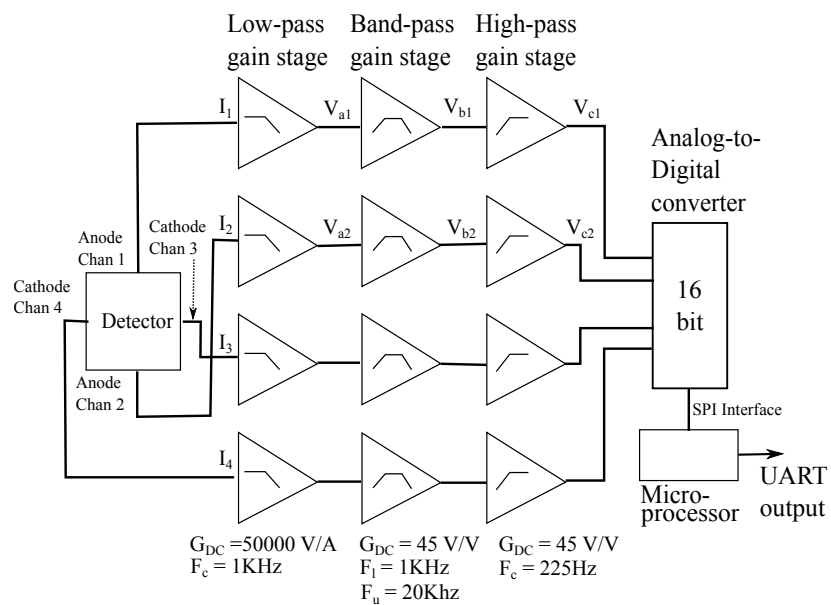


Figure 6.2: Optical reference vector sensor block diagram

of incident optical power on the detector. Tetra-lateral detectors have a single anode and four cathodes, and operate in a similar fashion with current split between the cathodes based on light intensity/position on the detector. Duo-lateral detectors are considered to have better linearity [26]. Various size detectors are available and larger detectors allow for a larger aperture size for a fixed field of view with simple optics.

Ignoring (the mild) non-linear effects the anode (or cathode) current pair produced by a light spot on the detector is given by [26]

$$I_1 = P_d R_d (1 - S/L) \quad (6.1)$$

and

$$I_2 = P_d R_d S/L, \quad (6.2)$$

where P_d is the incident spot power [W], R_d is the detector responsivity [A/W], L is the width of the detector [m] and S is the spot (centroid) distance from the anode/cathode 1 edge of the detector [m].

The position of the spot along one axis relative to the detector center may be determined as

$$X_m = \frac{I_1 - I_2}{I_1 + I_2} = 1 - 2(S/L) \quad (6.3)$$

where X_m is the derived measurement and ranges from 1 to -1 as S ranges from zero to L . Note that although the currents are linearly related to the spot position, and this derived measurement X_m is linearly related to the spot position, this measurement is a non-linear combination of the anode/cathode currents.

If the incident power on the detector is not a single spot, then the output currents may be approximated as those due to a superposition of spots. The earth-albedo reflection power is strong when operating in daylight conditions and the detector is subjected to background light power distributed across the detector. This unwanted light power (and the associated detector current) is often significantly larger than the desired signal even when using optical band-pass

filtering. It is not possible to directly recover the signal signature in this case due both to the measurement non-linearity with respect to currents (Equation (6.3)) and the unknown distribution of background light power on the detector. For this reason the beacon is time-modulated. The analog signal conditioning includes band-pass and high-pass stages which block the slowly varying background signal while allowing the higher-frequency modulated beacon signal to pass. After analog-to-digital conversion the anode/cathode signals are rectified (as discussed in detail below) to determine the signal amplitude for each channel for use in calculating the focused spot position with Equation (6.3). Using single lens optics the angle of the beacon beam in each axis may be calculated as the inverse tangent of the ratio of the spot position in that axis and the lens' focal length.

6.1.2 Link budget

Design of the sensor system involves trade-offs between detector size, aperture size, field of view, angle measurement accuracy, beacon power, beacon beamwidth, amplifier noise-figure and maximum range. The parameters chosen for the prototype sensor system, and the resultant performance described in sections 6.2 and 7.4, show this class of sensor has significant utility for use in sUAS. The prototype sensor system uses a 1 cm-square detector and optics providing a $\pm 15^\circ$ field of view. Using a 1 cm² detector and simple single-lens optics the maximum size of a lens with a focal length yielding a $\pm 15^\circ$ field of view limits the aperture size to 10 cm² (3.6cm diameter). Since image sharpness is not an issue here (the detector currents relate to the light spot centroid only), a lightweight Fresnel lens is used.

To determine the required signal-to-noise ratio in the analog channels the relation between measurement noise and channel noise is analyzed. If the channel 1 and 2 current signals are corrupted by uncorrelated Gaussian current noises η_1 and η_2 with variances $\sigma_{\eta_1}^2 = \sigma_{\eta_2}^2 = \sigma_\eta^2$ then

the measurement is

$$\hat{X}_m = \frac{I_1 + \eta_1 - I_2 - \eta_2}{I_1 + \eta_1 + I_2 + \eta_2} = \frac{P_d R_d (1 - 2S/L) + \eta_1 - \eta_2}{P_d R_d + \eta_1 + \eta_2}. \quad (6.4)$$

The measurement is unbiased as the expectation is $1 - 2S/L$. The variance of the measurement is

$$\sigma_{\hat{X}_m}^2 = E[(\hat{X}_m - \bar{\hat{X}}_m)^2] = E \left[\left(\frac{P_d R_d (1 - 2S/L) + \eta_1 - \eta_2}{P_d R_d + \eta_1 + \eta_2} - (1 - 2S/L) \right)^2 \right]. \quad (6.5)$$

Simplifying based on the assumption that the required signal power will be significantly larger than the noise, $1 (P_d R_d) / \sigma_{\eta_1} \gg 1$,

$$\sigma_{\hat{X}_m}^2 \approx E \left[\left((1 - 2S/L) + \frac{\eta_1 - \eta_2}{P_d R_d} - (1 - 2S/L) \right)^2 \right] = E \left[\left(\frac{\eta_1 - \eta_2}{P_d R_d} \right)^2 \right] = \frac{2\sigma_{\eta}^2}{(P_d R_d)^2}. \quad (6.6)$$

The standard deviation in the measurement is

$$\sigma_{\hat{X}_m} \approx \sqrt{2} \frac{\sigma_{\eta}}{P_d R_d}. \quad (6.7)$$

Noting that \hat{X}_m has a range of ± 1 , and $\frac{P_d R_d}{\sigma_{\eta}}$ is the signal-to-noise ratio (SNR) in the current signal channels, this equation directly relates the noise-induced measurement error to channel SNR. The prototype sensor system measurement error of less than 0.1° over a range of $\pm 15^\circ$ requires a signal-to-noise ratio greater than 49.5 dB ($(P_d R_d) / \sigma_{\eta_1} > 10^{4.95}$), which conforms to the assumption preceding Equation (6.6).

The expected signal-channel noise-power is determined using the model in Figure 6.3 that includes the detector and the first two gain stages; noise contributions from the third gain stage are negligible due to the large gain in the preceding stages. Using precision JFET op-amps (Analog Devices AD8625ARZ) for the gain stages and ensuring clean bias voltages the op-amp noise currents, I_{n2} and I_{n4} , and bias-voltage noise sources, V_{n1} and V_{n3} , are negligible over the bandwidths involved. For the remaining noise sources the RMS voltage noise at the output of the second stage is

$$V_{b.noise} = \left[\left(\frac{R_2}{R_{hp}} \right)^2 V_{n2}^2 B_2 + \left(\frac{R_f R_2}{R_{hp}} \right)^2 (I_{n1}^2 + I_{n3}^2) B_1 + \left(\frac{R_2}{R_2 + R_{hp}} \right)^2 V_{n4}^2 B_2 + R_{hp}^2 I_{n5}^2 B_2 \right]^{1/2}, \quad (6.8)$$

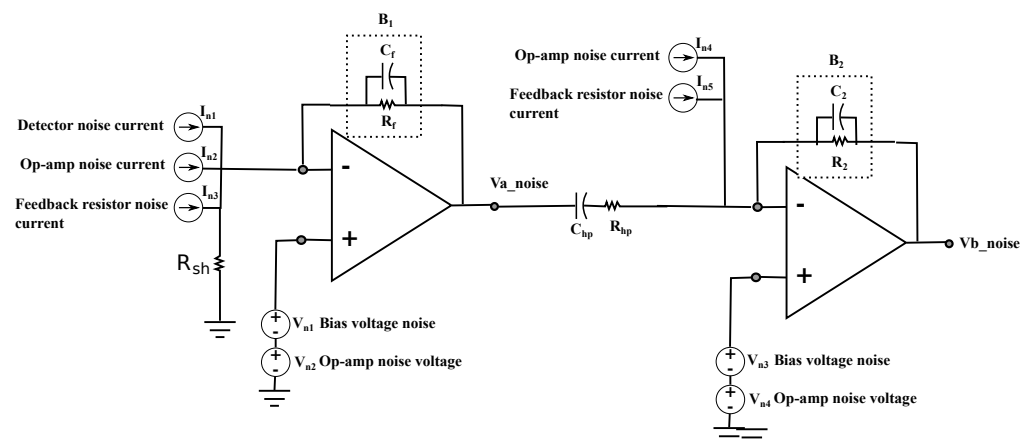


Figure 6.3: Front-end noise model

where V_{n2} and V_{n4} are the op-amp noise-voltage spectral densities, I_{n1} is the detector current-noise spectral density, I_{n3} and I_{n5} are the Johnson noise in the feedback and filter resistors expressed as current-noise densities and B_1 and B_2 are the bandwidths of the first and second stage feedback networks. For the components chosen for the prototype system the main noise contributors are the detector noise-current, the op-amp noise-voltages, and the first stage feedback-resistor noise-current, and Equation (6.8) evaluates to $300\mu V$ RMS at the output of the second stage. The signal-to-noise ratio in a second-stage output channel voltage may be calculated by relating the second-stage-output signal voltage to the incident signal power multiplied by the transimpedance gain G_{1-f_m} and the voltage gain G_{2-f_m} (at the modulation frequency), yielding

$$SNR = \frac{\phi_d A R_d G_{1-f_m} G_{2-f_m}}{V_{b_noise}} \quad (6.9)$$

where ϕ_d is the incident power density ($P_d = \phi_d A$) and A is the aperture area. Maintaining a 49.5 dB SNR, and considering the first (36000 V/A) and second (32 V/V) stage net gain at the modulation frequency, the detector responsivity of 0.63 A/W and the aperture diameter of 0.01 m, the minimum optical power density ϕ_d at the sensor aperture is $124 \mu W/m^2$.

Additional optical signal power may be required due to loss in the optical band-pass filter, as discussed below, but this calculated minimum optical power density shows the utility of this sensor system. If the beacon is lensed to produce a conical beam of (full width) θ_b [rad], the power density at the sensor as a function of range and total radiated power P_{rad} is

$$\phi_d = \frac{P_{rad} \alpha_L}{\theta_b^2 R^2}, \quad (6.10)$$

where α_L is a loss factor to account for optical band-pass filter loss, radiated power not included in the beam angle and for power density not being uniform across the beam angle. This equation yields the maximum range

$$R_{max} = \sqrt{\frac{P_{rad} \alpha_L}{\theta_b^2 \phi_{d,min}}}. \quad (6.11)$$

For a beacon with 1W radiated power, no optical losses ($\alpha_L = 1.0$), a beam width of 10° and the minimum optical power density derived above the maximum range for the system is 513 meters.

Using the prototype system as an example where $\alpha_L = 0.1$, the same radiated power, beam width and minimum optical power density result in a range of 162 meters.

6.1.3 Sensor module design




Figure 6.4: Sensor module front and back

The sensor module consists of a small (5 cm x 7.5 cm) printed-circuit-board with a 3D-printed assembly affixing the aperture lens and optical band-pass filter over the detector as shown in Figure 6.4. The 3D-printed assembly is somewhat translucent so it is covered with a foil-tape shield. The prototype sensor uses an external ARM-Cortex-M4 development board to provide the microprocessor functionality, but the sensor-module PCB has adequate area to incorporate a microprocessor in the future for a more integrated solution. Signal current levels from the detector are very small—on the order of nano-amperes. The first amplification stage, which operates as a transimpedance amplifier, uses a 20 volt supply. This voltage level is necessary to provide 4 volt and 16 volt anode/cathode bias levels while allowing amplification of currents due to background light without saturating the stage. The remainder of the board uses a 5 volt supply. 5V power is supplied by a two-cell lithium-polymer battery and linear regulator. A separate four-cell lithium-polymer boost battery and a 20V linear regulator provides the 20 volt bias and first stage supply. These supplies must be tightly regulated and filtered to avoid introducing noise into the circuit front end.

The first amplification stage employs a low-pass filter to limit noise contributions from the detector, op-amp voltage-noise, and transimpedance-gain feedback-resistor-noise, all of which are significant contributors to the front-end noise as discussed above. A band-pass second-stage amplifier reduces extraneous signal due to low-frequency background light, while further band-limiting the noise introduced by the other sources in the front-end. The beacon signal is square-wave modulated at 1 kHz. The 1 kHz signal allows for a 1 kHz lower-corner frequency of the second stage band-pass filter. This lower corner is high enough to block variation in the background signal produced by the combination of a heterogeneous background scene and sUAS platform motion. 1 kHz signal modulation is also low enough for oversampling the multi-channel signals with a low-power microprocessor. The analog filtering significantly alters the signal waveform as shown in Figure 6.5. Samples of the signal levels for an anode-channel pair and a cathode-channel pair are shown for three different reference-vector directions. The maximum practical sampling rate (limited by the SPI bus throughput) yields approximately 47X oversampling. This oversampling rate is sufficient to remove sensitivity of the amplitude determination to the unknown phase of the beacon signal and variation between the beacon and sensor clocks. Samples for a two-modulation-cycle period are collected and the channel amplitudes are calculated by determining the mean level of each channel and then summing the absolute values of the difference between the samples and the mean for each channel. Since the signals do not vary in shape, only amplitude, this effective calculation of the area under the rectified signals is proportional to their amplitude.

To avoid the need for automatic-gain-control a 16-bit ADC is used to provide a high dynamic range and the output of the second stage is also sampled when the third stage is saturated. The second stage sampling is not shown in the block diagram (Figure 6.2). Sampling two-modulation-cycle period segments for each of two channels per axis and for two axes, along with processing time in the microcontroller, still allows for a measurement rate of up to 200Hz, or 100Hz when sampling 8 channels to allow for potential third stage saturation.

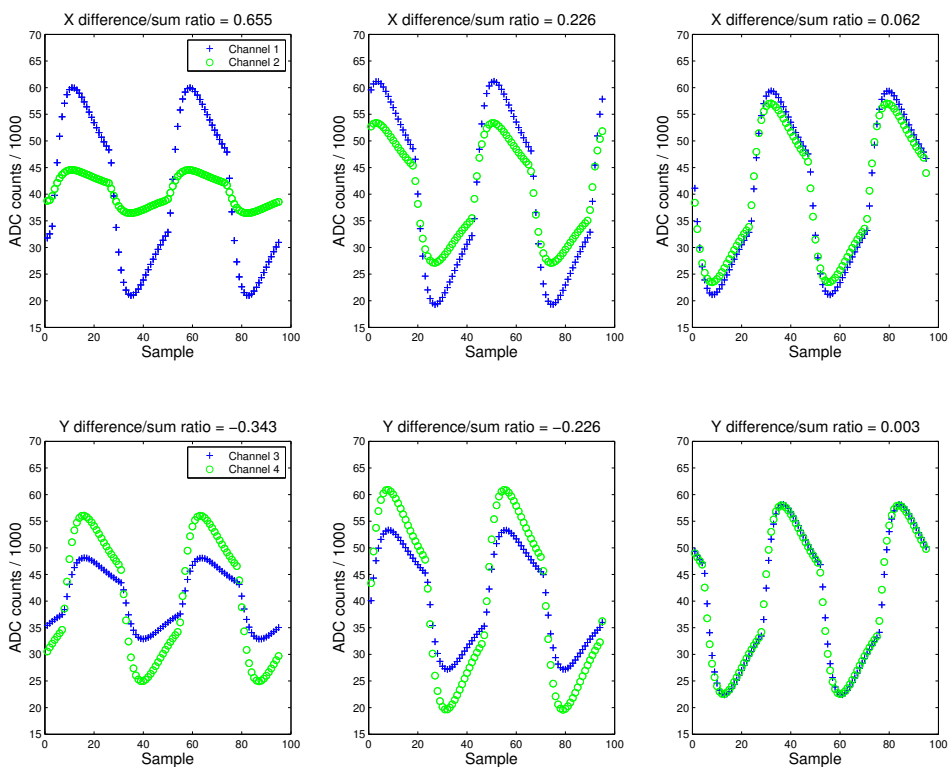


Figure 6.5: Signals after analog to digital conversion

6.2 Sensor characterization

6.2.1 Prototype system parameters

Parameteric values used in the prototype sensor module are summarized here:

- A 1 cm² position-sensing-photodiode detector with a responsivity of 0.63 A/W.
- Optics consisting of a 0.001 m² Fresnel aperture lens, chosen for its light weight, and a 656nm-center-wavelength hard-coated optical-density-OD4 10nm band-pass filter. While the optical focal length was chosen for a $\pm 15^\circ$ field of view, the rapid prototyped structure supporting the lens and optical band-pass filter left excessive space between the filter and detector resulting in the filter housing cutting the field of view to $\pm 12^\circ$. This problem can be corrected through a redesign of the optics support structure.
- The first amplification stage provides a DC transimpedance gain of 50,000 V/A with a 1 kHz-corner-frequency single-pole low-pass filter.
- The second and third amplification stages provide DC gain of 45 V/V each. The second stage includes a band-pass filter with corner frequencies of 1 kHz and 20Khz. The third stage includes a 225Hz-corner-frequency single-pole high-pass filter.
- Two prototype beacons were produced. The first uses seven deep-red (655 nanometer) LEDs with total average output power of 5 watts lensed to provide a 18 degree (HPFW) beam. The second uses four deep-red (655 nanometer) LEDs with combined average output power of 2 watts lensed to provide a 95 degree (HPFW) beam.
- Digital processing is performed in an ARM Cortex-M4 microcontroller with a 168 MHz clock frequency.

While the chosen optical band-pass filter and LED appeared well matched based on the LED manufacturer's stated typical peak wavelength and spectral width, they were not, in fact, well

matched. The LED specifications include both typical and minimum/maximum values. The typical values would provide a good match. However, the actual characteristics of the LED, while within the min/max specifications were not well represented by the typical values and not well matched to the optical band-pass filter. A significant power loss (10 dB) occurred due to this mismatch. The range of the prototype system is reduced by a factor of $\sqrt{10}$ relative to what would be achieved with a perfect match. Despite this unexpected loss a range of more than 100 meters is possible using the 18 degree beamwidth beacon without sacrificing measurement accuracy. Improved matching of the LED and filter may be achieved through sourcing LEDs with tighter control of the peak frequency and spectral width, using a laser for the beacon emitter, or by using a wider optical band-pass filter. The third option requires careful design of the first amplification stage to avoid saturation as more background light is allowed to reach the detector.

6.2.2 Calibration

While a duo-lateral PSD detector was chosen for good linearity, testing of the sensor module indicates that, although the linearity was good for spots in the central area of the detector, linearity was not sufficient to use a simple three-point linear calibration and meet measurement accuracy goals for beam angles in the outer portions of the field of view. A two-axis CNC positioning system was used to move a beacon relative to the sensor module to provide data for a more comprehensive calibration as shown in Figure 6.6. The beacon was fitted with a low-transmissivity neutral density filter, and operated at reduced power, to avoid saturating the sensor with only 3 meters range during the calibration. A large number of measurements were gathered and averaged at each of 625 grid points on a 25-by-25 point grid of equally spaced angles across the sensor's field of view. Using this data a calibrated measurement may be computed by determining which grid cell a particular measurement falls within and performing a two-dimensional linear interpolation using the 4 bounding grid points from the 625 point grid. This removes effectively all of the deterministic error due to detector non-linearity. The sensor accuracy for an incident power density level of 0.8

mW/m^2 is shown in Figure 6.7 prior to and after applying the calibration interpolation.

6.2.3 Noise performance

The range-versus-noise performance of the prototype system is measured in the lab using a neutral-density optical filter to reduce the beacon-signal power level. Interestingly, the lab fluorescent lighting contained noise within the pass-band of the sensor's front-end and this prevented the predicted performance from being achieved unless the indoor lights were turned off. The noise floor with the lab lighting turned off and the noise floor when operating outside were equal and allowed expected performance to be achieved.

Figure 6.8 shows measured noise performance as a function of beacon average power density at the sensor. Each point is the noise measured over a ≈ 400 point sample set. The plot shows that measurement noise below 0.1° RMS requires a power density of greater than approximately $0.7 \text{ mW}/\text{m}^2$. This result agrees well with the expected noise performance (plotted as 'Predicted performance') based on the link budget developed in section 6.1.2 and accounting for the 10dB loss in the optical band-pass filter discussed in section 6.2.1. At higher power levels the measurement noise is trending higher than predicted. This may be due to quantization noise and the effect of unknown signal phase on the amplitude determination in the digital processing. The implications are that the noise model developed is useful for predicting performance at the desired accuracy levels but as desired accuracy is increased unmodelled effects are becoming significant and the required power density at the instrument will be somewhat higher than calculated based on the noise model.




Figure 6.6: Sensor calibration setup using a 2-axis CNC positioning system to move beacon relative to the sensor

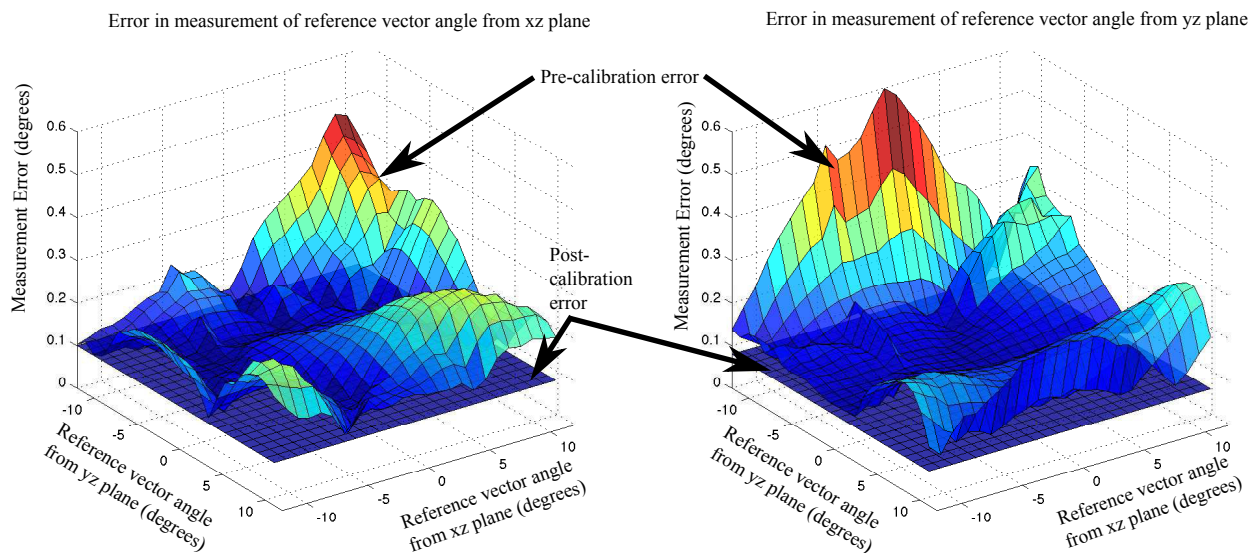


Figure 6.7: Pre-calibration error due to non-linearity of the sensor detector

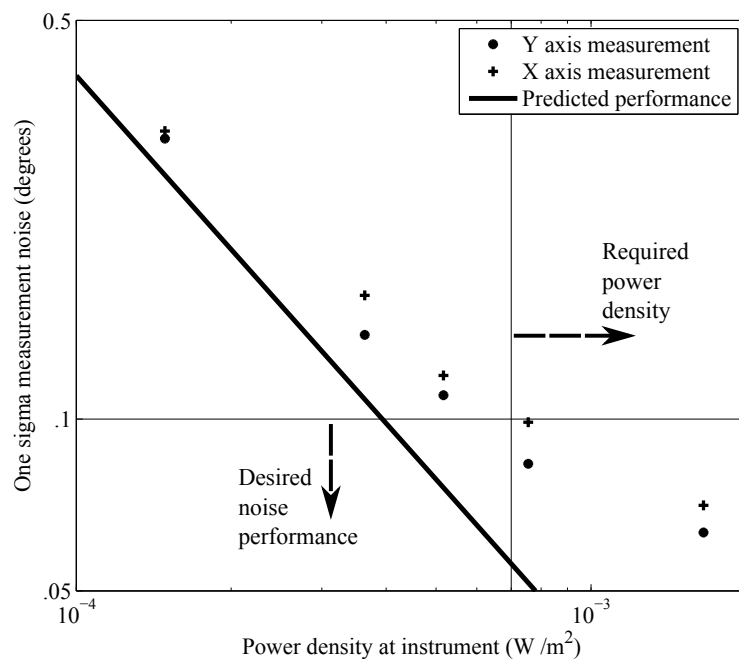


Figure 6.8: One sigma measurement noise versus power density at sensor

6.3 Potential adaptations

The prototype sensor system proved adequate for the intended use — providing an independent reference-vector measurement with accuracy on the order of 0.1 degrees for a sUAS in flight. Trade-offs between detector size, aperture size, field of view, angle measurement accuracy, beacon power, beacon beam angle, amplifier noise figure and maximum range allow significant flexibility in adopting the sensor for a variety of other uses. This system may be adapted to numerous purposes by making appropriate trade-offs in the design space, and may find applications in a variety of fields outside sUAS state estimation.

Since measurements may only be made while the sensor is in the beacon beam and the beacon is in the sensor's field of view, a large beacon beam width and sensor field of view are desirable to make measurements over a large geometric volume. But a large beacon beam width results in more spreading loss of the beacon power and a larger sensor field of view results in a smaller aperture due to short lens' focal lengths limiting available lens diameter. These effects may be mitigated with a stronger beacon, larger detector (allowing a longer focal length and therefore a larger aperture), or shorter range. Alternately the beacon may incorporate a pointing system and use a narrow beam width. This allows for a larger operational range at the expense of complexity. Range may be increased by relaxing the angle measurement accuracy requirement when the application allows. Range may also be increased by improving the receiver's noise figure. Switching to an infrared emitter in the beacon may facilitate improving the receiver's noise figure as the reduced background optical power at IR will allow a higher gain in the first stage, lessening the effect of second stage noise sources. However, first stage noise sources are most significant so this option can only provide an incremental improvement. Finally, reducing optical losses, such as the mismatch between the beacon emitter and the sensor's optical band-pass filter, also provide an opportunity for increasing range. This is particularly true for the prototype system which suffered from a large optical mismatch loss.

Although the sensor system is being used to validate an IMU-based attitude-estimation system, it may also be used as an additional sensor to augment that system. The optical sensor measurement may be treated as an additional observation in an extended Kalman filter in much the same manner as an accelerometer or magnetometer measurement. This is beneficial as this sensor provides a reference-vector direction measurement with significantly better accuracy than usually achieved with either an accelerometer or magnetometer sensor typically used in sUAS in an airframe in flight. Improvement in the IMU-based attitude-estimation system using this approach will be highly dependent on the operational scenario and the percentage of time that measurements from the optical sensor are available.

While a single beacon was employed here, potential exists to use multiple beacons and make multiple, simultaneous reference-vector direction measurements. Multiple beacons may employ different modulation so that their signals may be separated in the detector electronics. For example, code division multiple access (CDMA) may be employed by further modulating the beacons' square wave modulation, e.g. with a Gold code [24]. This allows multiple beacons to be individually recognized and processed in software, and would be particularly useful in special situations such as GPS denied environments. With a sufficient number of beacons, and resulting angular measurements, this system could provide a position estimate supplanting GPS for position measurement/estimation.

Chapter 7

Flight test analysis

Invariably, producing and testing a physical system expands on the knowledge gained through analysis of a problem. Certain concerns which are trivial from an analytic perspective may pose significant limitations in practice. This is certainly the case in sUAS state estimation. To better understand practical problems presented by the sort of sUAS state estimation described in Chapters 4 and 5 a sUAS state estimation system was constructed and flown. The main system components included:

- A Skywalker 1680 airframe with electric propulsion and radio control shown in Figure 7.1. This is an off-the-shelf airframe intended for radio control hobbyists. It's large payload capacity makes it a suitable choice for use as an sUAS. The airframe is expanded polyolefin foam with carbon fiber reinforcements, has a wingspan of 1.68 meters, and a flying weight of approximately four pounds. The power system uses a brushless electric motor, controlled with an electronic speed controller, turning a 9-inch diameter, 6-inch pitch propeller, and a 3 cell, 4400 milliamp-hour, lithium polymer battery. A standard 2.4 GHz hobby-grade model airplane transmitter, receiver, and servos were used for control.
- A sensor data collection system based on an autopilot under development at the University of Colorado at Boulder. For purposes of these flight experiments this autopilot provided a platform for a STM32F427 microcontroller and a micro-SD storage card. The microcontroller was programmed with firmware to collect data from sensors using standard UART

and I2C protocols, and to store the data onto the micro-SD card for post-processing.

- A STEVAL_MKI108V2 9-axis IMU module including L3GD20 and LSM303DLHC MEMS devices. The L3GD20 contains a 3-axis MEMS gyroscope. The LSM303DLHC contains a 3-axis MEMS accelerometer and a 3-axis magnetoresistive magnetometer. Both devices use I2C digital interfaces. While higher performance MEMS sensors are available these sensors were chosen as they are used in the development autopilot noted above.
- A U-blox LEA-6H GPS receiver module with a passive patch-antenna and logic-level UART digital interface.
- The prototype optical reference vector sensor described in chapter 6. Two different ground-based beacons were used; one with a average optical output power of 2W and a 95° half-power-full-width beam angle and a second with a average optical output power of 5W and a 18° half-power-full-width beam angle.

Sensor data is collected and stored on a micro-SD storage card and attitude estimation and attitude-estimate evaluation is performed with post-processing. This allows multiple filters and/or sets of filter parameters to be evaluated using the same flight data. Sensor data from the gyroscope, accelerometer and magnetometer is stored with a measurement frequency of 200 Hz. Sensor data from the GPS receiver is stored at 5 Hz. Sensor data from the optical reference vector sensor is stored at 50 Hz. Presented flight test results focus on data collected during a flight on December 5, 2014, at the U.S. Department of Commerce Table Mountain Field Site, a flight conducted March 12, 2105 at the CU Boulder South open space and a flight conducted on March 14, 2014 at the Arvada Associated Modelers Flying Field.



Figure 7.1: Test flight aircraft

7.1 Sensor calibration

Equations (4.1), (4.2) and (4.3) contain a variety of terms which introduce error into the gyroscope, accelerometer and magnetometer sensor measurements. Many of these terms are constant, or near constant, and by estimating these terms a calibrated sensor measurement can be calculated with much of the error removed. Depending on the quality of the sensors and the quality of the desired state estimate a variety of steps may be taken to determine these constants. Some of these terms may be treated not as constants but as values which vary with temperature or other parameters. This section describes both the steps taken to determine these values for analysis of test flight data and also steps which may be taken to improve the calibration.

One aspect of calibration not considered in this section is the determination of the relative attitude of sensors. For the flight-test case the gyroscope, accelerometer and magnetometer were rigidly mounted to a small PCB. By using the PCB edge and face to define the coordinate direction for an IMU frame any misalignment of the sensors relative to the PCB is included in the scaling and cross-coupling calibration terms. For other, similar systems the sensors may not be mounted in such a convenient way. They might even be mounted at different locations in an airframe and have misalignment relative to each other.

7.1.1 Gyroscope calibration

The single-axis gyroscope measurement model in Equation (4.1) includes a scaling term, S_x , two cross-coupling terms, M_{xy} and M_{xz} , and a zero-point bias term, B_{fx} . The zero-point bias term, as described in section 4.1.1, is generally time varying and is estimated in the sensor fusion algorithm. However, it is common practice to determine an initial value and use this initial value to initialize the filter algorithm rather than starting the filter with some value that differs significantly. If the latter approach is taken the filter will normally converge to a good estimate of the zero-point bias over time, but until that occurs the attitude estimate will have unnecessary error. Determining

the initial zero-point bias is a simple matter of collecting data from the gyroscope while it is at rest. An airframe at rest with respect to the earth is still rotating in an inertial frame due to the rotation of the Earth. However for contemporary MEMS gyroscopes the Earth's rotation rate is typically near or below the noise floor of the gyroscope and determining the initial zero-point to an accuracy comparable to the Earth's rotation rate is unnecessary. A number of samples may be collected and averaged with the result used as the initial zero-point bias.

For many MEMS gyroscopes and applications the built-in factory calibration of gyroscope scaling and cross-coupling is satisfactory and S_x , M_{xy} and M_{xz} may be taken as 1, 0 and 0, respectively. This is the approach taken for analysis of the test-flight data. If a more accurate calibration is desired one of two approaches may be taken. First, calibration values may be produced through bench-testing. The procedure is to rotate the gyroscope by a known amount around a known axis while collecting the gyroscope output. The gyroscope output is integrated to produce an amount of rotation. With data collected for all three axes and for more than three rotations (with non-colinear axes) an overdetermined set of data is available and the calibration constants may be determined by a least-squares optimization. The second approach is to collect data during flight and perform an optimization with post-processing. In this case known rotations are not available so the optimization must use data from other sensors. For example as the sUAS platform undergoes a rotation in flight the rotation rate may be observed both with the gyroscope and through the time derivatives of the accelerometer and magnetometer measurements. Although this approach uses data from other sensors which may be noisy a large amount of data may be collected over time and used in the optimization.

7.1.2 Magnetometer calibration

For convenience, the simple form of Equation (4.3) is restated:

$$\hat{h} = \mathbf{A}_e[\vec{h}_{earth}]_B + b_e + \epsilon. \quad (7.1)$$

A calibrated magnetometer measurement with the effects of \mathbf{A}_e and b_e removed is calculated using the calibration matrices \mathbf{A}_{cal} and b_{cal} as

$$\hat{h}_{cal} = \mathbf{A}_{cal}(\hat{h} + b_{cal}). \quad (7.2)$$

Note that \mathbf{A}_{cal} is a best estimate of \mathbf{A}_e^{-1} and b_{cal} is a best estimate of $-b_e$.

If a 3 coefficient (only) calibration is desired to correct for zero-point bias and hard-iron effects then a simple calibration technique to determine the calibration column vector, b_{cal} , may be constructed by assuming that the distortion matrix, \mathbf{A}_e , is near unit norm. The author is unaware of others using this approach, though it may not be novel. The Earth's magnetic field vector is essentially constant over the time required for the calibration procedure. If \mathbf{A}_e , has near unit norm then $|\mathbf{A}_e[\hat{h}_{earth}]_B|$ is nearly constant regardless of the magnetometer orientation. Now consider two magnetometer measurements, \hat{h}_1 and \hat{h}_2 made in two different orientations. Ignoring the noise term, $\vec{\epsilon}$, and taking $|\mathbf{A}_e[\hat{h}_{earth}]_B|$ as constant leads to

$$(\hat{h}_1 - b_e) \cdot (\hat{h}_1 - b_e) = (\hat{h}_2 - b_e) \cdot (\hat{h}_2 - b_e) \quad (7.3)$$

which yields

$$|\hat{h}_1|^2 - |\hat{h}_2|^2 = 2(\hat{h}_1 - \hat{h}_2) \cdot b_e \quad (7.4)$$

Now if many pairs of measurements are used, and the $|\hat{h}_1|^2 - |\hat{h}_2|^2$ and $2(\hat{h}_1 - \hat{h}_2)$ terms for the pairs of measurements are collected in a column vector and a matrix, respectively, then determination of the calibration vector, b_e , becomes a classic least-squares optimization for an overdetermined set of linear equations.

Prior to flight the flight-test aircraft was rotated through 360° of azimuth at each of five pitch angles providing a diverse set of attitudes. Magnetometer measurements were recorded every 1/2 second during this period providing a set of data which was arranged into pairs and b_e was determined as described above. Figure 7.2 shows the result of this determination of the b_{cal} calibration vector. The curves represent the magnitude of magnetometer measurements with the

red curve being the uncorrected measurement and the blue curve being the measurement corrected with the determined value of b_{cal} (and $\mathbf{A}_{cal} = \mathbf{I}$). During the time period depicted in the figure the airframe is rotated through 360° of azimuth at each of five pitch angles providing a diverse set of attitudes. As can be seen application of the b_{cal} calibration vector determined with this technique greatly reduces the variation in the the measured magnetic field magnitude. Remaining variation is due to \mathbf{A}_e differing from \mathbf{I} .

Although the 3-coefficient calibration significantly improves the magnetometer measurement, a calibration including correction terms for scaling factor error, axial misalignment and cross-coupling, and soft-iron effects will further improve the magnetometer measurement accuracy and should be used when a high-accuracy attitude solution is required. Gegre-Egziabher et al [23] describe a calibration method which will correct for scaling factor error and some soft-iron effects; their calibration includes six terms that include b_{cal} and the diagonal terms of \mathbf{A}_{cal} . They describe this method as suitable for UAS because "aircraft structures are normally made from aluminum or composite materials", with the implication that soft-iron effects should not be present and the off-diagonal terms of \mathbf{A}_{cal} should be near zero. Analysis of magnetometer measurements collected during the sUAS flight test does not support this assertion. Use of Gegre-Egziabher et al's method failed to significantly improve the magnetometer measurement accuracy beyond the improvement provided by the simple 3-coefficient calibration described above. However, a full 12-coefficient calibration method described below was effective.

The method employed by Gegre-Egziabher et al may be expanded to estimate the full b_{cal} and \mathbf{A}_{cal} matrices. In the 3-coefficient calibration algorithm above the local Earth magnetic field strength was not needed as differencing measurements removed that term. Gegre-Egziabher's method, and the method described here use the local Earth magnetic field strength but if the value is incorrect the effect is only to scale the \mathbf{A}_{cal} matrix. Often the magnetometer measurement is normalized before use so this is not problematic. Since the Earth's magnetic field strength is essentially constant the magnitude of the corrected magnetometer measurement should also be




Figure 7.2: Magnetometer offset calibration results

constant regardless of orientation. Proper calibration constants may be determined by finding the constants which minimize the variation in the measured field strength across a set of measurements in diverse orientations. The squared magnitude of the corrected magnetometer measurement is

$$|\hat{h}_{cal}|^2 = \sum_{i=1}^3 \sum_{j=1}^3 \left(\mathbf{A}_{cal,i,j} (\hat{h}_j + b_{cal,j}) \right)^2. \quad (7.5)$$

Now consider a perturbation of this squared magnitude based on a perturbation of the calibration constants. Arranging perturbations of the elements of the calibration matrices in a column vector as

$$\delta x = [\delta \mathbf{A}_{cal,1,1} \ \delta \mathbf{A}_{cal,1,2} \ \delta \mathbf{A}_{cal,1,3} \ \delta \mathbf{A}_{cal,2,1} \ \dots \ \delta b_{cal,1} \ \delta b_{cal,2} \ \delta b_{cal,3}]^T \quad (7.6)$$

the perturbation of the squared magnitude may be written as a linear, first-order approximation as

$$\delta |\hat{h}_{cal}|^2 = \mathbf{J}(|\hat{h}_{cal}|^2) \delta x, \quad (7.7)$$

where \mathbf{J} is the Jacobian operator. Using a set of k magnetometer measurements, taken in diverse attitudes, Equation (7.7) is used to produce a set of linear equations as

$$\begin{bmatrix} |\vec{h}_{earth}|^2 - |(\hat{h}_{cal})_1|^2 \\ |\vec{h}_{earth}|^2 - |(\hat{h}_{cal})_2|^2 \\ \vdots \\ |\vec{h}_{earth}|^2 - |(\hat{h}_{cal})_k|^2 \end{bmatrix} = \begin{bmatrix} \mathbf{J}(|(\hat{h}_{cal})_1|^2) \\ \mathbf{J}(|(\hat{h}_{cal})_2|^2) \\ \vdots \\ \mathbf{J}(|(\hat{h}_{cal})_k|^2) \end{bmatrix} \begin{bmatrix} \delta \hat{x} \end{bmatrix}. \quad (7.8)$$

Equation (7.8) is in the form $y = Ax$ and may be solved using standard linear algebra techniques.

Now the set of calibration constants may be determined iteratively following:

- (1) Chose an initial guess for the calibration constants. This may simply be $\mathbf{A}_{cal} = \mathbf{I}$ and $b_{cal} = [0 \ 0 \ 0]^T$.
- (2) Using the set of k magnetometer measurements solve Equation (7.8) to determine $\delta \hat{x}$.
- (3) Update the estimated calibration constants using $\delta \hat{x}$.
- (4) Repeat steps 1-3 and monitor for convergence.




Figure 7.3: 12-coefficient Magnetometer calibration results

Using the same set of points as used for the 3-coefficient calibration this algorithm converges to double-floating-point numeric precision in less than ten iterations. Figure 7.3 show the results of this calibration in the same format and for the same time period as Figure 7.2. Note that only a small subset of the measurements shown in this figure is used to perform the calibration but the calibration results are quite good over the larger measurement set. The corrected measurements can now be seen to have a constant magnitude, excepting measurement noise. It is interesting to note that the values of b_{cal} determined with this 12-coefficient calibration algorithm differ significantly from those determined with the 3-coefficient calibration algorithm above, and that the values of b_{cal} determined with this 12-coefficient calibration algorithm do not work well on their own (e.g., when using $\mathbf{A}_{cal} = \mathbf{I}$). This indicates that a simple 3-coefficient calibration fails to include a significant amount of information that is required to minimize residual error. The more accurate magnetometer measurements provided by the proposed 12-coefficient calibration algorithm allow for a significant improvement in sUAS attitude-estimation accuracy.

7.1.3 Accelerometer calibration

Accelerometer calibration bears close similarity to magnetometer calibration. For many cases a simple 3-coefficient calibration accounting for zero-point bias may be sufficient. For highest fidelity a 12-coefficient calibration may be performed. However, since there is no error analogous to soft-iron errors the off-diagonal terms of the \mathbf{A}_{cal} matrix may be negligible for some devices and a 6-coefficient calibration may suffice. The advantage of a 6-coefficient calibration is that fewer measurements are required to produce a good calibration. This is significant because, unlike the magnetometer, taking calibration measurements while in motion can be problematic for the accelerometer. Motion that produces an inertial acceleration at the sensor will influence the measurement and corrupt the calibration. Potentially, calibration measurements could be collected during a period where the platform is rotated through various orientations with the center of rotation remaining near the accelerometer so that any linear acceleration experienced by the sensor is small. However, this may be operationally difficult. A better approach is to collect calibration measurements while the sUAS is held still in a variety of poses.

Working with the calibration algorithms showed an important point about calibration measurements. While a measurement set with diversity in pitch and yaw produced an acceptable set for magnetometer calibration the same set was unsuitable for accelerometer calibration. Without roll diversity the measurement set provided adequate observability of the magnetic field vector due to the local magnetic inclination (66°), but lacked full observability of the gravity vector. A set of calibration measurements collected while the airframe is held motionless in six attitudes (level, upside-down, right wing down, right wing up, nose down, nose up) provides adequate observability for both magnetometer and accelerometer calibration.

7.1.4 GPS latency calibration

While a calibration to reduce errors in a COTS GPS receiver module's velocity measurement is generally not practical, there is one aspect of these receiver's behavior that should be characterized for high accuracy state estimation. That aspect is the latency of the velocity measurements. This means characterizing the delay between the time at which the velocity is effectively measured and the time at which that measurement is reported.

The GPS velocity measurement is formed by measuring the carrier-phase-change between two measurement epochs. At the second epoch the change in carrier-phase since the first epoch is converted to a distance and then divided by the time between epochs. This essentially yields the average speed between epochs. Time is required by the receiver to process and report the measurement after the second of the two measurement epochs and this time is the latency of the measurement. Knowing this time is necessary to match the GPS velocity measurement to the correct accelerometer measurements for specific force correction in attitude estimation and to match the GPS velocity measurement to the correct velocity estimates to produce the feedback signal in velocity and position estimation.

Estimating the latency is generally performed with a filter or by direct comparison of the GPS velocity and the velocity derived by some other means over some time period where there is variation in the velocity. For the flight test analysis the latency was measured to be 0.10 seconds using data from the GPS testing described in section 5.2.2 by direct comparison of the GPS velocity output and the known velocity of the GPS based on the arm position and rotation rate.

7.2 Effect of platform propulsion system on sensor measurements

Physical quantities to be measured by state-estimation sensors in a flying sUAS may differ from what is desired to be measured, particularly due to the influence of the sUAS propulsion system. For example, the motion of interest (e.g., the bulk motion of the sUAS flying through the

sky with accompanying changes in attitude, velocity and position) will differ from the fine-scale motion experienced by the gyroscope and accelerometer due to vibration caused by small scale turbulence and by the sUAS propulsion system. Also, many sUAS use electric propulsion systems. Currents flowing in the propulsion system produce magnetic fields which contribute to the total magnetic field at the magnetometer.

7.2.1 Vibration and accelerometer measurements

The most pronounced effect of the propulsion system on the state-estimation sensors flight-test measurement data is seen in the accelerometer measurements. The propulsion system produces unwanted high-frequency linear acceleration through unbalanced rotating components and aerodynamic forces associated with the propeller. The bulk motion of the sUAS airframe in flight is seen to be confined to the lowest frequency bins in these spectra. However, there is considerable power at higher frequencies due to fine-scale motion produced by the propulsion system. Depending on the application and the particular system this fine scale motion may be of interest or may be problematic. For example if wind measurements are desired at 1 meter spacing using an airframe flying at 10 m/s, then a 10 Hz measurement rate is required and the fine scale motion due to vibration would be aliased in the wind measurement if not suppressed in the velocity estimate. Conversely if high-bandwidth wind measurements are desired then the fine scale motion must be accurately estimated.

Figure 7.4 shows the frequency spectra of the accelerometer measurements during the December 2014 flight. The first column is contour plots of the frequency spectra and time with color (red high, blue low) representing the power level at a particular frequency and time. Although it is not easily visible in the contour plots there is considerable power in the 0 Hz frequency bin, particularly for the z axis. The middle column shows the power content of the 0 Hz bin versus time. This can also be seen in Figure 7.5 which shows a three dimensional representation of the x-axis accelerometer spectra as a function of time - the same information presented in the upper-left plot

of Figure 7.4 . The tall, red spikes along the upper right edge of Figure 7.5 are the peaks in the top-center plot in Figure 7.4, and represent the accelerometer information relating to bulk motion of the aircraft. Finally, the right-hand column of Figure 7.4 show the accelerometer-measurement frequency spectra at a particular time (379 seconds) mid-flight.

Although the motor speed is not recorded during the flight test it may be inferred from Figure 7.4. At each time, there are three clear frequency ranges of high power in the x-axis accelerometer measurement spectra. The spectra at each time shown is produced using ≈ 2.5 seconds of data and there may be changes to the motor speed during this time producing a broad 'peak', sometimes with multiple maxima. The first frequency range with significant power content is at/near 0 Hz and is produced by the bulk motion of the aircraft. The second is due to the propulsion system and corresponds to the motor speed. The third is also due to the propulsion system and corresponds to twice the motor speed. The motor has a measured peak static speed of ≈ 9500 RPM, and typical motor speeds in flight are significantly less. The time evolution of the 'peak' corresponding to the motor speed is near 60 Hz (corresponding to 3600 RPM) at time 250 seconds, and increases to near 80 Hz (4800 RPM) at time 700 seconds. The time evolution of the 'peak' corresponding to twice the motor speed is near 80 Hz at time 250 seconds, and decreases to near 40 Hz at time 700 seconds. As the sampling frequency is 200 Hz a signal at 120 Hz is aliased down to 80 Hz and a signal at 160 Hz is aliased down to 40 Hz, so the two peaks are consistent with respect to motor speed and twice motor speed when aliasing is considered. Note that, as also expected, the frequency variation in the time evolution of this 'twice motor speed' peak is double that of the frequency variation in the time evolution of this 'single motor speed' peak. It is of interest that the effect of propulsion system induced fine-scale motion is predominantly at the frequency corresponding to twice the motor speed, or at the frequency of one of the two propeller blades passing a particular location. With the test flight airframe this may be accentuated due to the "pusher" propeller arrangement and the fact that the propeller blades experience disturbed air at two locations - behind each wing.

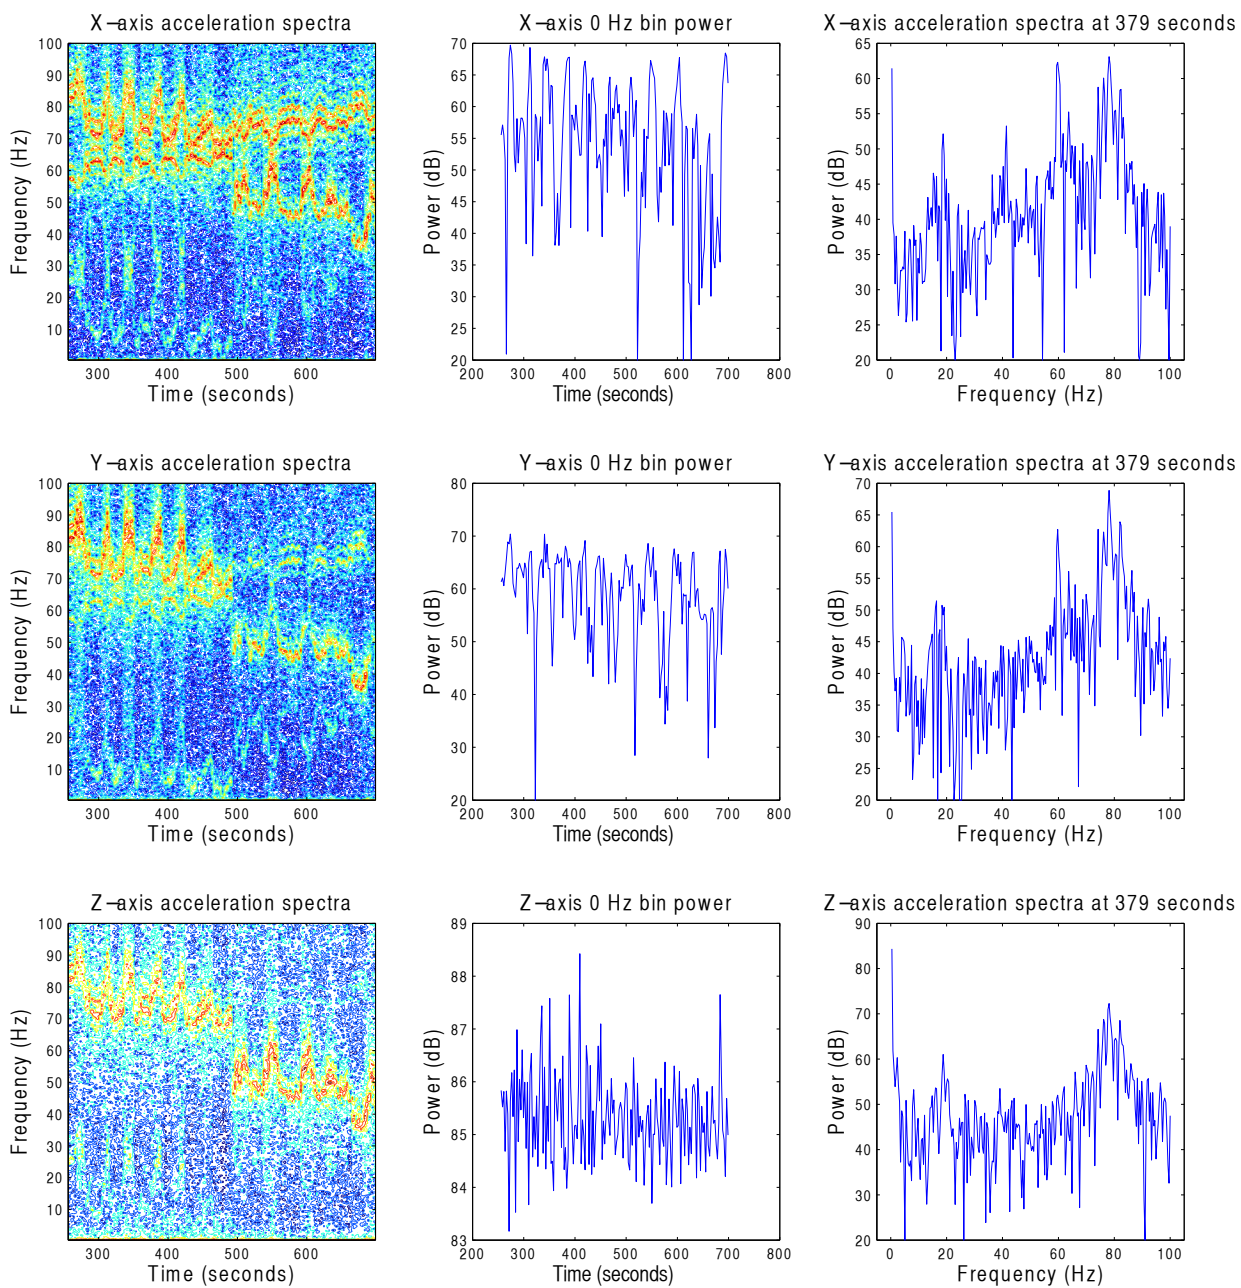


Figure 7.4: Accelerometer measurement spectra




Figure 7.5: X-axis accelerometer measurement spectra

The consequence of the significant high-frequency power content in the accelerometer signal differs with the application of the accelerometer information. When used for observation of the gravity vector which is used as a feedback correction signal in an attitude filter the low-frequency content is of interest. High-frequency content may be filtered, either with a low-pass filter applied to the accelerometer signal before it is used in the attitude filter, by the filter structure itself, or both. When used as a high-frequency sensor for velocity estimation, the high-frequency content in the accelerometer signal represents real high-frequency motion experienced by the sensor. If the accelerometer is rigidly connected to the relative wind sensor, so that both sensors are experiencing the same high-frequency motion, then the high-frequency content of the accelerometer signal represents important information about platform motion and should be preserved. However, the aliasing of the motion occurring at twice the motor speed (frequency) represents a corruption of this information. If an accurate high-frequency velocity estimate is required, then it is necessary to avoid this aliasing either by sampling at a higher frequency or operating at lower motor speeds.

A higher sampling rate was used for subsequent test flights to avoid aliasing the motor vibration signals.

7.2.2 Vibration and gyroscope measurements

Figures 7.6 and 7.7 show the gyroscope measurement spectra in the same formats as Figures 7.4 and 7.5 showed the accelerometer measurement spectra. In Figure 7.6 the vibration signature at the motor frequency can be clearly seen, and there is no apparent vibration signature at twice the motor frequency in the gyroscope measurements. There is also an interesting phenomena visible in the spectra; the vibration signature from the motor may also be seen replicated and at reduced power at 30 Hz below the real signature and at 30 Hz above the real signature. A replicated signature is also present at 60 Hz below the real signature, although this replicated signal is obscured by the real low-frequency content of the signal. These replicated vibration signals are due to (improper) resampling of the gyroscope measurement. This unintentional resampling occurred as the micro

controller firmware obtained measurement samples from the gyroscope at 200 Hz, but the gyroscope produced internal samples at 380 Hz. Proper methods exist to change the sampling frequency. If the input and output sample frequencies are known then an interpolate/decimate scheme may be used [58]. Approximate methods have also been developed for cases where one or both rates are unknown which maintain high signal-to-noise ratios [8]. However, Osborn [54] suggests that while asynchronous resampling may be necessary in some cases, such as internal to a 3 axis gyroscope to combine signals from the three gyroscopic elements each operating with a sample rate equal to its own natural frequency, it is preferable to design a system to avoid asynchrony. This may be impractical for a sUAS state estimation system where sensor measurement rates may be driven by separate sensor clocks, but may be accommodated with most sensors using a polling or interrupt scheme.

If an analog sensor signal is band limited, then the Shannon sampling theorem [68] allows for theoretical perfect reconstruction of the signal from a set of sampled measurements, provided that the samples are taken at high enough frequency to satisfy the Nyquist criterion. A signal sampled at a particular rate may be changed to be sampled at a different rate through a process of reconstruction and sampling the reconstructed signal, (theoretically) without loss of information or introduction of noise. Signal reconstruction may be effected by passing impulses, weighted by the discrete-time measurement values, through an ideal low-pass filter. This process may be approximated wholly in the digital domain to arrive at a signal with a higher sample rate. By adding $M-1$ zero values between measurements in a sample series and passing the augmented series through an ideal low-pass filter produces a new (intermediate) measurement series at a sample rate equal to M times the original sample rate. Low-pass filtering is required as the augmented series spectra will contain the original spectra replicated at higher frequency. If this intermediate series is then subsampled by taking every L 'th sample, a final measurement series is produced at a sampling rate differing from the original by the rational ratio M/L . Assuming this final sampling rate respects the Nyquist criteria no information is (ideally) lost. The difficulty arises, of course,

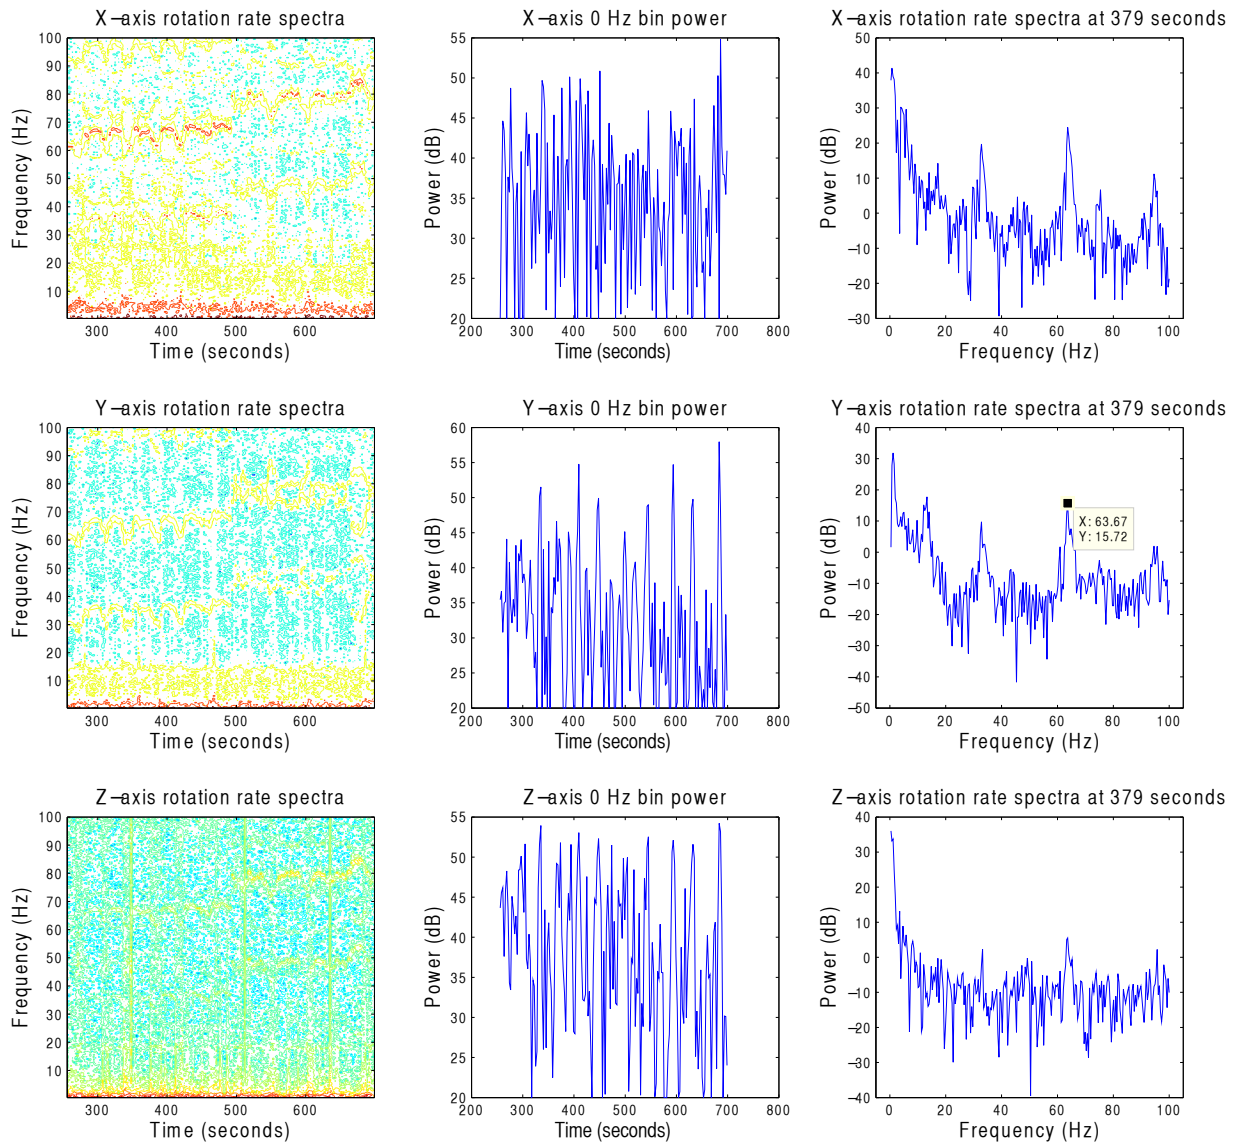


Figure 7.6: Gyroscope measurement spectra

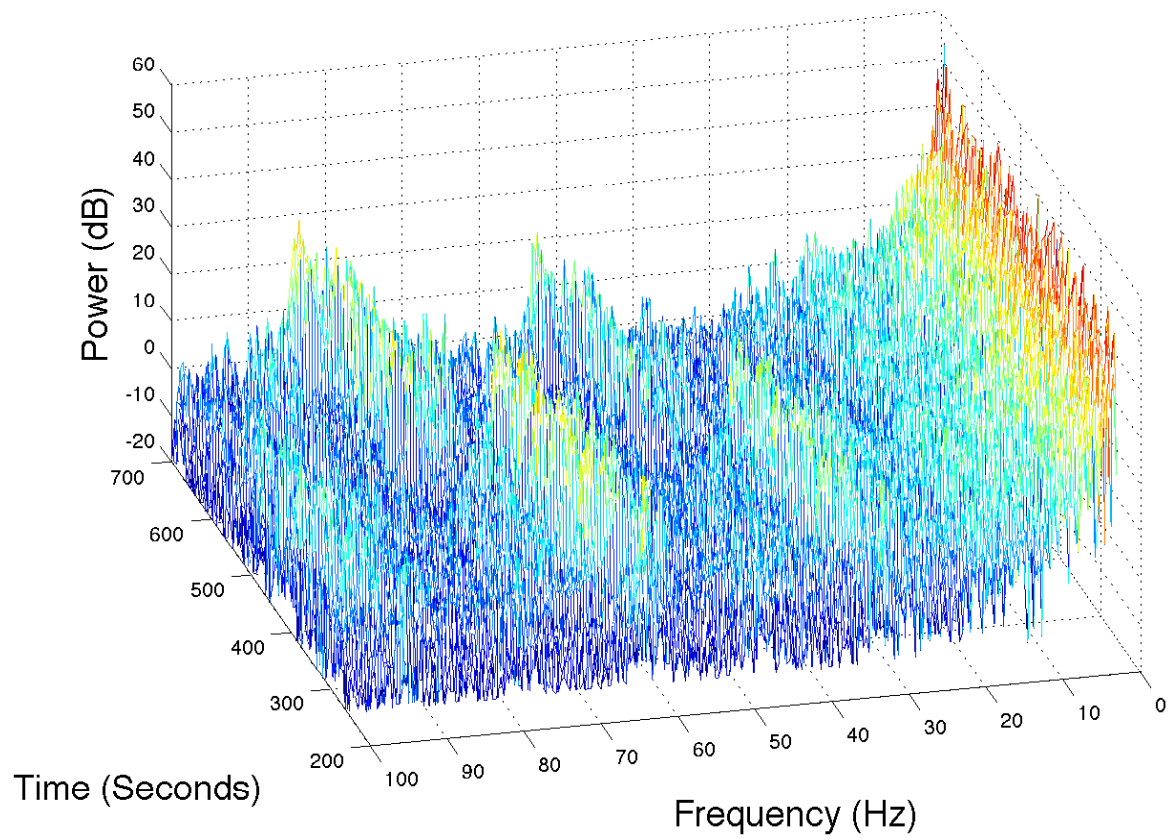


Figure 7.7: X-axis Gyroscope measurement spectra

due to the lack of an ideal low-pass filter. Regardless, this method is frequently used to resample digital signals and performance of the method is dependent on the input signal spectra and the low-pass filter characteristics.

A similar method allows for resampling when the desired sample rate differs from the original rate by a non-rational factor, or when the value of the original signal at an arbitrary time is required. This method, published by Smith [69] as a public-domain algorithm, involves convolution of a low-pass filter impulse response with the original time-series. The impulse response of an ideal low-pass filter is a sinc function extending from $t = -\infty$ to $t = \infty$. Smith truncates this impulse response, for example at the fifth zero crossing to the left and right of the origin, and applies a Kaiser window taper. This results in good performance, for example -80 dB stop-band rejection for truncation at the fifth zero crossing to the left and right of the origin, but does result in significant filter delay.

An alternative when less performance is required is to use a simple two-point linear interpolation filter. Linearly interpolated fractional delay is equivalent to filtering and resampling a weighted impulse train (the input signal samples) with a continuous-time filter having the simple triangular impulse response [69] and the frequency response of the interpolation is given by

$$H(f) = \text{sinc}^2(fT). \quad (7.9)$$

Note that the main lobe of this frequency response extends up to the sampling frequency. This theoretical analog output is also sampled, thereby aliasing through the sides of the the sinc^2 frequency response main lobe and through the side lobes. As an example of the effects of this process consider a signal with components at 0 Hz and 5 Hz (representing bulk motion of the airframe), 80 Hz (representing vibration motion at the motor frequency) and 160 Hz (representing vibration motion at twice the motor frequency). Figure 7.8 shows a discrete Fourier transform of this signal sampled at 500 Hz, as well as the discrete Fourier transform of this signal resampled at 400 Hz using linear interpolation. It can be seen that attenuated, aliased versions of the signal components are present, but are attenuated by at least 18 dB. When resampling the signal without interpolation, as in the flight test data collection presented above, these attenuated, aliased versions of

the signal components are present at the same frequencies, but are less attenuated, with the largest attenuated by 10 dB and with a lesser rolloff on the attenuation of other components relative to the interpolated results.

The data collection and digital signal processing scheme was changed to use a sensor-driven-interrupt driver for subsequent test flights to provide synchronized measurement and retrieval of data and remove the resampling issue.

7.2.3 Electrical fields and magnetometer measurements

The test flight airframe used an electric propulsion system and magnetic fields produced by currents flowing in the airframe is a potential concern. However, by examining the magnetometer measurement data during the pre-flight period (motor off) and during flight it was determined that the airframe electrical systems had a minor effect on the magnetometer measurements. The propulsion system involves both direct current, flowing from the battery to the electronic speed controller, and alternating current, flowing from the speed controller to the motor. While the switching rate for the electronic speed controller used is not known, switching rates for this type of controller is typically in the range of one kilohertz to tens of kilohertz. Magnetic fields produced by alternating currents in those frequency ranges appears as high frequency noise in the magnetometer data. The high frequency noise observed in the magnetometer measurement magnitude during flight is roughly double that observed during pre-flight (motor off). Magnetic fields produced by direct currents have a fixed direction in the airframe body frame, and will vary with the current (motor power setting). These fields will appear as a variation in the observed magnetometer measurement magnitude which is dependent on the airframe orientation with respect to the Earth's magnetic field vector. Such variation in the flight test data has a peak to peak value of approximately 0.015 Gauss, which equates to a maximum directional error of approximately ± 0.6 degrees. This error may be reduced by performing the magnetometer calibration while the propulsion system is producing currents typical of cruising flight conditions, although doing so in the field may involve

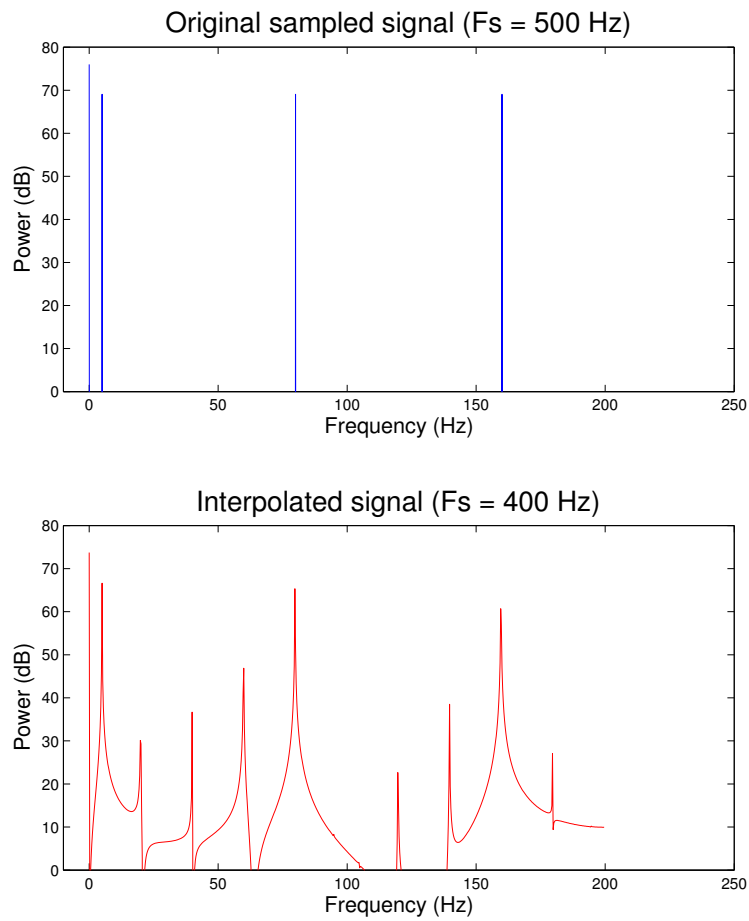


Figure 7.8: Linear interpolation resampling example frequency response

the hazard of performing the calibration while the propeller is turning.

7.3 Sensor fusion using actual flight data

Attitude-estimate sensor fusion was performed with the first stage of a cascaded filter **Type** using a 7-state EKF filter **Structure**. A time-differenced GPS-velocity-measurement specific force correction is employed as described in section 4.2.4. The sensor fusion algorithm is implemented in MATLAB and relies on Equations (4.7), (4.8), (4.10), (4.16), (4.17), (4.34) and (4.35). Sensor-fusion analysis of the test flight data is restricted to the first (AHRS) stage as the optical reference vector sensor, described in chapter 6, is used for validation of the attitude estimate in section 7.4. No similar outside reference capability is available for validation of a velocity/position estimate.

Performance of the attitude-estimation filter is driven by the quality of the sensor signals. As noted at the beginning of the chapter the sensors used in the test flights are not the highest quality available, but are of interest due to their potential use in a University of Colorado autopilot. The noise performance of these sensors measured in the field during pre-flight data collection was:

- Gyroscope 1- σ noise: 0.0055, 0.0045, 0.0056 (rad/sec) x/y/z axis
- Accelerometer 1- σ noise: 0.240, 0.232, 0.239 (m/sec²) x/y/z axis
- Magnetometer 1- σ noise: 0.0015, 0.0017, 0.0015 (Gauss) x/y/z axis
- GPS velocity 1- σ noise: 0.04, 0.03, 0.07 (m/s) north/east/up

It is interesting to note that while the gyroscope's measured noise performance was within the device's data sheet specification, the accelerometer had significantly more noise than might be inferred from that device's data sheet. The accelerometer data sheet specifies the acceleration noise density only when operating with a full scale value of ± 2 g, but the accelerometer was operated with a full scale value of ± 4 g. While the noise density when operating with a full scale value of ± 4 g might be expected to be twice the noise density when operating with a full scale value of ± 2 g it was in fact greater by a factor of 8. The magnetometer's data sheet provides no specification for

measurement noise, but the noise measured in the field during pre-flight data collection was close to the specified magnetic field resolution.

Expected attitude-estimation results must consider the sensor noise levels described above and also additional noise due to propulsion system effects on sensor measurements. The accelerometer noise in flight is assumed to be four times greater than during pre-flight as a conservative estimate in view of the considerable amount of aliased signal power visible in Figure 7.4. Although the sampling rates were raised and the resampling issue corrected for March 2015 flights there is still a significant amount of power attributable to higher order vibration harmonics. To deal with this aliased signal power the gyroscope noise in flight is also treated as four times greater than during pre-flight. While the aliased signal power is not white noise, raising the effective noise used in calculation of the filter gain decreases the effect of this spurious signal. Section 7.2.3 described that the high frequency noise seen in the magnetometer measurement during flight was twice that seen during pre-flight measurements. Uncalibrated hard iron effects caused by motor currents were also observed and described in Section 7.2.3, so a magnetometer bias of 0.0075 Gauss is included in the simulation to account for this error source.

Using these adjusted noise values the simulation described in Appendix 4 and used for performance comparisons in chapter 4 is used to predict the attitude estimation accuracy for the sensor fusion algorithm operating on the test flight data. The simulation was run for 500 Monte Carlo samples using the 3 m/s wind truth data from chapter 4 and predicted r.m.s. error levels are 0.35° roll error, 0.55° pitch error and 0.58° yaw error.

7.4 Validation using optical reference vector sensor

Validation of the expected attitude-estimate accuracy was a prime motivation for flight tests and data from the optical reference vector sensor described in chapter 6 was also collected during test flights for this purpose. The measured reference vector direction is compared to the reference

vector direction estimated with the estimated attitude and position, and this comparison provides information about the quality of the attitude estimate. However analysis of the flight test data, and subsequent testing of the U-blox LEA-6H GPS receiver module, showed that uncertainty in the comparison due primarily to error in the GPS position measurement does not allow for the attitude-estimate error to be accurately constrained. It is still possible, however, to make the statement that the difference between the measured and estimated reference vector directions is no larger than expected when accounting for the measurement error in the optical reference vector direction, the GPS position measurement error, and the expected attitude-estimate error. This distinction is clarified below.

Figure 7.9 shows the optical reference vector measurements and estimates as angles in the body-xz plane and the body-xy plane. Each row in this figure is data collected during a different overhead pass where a significant number of valid optical reference vector sensor measurements are available. The times these passes occurred (after take-off) is listed on the left. The first column shows the reference vector angle in the x-z plane and the second column shows the reference vector angle in the y-z plane. The third (rightmost) column shows the difference between the measured and estimated angles in both the x-z plane and y-z plane. Figures 7.10 and 7.11 show enlarged views of the subplots in the fourth row of Figure 7.9. In the first two columns the measured vector direction angle is plotted as a red line. This line is surrounded with a red band indicating the measurement accuracy. The thickness of the band varies depending on the range from the ground-based beacon to the sUAS, and the resulting beacon signal-to-noise ratio. The estimated reference vector angle is plotted as a black line. This estimate is computed using the attitude estimate and the GPS position measurement. The GPS position measurement error causes an uncertainty in the estimated angle and this uncertainty is shown as the grey band. The thickness of this band varies with the height at which the pass was made.

Prior to analyzing flight test data the GPS position measurement error (using a U-blox LEA-6H receiver) was not believed to be a limitation to validating the attitude estimate by this method.

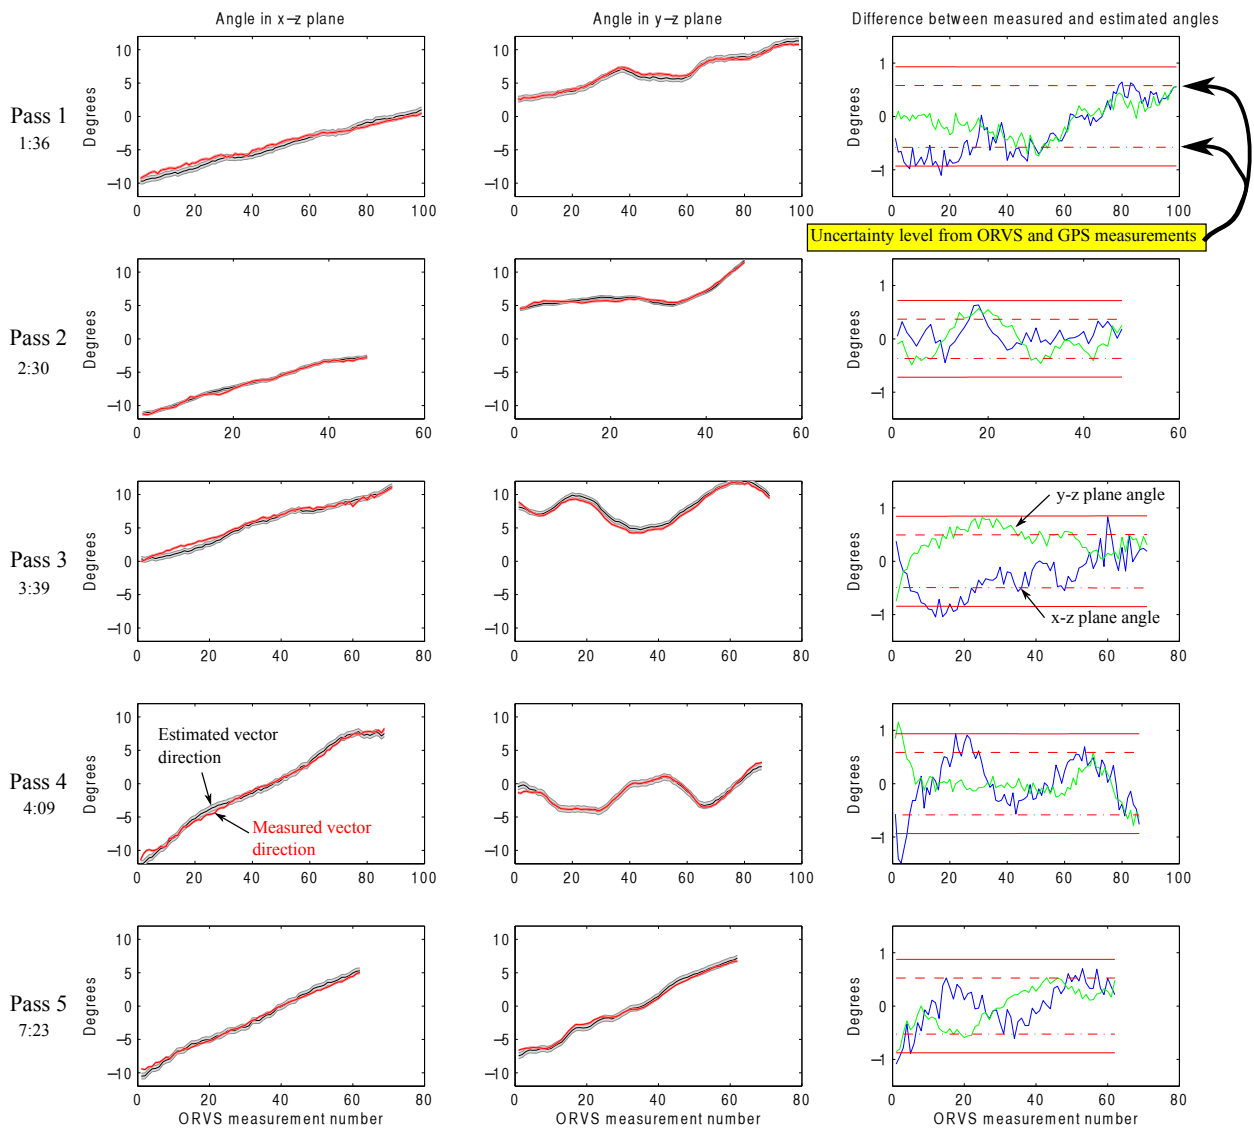


Figure 7.9: ORVS measurement versus estimated reference vector based on attitude estimate and GPS position for 5 overhead passes during the March 14, 2015 flight.

However analysis of the data called this assumption into question. Two identical LEA-6H receivers with identical settings were tested over a 10 minute period while co-located and stationary on the ground. Significant drift between their position measurements was noted over this time period and although drift in the position measurement was expected the fact that the two receivers drifted in a generally uncorrelated manner was surprising. If not for this effect then a stationary ground-based receiver could be used to correct for drift in the position measurement, significantly reducing uncertainty in the comparison made above. Unfortunately this does not appear to be valid approach with regular COTS GPS receivers; this approach could be pursued with more sophisticated GPS RTK receiver sets suitable for sUAS now becoming available. The data presented in Figure 7.9 was produced by using a gradient descent optimizer to find the most likely GPS position measurement drift at the time of each pass. Testing the U-blox LEA-6H GPS receiver showed that the position measurement drifted as a rate as high as 0.4 m/s while stationary; it is reasonable to assume that the receiver position measurement may drift at this rate while in flight as the dynamic motion's affect on the receiver tracking loops would be expected to increase noise and drift. The width of the grey bands in Figures 7.9 and 7.10 was computed using a 0.4 m/s drift rate over the time-length of each overflight path and the relevant geometry to calculate the angular uncertainty.

Figure 7.11 shows the difference between the measured and estimated reference vector angles in the x-z and y-z planes. Red dotted lines are included that bound the difference in the angles based on the (1σ) uncertainty in the optical reference vector sensor measurement and the potential drift in the GPS position measurement. Solid red lines are also plotted 0.35 degrees outside the dotted red lines and represent the maximum expected difference between the measured and estimated angles (based on an expected attitude-estimate accuracy of 0.35 degrees in roll - see section 7.3). The maximum expected difference of the angle in the x-z plane would be somewhat higher due to the expected pitch accuracy of 0.55 degrees. Figure 7.9 only shows excursions of the y-z plane angle difference (green line) at the beginning of the fourth pass. This could occur because of the attitude-estimate error exceeding a 1σ level. It may also be the case that the ORVS measurement

has unusually high error; the measurement is at the extreme edge of the calibrated field of view and there may be an issue with a reflection in the optics in this case.

Aside from these few measurements at the beginning of the fourth pass there are no other measurements which would indicate that the attitude-estimation error is above predicted levels. In fact, the minimal excursions outside the predicted 1σ error bound suggest that this 1σ level is conservative. The method used to estimate the error contribution of GPS-position-measurement-error drift (based on the highest levels seen when the receiver was stationary) may be too conservative.

While the flight test setup used herein could not provide a validation of the attitude estimate to a desired level of performance (e.g., to near the predicted level of performance), the analysis clearly shows that the methods used here may be adapted to do so. If the dotted red error bound in Figure 7.11 were significantly reduced so that the error due to the attitude estimate were the dominant error in the vector comparison, then the attitude estimate can be evaluated without excessive uncertainty. For example in the case above, if the uncertainty in the vector comparison due to the optical-reference-vector measurement error and due to the GPS position measurement error was reduced to 0.05 degrees then the attitude-estimate could be bound using the flight test data to approximately 0.4 degrees versus a predicted 0.35 degrees. This would be of much greater benefit that the claim that can be made from Figure 7.11 - that the attitude-estimate error is less than or equal to 0.9 degrees.

Reducing the optical-reference-vector measurement error and the GPS position measurement error to these levels may be reasonable. This would be a straightforward matter for the optical reference-vector sensor system. If the optical beacon power was increased by an order of magnitude, or the optical beacon was reconfigured to use a steered, narrow beam such that the power density arriving at the aircraft was an order of magnitude larger, then the error in the optical-reference-vector measurement would be reduced by nearly an order of magnitude. Reducing the GPS position-measurement error by an order of magnitude with COTS GPS receivers is currently impractical. However relatively inexpensive carrier-phase real time kinematics GPS (RKT-




Figure 7.10: Enlarged view of Pass 4 measured and estimated angles from Figure 7.9




Figure 7.11: Enlarged view of Pass 4 difference between measured and estimated angles from Figure 7.9

GPS) systems may become available which are suitable for sUAS and may provide centimeter-level position-measurement accuracy. Currently the Piksi system from Swift Navigation has been released as a beta-development product. Swift Navigation has not released detailed specifications at this time, but the current Piksi data sheet [1] indicates centimeter accurate relative positioning. Other products, such as the Trimble BD920 receiver, while somewhat large and expensive for sUAS use, show a trend towards sUAS appropriate RTK-GPS receivers.

7.4.1 Filter consistency

Measurements from the optical reference vector sensor allow the error in the attitude solution to be bounded, at least during those times when measurements are made by the optical reference vector sensor. Another method was used to look for any potential problems in the attitude filter performance during other times in the flight. The phenomena of divergence in an extended Kalman filter is well known and is usually the product of modeling errors (particularly poor choice for process noise or sensor error covariance, both of which are often treated as tuning parameters), numerical errors or programming errors. A statistical characterization of the state-estimate error is desirable; a state estimation filter is said to be consistent when the state errors have zero mean and magnitude commensurate with the state covariance as yielded by the filter [5]. For the filter run on real flight data without a source of truth data related criteria may be used to judge consistency: The observation measurement innovations should be zero mean, commensurate with the sensor error covariance, and acceptable as white.

The two observation measurement sensors (the accelerometer and magnetometer) are considered separately. First, the magnetometer innovations across the time period comprising the bulk of the flight are shown in Figure 7.12, have a mean value of $[6 * 10^{-4} \quad -3 * 10^{-4} \quad 1.5 * 10^{-3}]$ Gauss and standard deviation of $[0.0056 \quad 0.0058 \quad 0.0042]$ Gauss. The mean value, while non-zero, is close to zero and likely the result of residual calibration errors due to hard-iron effects from the operating electric propulsion system. The standard deviations are of the order expected but somewhat large.

This is likely a further effect of the propulsion system on the magnetometer and perhaps expected when viewing the data of Figure 7.12 over shorter timespans where some apparent structure is apparent. Examining the frequency spectra of the innovations, for example Figure 7.13, there is significant power concentrated in several frequency bands.

Similarly, but more dramatically, the accelerometer innovations clearly show the effect of vibratory motion due to the propulsion system captured by the accelerometers. This is quite expected considering the significant amount of vibration energy seen in the accelerometer measurements as shown in Figure 7.4.

The consistency tests point out that the accelerometer and magnetometer measurements have error characteristics which are not simply sensor noise. Failure of these statistical tests does not indicate that chosen filter is a poor choice for this problem; rather, there are known but un-modeled effects which do not fit well with the base assumptions leading to development of the EKF. Still the EKF does well in this application. This can be seen in the relatively benign behavior of the magnetometer innovations over the course of the flight. There is no time period where the innovation characteristics change significantly which would be the case if there were a trend towards divergence of the filter.

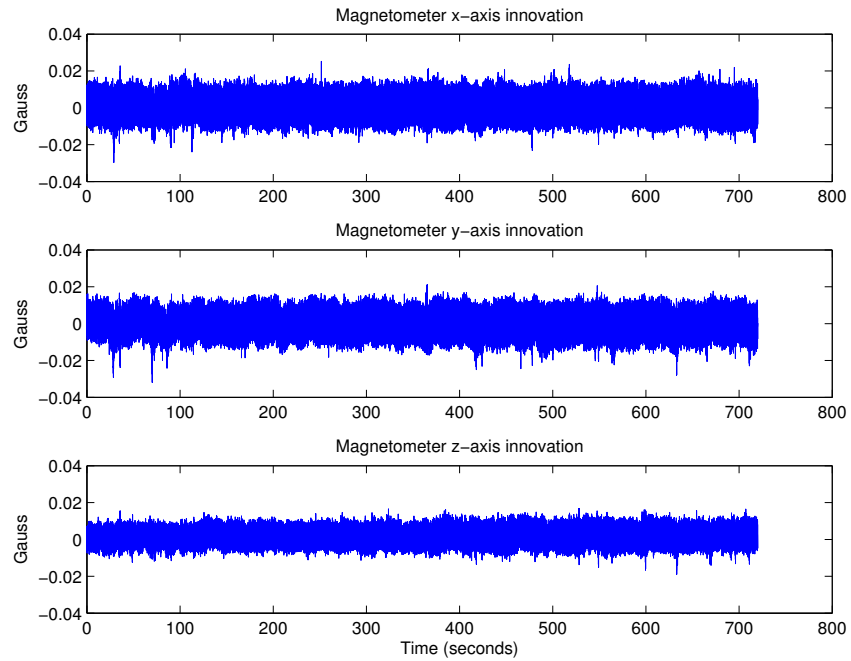


Figure 7.12: Magnetometer measurement innovations during flight




Figure 7.13: Frequency spectra of the x-axis magnetometer measurement innovations during flight

7.5 Flight test analysis summary

Analysis of the data collected during flight tests produced several noteworthy results that are summarized here.

- The raw magnetometer data showed significant need for calibration corrections. Using a 3-coefficient calibration to correct for zero point bias and hard-iron effects greatly improved the magnetometer data, but still left considerable residual error. A 6-coefficient calibration algorithm, suggested by some researchers as sufficient for sUAS, also proved unsatisfactory. An extension to Gegre-Egziabher et al's [23] 6-coefficient calibration algorithm that provides a 12-coefficient calibration is presented in section 7.2 which appears quite effective. This calibration algorithm may also be used for the accelerometer. If the airframe uses an electric propulsion system and calibration is performed while the propulsion system is off, then some hard-iron distortion is expected in flight due to electric fields produced by the propulsion system
- A significant amount of high frequency motion is induced in the airframe by the propulsion system. Acceleration signals at the motor speed frequency and twice the motor speed frequency, and rotation rate signals at the motor speed frequency are prominent in the inertial sensors' data. Two problems were encountered with respect to these high frequency motion signals with respect to the inertial sensors' data collected during the flight tests. First, the sampling rate was too low during early flight tests, such that the signal at twice the motor frequency was aliased in the sampled data. Since the signal was aliased and no longer represented true motion of the airframe it could only be treated as noise. Second, during early flight tests the data from the inertial sensors was unintentionally resampled at a rate different from the rate at which it was produced. This resampling produced attenuated, aliased copies of the high frequency signals at multiple locations in the (re)sampled signal spectra. Again, as these aliased signals do not represent true motion of the airframe they

can only be treated as noise. Proper system design can circumvent both of these problems.

- An attitude-estimation filter run on the collected sensor data produced a sensible attitude estimate. Quality of the attitude estimate was predicted based on simulation of the attitude-estimation filter using noise levels representative of those in the test flight data. Evaluation of the attitude estimate by comparing a reference vector measurement made with an optical reference vector sensor to an estimate of the reference vector produced with the attitude estimate, and a position estimate based on GPS position measurements, bound the attitude-estimate error at 0.9 degrees and did not provide any indication that the attitude-estimate error was larger than the 0.35 degree (roll) predicted accuracy level.
- Use of the optical reference-vector sensor in a sUAS was demonstrated and shown to be useful for validating sUAS attitude-estimation systems. With a higher-accuracy GPS system the ORVS may be used to validate sUAS attitude-estimation systems to 0.1 degree level accuracy.

Chapter 8

Summary and conclusions

Small unmanned aircraft systems are an emerging technology and offer an attractive platform for wind measurement. sUAS have many cost, safety and operational advantages over manned aircraft and consequently are ideally suited to collecting wind measurements in many situations for a variety of fields of research. Wind measurement from sUAS has the same sensitivity to knowledge of the aircraft velocity, attitude and attitude rate as other forms of airborne wind measurement, but it is not practical for sUAS to carry navigation grade sensors like manned aircraft to accurately determine these quantities. This difficulty leads to the primary research questions addressed in this work. In particular, this work focused on three research areas each with direct applicability to achieving high-accuracy from sUAS-based wind measurement systems:

- Attitude estimation using small, low-cost sensors,
- Velocity estimation using small, low-cost sensors, and
- Development of a new type of sensor suitable for inclusion in a sUAS attitude-estimation system to improve accuracy or for use validating a sUAS attitude-estimation system.

In addition, a state-estimation system and a prototype optical reference-vector sensor system were constructed and tested during flights in a sUAS; the experimental data supports results from the analytic research and simulation, and provides insight into practical implementation issues. The following sections contain summaries and conclusions from each of these areas of research.

8.1 Summary and conclusions: sUAS attitude estimation

Analysis of wind-measurement sensitivity to attitude knowledge shows that UAS-based wind-measurement accuracy on the order of 1 cm/s requires attitude-estimation accuracy on the order of 0.1° . Sensors suitable for use in sUAS, while quite good considering their size and cost, are poor compared to larger, navigation-grade sensors. Consequently, the choice of sensor-fusion algorithm is quite important. Based on the typical set of contemporary sUAS attitude-estimation sensor-fusion sensors (MEMS gyroscope and accelerometer, magnetoresistive magnetometer, and single-frequency GPS receiver) the analyses in Chapter 4 produce three main conclusions.

- There are a variety of attitude-estimation sensor-fusion **structures**. Some commonly used for sUAS attitude estimation are complimentary filters, extended Kalman filters and unscented Kalman filters. These different filter **structures** all integrate the gyroscope measurements and use a feedback-correction signal calculated using information from the accelerometer and magnetometer. Filter **structures** differ in the way this feedback-correction signal is calculated and this introduces different levels of error into the attitude estimate. Considering a hierarchy of filters, it is found that at some point no significant improvement may be gained by using a more complex and computationally expensive filter as the difference in error introduced by the choice of filter **structure** is negligible compared to the error in the attitude estimate introduced by the sensor errors. For the case of sUAS attitude-estimation systems with contemporary sensors the extended Kalman filter produces better results than less complex filters (e.g., complimentary filter), but no meaningful improvement is gained by using an unscented Kalman filter or other more computationally expensive filters.
- There are also two main **types** of sUAS attitude-estimation filters; those that produce an estimate of the position and velocity as well as the attitude in a single stage and those that produce an estimate of the attitude in a first stage and estimate position and velocity in

a second stage. It is shown analytically, based on the sUAS sensor set mentioned above, that a filter which estimates position, velocity and attitude will necessarily produce a less accurate attitude estimate than a filter estimating attitude only. This is due to the differing way that the two filter **types** use the accelerometer information. In a filter estimating position, velocity and attitude the accelerometer information has only a weak influence (or, depending on the filter **structure**, no influence) on the attitude estimate. In a filter estimating attitude only the accelerometer measurement is a significant part of the feedback correction, and significantly improves performance. It is further shown that for sUAS attitude estimation with contemporary sensors and flight in wind the difference between these approaches is very significant with difference in accuracy levels as high as an order of magnitude.

- Direct observation of the gravity vector is desirable for attitude estimation but accelerometers measure specific force that includes both gravity and translational acceleration. The method of handling the translational acceleration in the attitude-estimation filter is often overlooked but is a critical consideration for high-accuracy attitude estimation for sUAS flown in wind. Traditional methods of estimating the translational acceleration are particularly vulnerable to wind-gust-induced rotation. It is shown that time-differenced GPS velocity measurements are an appropriate method of estimating translational acceleration for sUAS attitude estimation and can reduce error by as much as an order of magnitude versus traditional methods for sUAS flown in medium and high-turbulence conditions.

Using the recommended approach for a sUAS attitude-estimation filter (an EKF AHRS with time-differenced GPS velocity measurements used to estimate translational acceleration) simulation results show that attitude-estimation accuracy on the order of 0.1° is possible using contemporary sensors suitable for sUAS. This is a key result leading towards the ability to make high-accuracy wind measurements from sUAS.

8.2 Summary and conclusions: sUAS velocity estimation

Airborne wind-measurement accuracy has a 1-to-1 sensitivity to error in knowledge of the platform velocity. sUAS velocity estimation is similar to sUAS attitude estimation, but where attitude estimation relies on integration of gyroscope measurements for high-frequency information, velocity estimation relies on integration of accelerometer measurements. GPS is the only appropriate sensor available to provide a strong feedback signal for accelerometer integration drift correction in sUAS velocity estimation and GPS velocity-measurement accuracy is a significant limitation to sUAS velocity-estimation accuracy. Using contemporary MEMS accelerometers, $O(1 \text{ cm/s})$ velocity-estimation accuracy requires GPS velocity-measurement accuracy of $O(10 \text{ cm/s})$.

sUAS velocity-estimation systems generally rely on commercial off-the-shelf GPS receivers and detailed information about the operation and performance of these receivers is not available. GPS accuracy specifications are usually made in a statistical manner; for example a specification of velocity measurement accuracy might be that the accuracy will be better than 0.01 m/s at least 50% of the time, without any information about what conditions might cause a particular measurement to be included or excluded from the 50% group. Consequently, a deeper understanding of the performance available requires characterization of a particular GPS receiver model.

The ability of a GPS receiver to provide high-accuracy velocity measurements while operating in a sUAS-in-flight environment is the critical factor for high-accuracy sUAS velocity estimation. A particular GPS receiver (U-blox 6T) was tested to evaluate the velocity-measurement accuracy while the receiver was in accelerating motion. Testing showed that the velocity-estimation error is highly dependent on the number of satellite channels with phase-locked versus frequency-locked Doppler tracking. Further, the testing showed that when the receiver is subject to a sufficient level of acceleration the tracking loops are not able to maintain phase lock, so the velocity-measurement accuracy is affected by accelerating motion. However, correlation between acceleration level at modest levels of acceleration and velocity-measurement accuracy is poor. A much better predictor

of a satellite channel having phase lock (so long as acceleration is modest) is the channel's signal-to-noise ratio. For the tested commercial-off-the-shelf GPS receiver 10 cm/s velocity-measurement accuracy requires phase-lock for a strong majority (95+%) of individual satellite channel measurement epochs. However, during flight tests this did not appear to be problematic, and this criteria can be monitored during flight to validate the resulting wind-measurement accuracy.

An additional result was the determination that for platforms with high dynamic motion, the GPS receiver's internal filtering of the velocity measurement can add significant error. It was demonstrated that for receivers that provide raw measurement data a velocity solution may be computed externally with significantly less error than the GPS-receiver-produced (filtered) measurement.

8.3 Summary and conclusions: Optical reference-vector sensor

Contemporary sensors used in sUAS are very good considering their size and cost, but their performance is a limitation to high-accuracy state estimation. Two uses motivated development of a new sensor system suitable for use in sUAS. The first motivation is to provide a high-accuracy sensor with (partial) observability of the platform attitude for use in validating the sUAS attitude-estimation system while in a sUAS-in-flight environment. The second motivation is to provide an additional sensor which may be incorporated in sUAS state-estimation systems to reduce the state-estimate error, particularly with respect to the attitude estimate.

The result of this development effort was the production of a prototype optical-beacon sensor system. This system provides a reference-vector-direction measurement which is unaffected by platform motion, including vibration, and other effects which are problematic for other sUAS sensors. The system is able to operate in full daylight conditions over a range of hundreds of meters. The reference vector direction can be measured with an accuracy of better than 0.5% of the sensor's field-of-view, resulting in 0.1° accuracy over a 30 degree range, and a measurement

update rate of up to 200 Hz. Even higher accuracy is available at shorter ranges. The prototype system was used in the field tests described in Chapter 7.

While the prototype sensor was developed specifically to support attitude-estimation systems for sUAS being flown outdoors in wind, the sensor system also has utility for a wide variety of other purposes. Numerous trade-offs were identified during the development process which allow the design to be changed to meet other performance criteria. With further development the system can even be expanded to provide a full pose estimate without other sensors.

8.4 Summary and conclusions: field (flight) tests

sUAS attitude-estimation systems have generally been poorly characterized heretofore in that they have been evaluated by one of two methods, neither of which is wholly suitable. Some have compared sUAS attitude-estimation systems against high-accuracy attitude-estimation systems when flown in larger UAS or manned aircraft. However, the analyses in Chapter 4 show that sUAS attitude-estimation-system accuracy is dependent on the motion environment and larger, heavier aircraft have less wind-gust-induced motion. Others have evaluated sUAS attitude-estimation systems using simulation only, and often with optimistic sensor models.

Field tests were conducted where a sUAS was flown while instrumented with typical sUAS state-estimation-system sensors and the prototype optical reference-vector sensor described above. All data was recorded for post-flight analysis. The primary purpose of collecting and analyzing field test data was demonstrating sUAS-attitude-estimation capabilities predicted in Chapter 4, as detailed in section 7.4. Two measures of the vector direction in the body frame of the vector from the sUAS to the ground-based beacon are produced. The first is produced directly by the optical reference-vector sensor. The second is produced using the known, fixed position of the beacon and the GPS-measured position of the airframe to calculate the vector in the inertial frame. This vector is then rotated to the body frame using the attitude estimate. By comparing the two

vector directions and accounting for the error introduced by the optical reference-vector sensor and the GPS position measurement a bound may be placed on the attitude-estimate error. While the uncertainty introduced by the GPS position measurement did not allow the attitude-estimate accuracy to be bound as tightly as desired this comparison did show good agreement between the predicted and observed performance.

Also, the process of producing and flying a physical sUAS attitude-estimation system and analyzing its performance produces several interesting results pointing out the difficulties inherent in actual sUAS state estimation. In particular:

- Although suggested otherwise by the literature, sUAS airframes may have sufficient ferrous material to require a magnetometer calibration correcting soft-iron distortion effects. This was the case for the flight-test system. A new 12-coefficient magnetometer-calibration algorithm suitable for use with sUAS was developed and shown to significantly reduce magnetometer-measurement error when compared to 3-coefficient and 6-coefficient magnetometer-calibration algorithms.
- A particularly interesting problem attributed to a combination of significant propulsion system vibration coupled with improper digital signal processing was identified. Failing to synchronize digital measurement samples passed from the gyroscope sensor to the data-collection micro controller resulted in unintentional resampling of the data. The consequence was attenuated replication of the signal spectra at multiple frequency offsets and this was particularly problematic as it effectively raised the sensor noise levels. This problem was rectified prior to further flight testing.
- While the optical reference-vector sensor system performed well, the ability to verify and validate the accuracy of the attitude estimation system in flight was hindered by the available position measurement accuracy of the GPS. Error analysis of the comparison between the measured and estimated reference vectors indicated that the available position mea-

surement accuracy introduced too much error into the estimated reference vector so that the error in the attitude estimate could not be constrained as tightly as desired.

8.5 General conclusions

This work considered in detail the potential accuracy of sUAS-based wind measurement. The main error contributors are error in measurement/estimation of the relative wind (vector) velocity, the sUAS attitude, the sUAS velocity, and the sUAS attitude rate. For contemporary sUAS-based systems none of these error sources are dominant. Three of these four sources relate to sUAS state-estimation and a careful investigation of contemporary sUAS state-estimation was the main focus of this work. A combination of analytic analyses, simulation and analysis of field data was used to thoroughly investigate contemporary sUAS state estimation, with particular attention given to the effects of wind-gust-induced motion. Based on this investigation it is concluded that well-executed contemporary sUAS state estimation is capable of attitude-estimate accuracy on the order of 0.1° and velocity-estimate accuracy on the order of 0.01 m/s. These accuracy levels support a wind-measurement accuracy level on the order of 0.01 m/s.

8.6 Future work

The work presented herein leaves open several areas of inquiry which may be further pursued.

- Inertial MEMS sensor, small magnetometer and GNSS receiver technology is rapidly evolving due to inclusion of these devices in high-volume consumer electronics. New devices may offer significant performance gains in sUAS state estimation but the sUAS operating environment may be quite different from that for which these devices were intended. Monitoring development and release of these devices, as well as analyzing their suitability to application in sUAS, will allow new sUAS state-estimation systems to make the best use of this evolving technology.

- This work did not examine the question of how the fine-scale motion of the relative-wind sensor relates to the fine-scale motion of the state-estimation sensors. This question is closely tied to particular implementations of sUAS-based wind-measurement systems. It is expected that there will be differences between these fine-scale motions in systems which do not have sufficient rigidity between the relative-wind sensor and the state-estimation sensors. For systems making high-bandwidth wind measurements this difference in fine-scale motion may be of great significance and should be further investigated.
- The optical reference-vector sensor presented herein shows promise for use in a variety of applications. Further development of this sensor system may adapt it to a wide variety of uses as described in section 6.3.
- Sensor calibration methods for the work performed herein relied on calibration performed on the ground with the airframe propulsion system turned off. A more representative calibration, particularly for the magnetometer, might be performed using data collected while in flight. Further work may develop appropriate methods and algorithms to accomplish in-flight sensor calibration.

Bibliography

- [1] Piksi Datasheet, version 2.3.1. Technical report, Swift Navigation, San Francisco, CA, 2013.
- [2] Martin Alkeryd. Evaluation of Position Sensing Techniques for an Unmanned Aerial Vehicle. PhD thesis, Linköpings universitet, 2006.
- [3] T. Aruga and T. Igarashi. Three-Axes Attitude Determination of Spacecraft Using a Laser. IEEE Transactions on Aerospace and Electronic Systems, AES-13(5):473–479, September 1977.
- [4] R. Banta, Y. Pichugina, and A. Brewer. Turbulent Velocity-Variance Profiles in the Stable Boundary Layer Generated by a Nocturnal Low-Level Jet. Journal of the Atmospheric Sciences, 63:2700–2719, 2006.
- [5] Yaakov Bar-Shalom, X. Rong Li, and Thiagalingam Kirubarajan. Estimation with Applications To Tracking and Navigation. John Wiley and Sons, Inc., 2001.
- [6] Jeffrey D Barton. Fundamentals of Small Unmanned Aircraft Flight. Johns Hopkins APL Technical Digest, 31(2):132–149, 2012.
- [7] R. Beard and T.W. McLain. Small Unmanned Aircraft Theory and Practice - Appendix E. Princeton University Press, 2012.
- [8] Paul Beckmann and Timothy Stilson. AN EFFICIENT ASYNCHRONOUS SAMPLING-RATE CONVERSION ALGORITHM FOR MULTI-CHANNEL AUDIO APPLICATIONS. In Audio Engineering Society 119th Convention, New York, New York, USA, 2005.
- [9] G. Buzorius, Ü. Rannik, J.M. Mäkelä, T. Vesala, and M. Kulmala. Vertical aerosol particle fluxes measured by eddy covariance technique using condensational particle counter. Journal of Aerosol Science, 29(1):157–171, 1998.
- [10] Campbell Scientific. CSAT3 Three Dimensional Sonic Anemometer Instruction Manual. Technical report, Campbell Scientific, Logan, UT, 2012.
- [11] N.A. Carlson. Federated filter for computer-efficient, near-optimal GPS integration. In Proceedings of Position, Location and Navigation Symposium - PLANS '96, pages 306–314. IEEE, 1996.

- [12] Andreas Christen, Eva van Gorsel, and Roland Vogt. Coherent structures in urban roughness sublayer turbulence. International Journal of Climatology, 27(14):1955–1968, November 2007.
- [13] D. Comotti. Orientation estimation based on gauss-newton method and implementation of a quaternion complementary filter. Technical report, Università degli studi di Bergamo, 2011.
- [14] Martin Cope, Dale Hess, Sunhee Lee, and Et. All. Traffic and Meteorological Impacts on Near-Road Air Quality: Summary of Methods and Trends from the Raleigh Near-Road Study. Journal of the Air & Waste Management Association, 58(7):865–878, July 2008.
- [15] John L. Crassidis, F. Landis Markley, and Yang Cheng. Survey of Nonlinear Attitude Estimation Methods. AIAA Journal of Guidance, Control, and Dynamics, 30(1):12–28, May 2012.
- [16] Philippe Drobinski, Pierre Carlotti, Jean-Luc Redelsperger, Valery Masson, Robert M. Banta, and Rob K. Newsom. Numerical and Experimental Investigation of the Neutral Atmospheric Surface Layer. Journal of the Atmospheric Sciences, 64(1):137–156, January 2007.
- [17] Jack Elston, Brian Argrow, Maciej Stachura, Doug Weibel, Dale Lawrence, and David Pope. Overview and Comparison of Wind Estimation Using Small Fixed-Wing Unmanned Aircraft. Journal of Atmospheric and Oceanic Technology, (Accepted for publication), 2014.
- [18] Scott M. Ettinger, Michael C. Nechyba, Peter G. Ifju, and Martin Waszak. Vision-guided flight stability and control for micro air vehicles. Advanced Robotics, 17(7):617–640, January 2003.
- [19] Mark Euston, Paul Coote, and Robert Mahony. A Complementary Filter for Attitude Estimation of a Fixed-Wing UAV. Intelligent Robots and, 2008.
- [20] T. Fiorenzani, Costanzo Manes, G. Oriolo, and P. Peliti. Comparative Study of Unscented Kalman Filter and Extended Kalman Filter for Position/Attitude Estimation in Unmanned Aerial Vehicles. Research Report Series of IASI-CNR, Rome, Italy (ISSN: 1128-3378), 08-08, 2008.
- [21] R. Fletcher and M. J. D. Powell. A Rapidly Convergent Descent Method for Minimization. The Computer Journal, 6(2):163–168, August 1963.
- [22] R.C. Foster, F. Vianey, P. Drobinski, and P. Carlotti. Near-surface coherent structures and the vertical momentum flux in a large-eddy simulation of the neutrally-stratified boundary layer. Boundary-layer meteorology, 2006.
- [23] Demoz Gebre-Egziabher, Gabriel H. Elkaim, J. David Powell, and Bradford W. Parkinson. Calibration of Strapdown Magnetometers in Magnetic Field Domain. Journal of Aerospace Engineering, 19(2):87–102, February 2006.
- [24] R. Gold. Optimal binary sequences for spread spectrum multiplexing (Corresp.). IEEE Transactions on Information Theory, 13(4):619–621, October 1967.
- [25] Tronics Group. An Overview of MEMS and non-MEMS High Performance Gyros. Technical report, tronics Group, 2013.

- [26] Hamamatsu Photonics K.K. PSD Technical Information. Technical report, Hamamatsu Photonics K.K., Hamamatsu City, Japan, 2013.
- [27] William Rowan Hamilton. II. On quaternions; or on a new system of imaginaries in algebra. Philosophical Magazine Series 3, 25(163):10–13, July 1844.
- [28] M. Maureen Hand and Mark J. Balas. Blade load mitigation control design for a wind turbine operating in the path of vortices. Wind Energy, 10(4):339–355, July 2007.
- [29] Y. Harazono, J. Kim, A. Miyata, T. Choi, J.-I. Yun, and J.-W. Kim. Measurement of energy budget components during the International Rice Experiment (IREX) in Japan. Hydrological Processes, 12(13-14):2081–2092, October 1998.
- [30] Carl Hartung, Richard Han, Carl Seielstad, and Saxon Holbrook. FireWxNet. In Proceedings of the 4th international conference on Mobile systems, applications and services - MobiSys 2006, page 28, New York, New York, USA, June 2006. ACM Press.
- [31] Isaac M. Held and Arthur Y. Hou. Nonlinear Axially Symmetric Circulations in a Nearly Inviscid Atmosphere. Journal of the Atmospheric Sciences, 37(3):515–533, March 1980.
- [32] M.L. Hill. INTRODUCING ELECTROSTATIC AUTOPILOT. Astronautics & Aeronautics, 1972.
- [33] R.J. Hill. Corrections to Taylor’s frozen turbulence approximation. Atmospheric Research, 40(2-4):153–175, May 1996.
- [34] J. Horel, T. Potter, L. Dunn, W. J. Steenburgh, M. Eubank, M. Splitt, and D. J. Onton. Weather Support for the 2002 Winter Olympic and Paralympic Games. Bulletin of the American Meteorological Society, 83(2):227–240, February 2002.
- [35] Young Meteorological Instruments. Instructions: Ultrasonic Anemometer Model 81000. Technical report, Young Meteorological Instruments, Traverse City, MI, 2014.
- [36] Arno Solin Jouni Hartikainen. Optimal filtering with Kalman filters and smoothers - a Manual for Matlab toolbox EKF/UKF. Technical report, Aalto University School of Science, Aalto, Finland, 2011.
- [37] S.J. Julier and J.K. Uhlmann. Unscented Filtering and Nonlinear Estimation. Proceedings of the IEEE, 92(3):401–422, March 2004.
- [38] D. Jung and P. Tsiotras. Inertial Attitude and Position Reference System Development for a small UAV. AIAA Infotech at aerospace, 2007.
- [39] G. Kocer, M. Mansour, N. Chokani, R.S. Abhari, and M. Muller. Full-Scale Wind Turbine Near-Wake Measurements Using an Instrumented Uninhabited Aerial Vehicle. Journal of Solar Energy Engineering, 133(4):041011, 2011.
- [40] J.B. Kosmatka. Development of a long-range small UAV for atmospheric monitoring. AIAA Paper 2234, 2007.
- [41] Jack. B Kuipers. Quaternions and Rotation Sequences: A Primer With Applications... Princeton University Press, 2002.

- [42] D. Lawrence and B. Balsley. High-resolution atmospheric sensing using a low-cost miniature autonomous arial vehicle. Journal of Atmospheric and Oceanic Technology, 2011.
- [43] Q.S. Li, J.R. Wu, S.G. Liang, Y.Q. Xiao, and C.K. Wong. Full-scale measurements and numerical evaluation of wind-induced vibration of a 63-story reinforced concrete tall building. Engineering Structures, 26(12):1779–1794, 2004.
- [44] S.O.H. Madgwick. An Efficient Orientation Filter for Inertial and Inertial/Magnetic Sensor Arrays. Report x-io and University of Bristol (UK), 2010.
- [45] F. Landis Markley. Attitude Determination Using Vector Observations: A Fast Optimal Matrix Algorithm. The Journal of the Astronautical Sciences, 41(2):261–280, 1993.
- [46] Sabrina Martin, Frank Beyrich, and Jens Bange. Observing Entrainment Processes Using a Small Unmanned Aerial Vehicle: A Feasibility Study. Boundary-Layer Meteorology, December 2013.
- [47] Matthew Fonatine Maury. Wind and Current Chart of the North Atlantic. Technical report, U.S. Naval Observatory, 1852.
- [48] Matthew Fontaine Maury. Physical Geography Of The Sea. Harper, 1855.
- [49] Metek GmbH. Ultrasonic Wind Sensor uSonic-3 Class A Datasheet. Technical report, Metek GmbH, Elmshorn, Germany, 2013.
- [50] T. Mikkelsen. Remote Sensing of Wind. Technical report, Technical University of Denmark, 2012.
- [51] Pratap Misra and Per Enge. Global Positioning System: Signals, Measurements and Performance. Ganga-Jamuna Press, 2001.
- [52] Cleve Moler and Charles Van Loan. Nineteen Dubious Ways to Compute the Exponential of a Matrix, Twenty-Five Years Later. SIAM Review, 45(1):3, 2003.
- [53] L. Orf, E. Kantor, and E. Savory. Simulation of a downburst-producing thunderstorm using a very high-resolution three-dimensional cloud model. Journal of Wind Engineering and Industrial . . ., 2012.
- [54] C.S. Osborn. Gyroscope Angular Rate Processing Achross Asynchronous Clock Domains. PhD thesis, Massachusetts Institute of Technology, 2002.
- [55] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. Numerical Recipes in C. University Press, Cambridge University, 1992.
- [56] J. M. Prospathopoulos, E. S. Politis, K. G. Rados, and P. K. Chaviaropoulos. Evaluation of the effects of turbulence model enhancements on wind turbine wake predictions. Wind Energy, 14(2):285–300, March 2011.
- [57] J. Quindlen and J. Langelaan. Flush Air Data Sensing for Soaring-Capable UAVs. In 51st AIAA Aeospace Sciences Meeting, 2013.
- [58] L.R. Rabiner and R.E. Crochiere. Interpolation and decimation of digital signalsA tutorial review. Proc. IEEE, 69(3):300–331, 1981.

- [59] H. E. RAUCH, C. T. STRIEBEL, and F. TUNG. Maximum likelihood estimates of linear dynamic systems. AIAA Journal, 3(8):1445–1450, August 1965.
- [60] Valérie Renaudin, Muhammad Haris Afzal, and Gérard Lachapelle. Complete Triaxis Magnetometer Calibration in the Magnetic Domain. Journal of Sensors, 2010:1–10, 2010.
- [61] Mathias W Rotach. On the influence of the urban roughness sublayer on turbulence and dispersion. Atmospheric Environment, 33(24):4001–4008, 1999.
- [62] J.R. Rowland. Intensive probing of clear air convective fields by radar and instrumented drone aircraft. In Radar Meteorology Conference, pages 321–326, Boston, January 1972. NASA Technical Report, American Meteorological Society.
- [63] S. Salcudean. US Patent US5059789 A: Optical position and orientation sensor, October 1991.
- [64] Carlos Santos, Felipe Espinosa, and Héctor García de Marina. Adaptive UAV Attitude Estimation Employing Unscented Kalman Filter, FOAM and Low-Cost MEMS Sensors. Sensors, 12(7):9566–9585, 2012.
- [65] Paul G Savage. Strapdown Analytics, Volume 1. Strapdown Associates, Incorporated, 2000.
- [66] G. Sepp. Earth laser beacon sensor for earth-oriented geosynchronous satellites. Applied optics, 14(7):1719–26, July 1975.
- [67] SgurrEnergy Ltd. Galion LIDAR Product Brochure. Technical report, SgurrEnergy Ltd., Glasgow, United Kingdom, 2013.
- [68] Claude E. Shannon. Communication in the presence of noise. Proceedings of the IEEE, 86(2):447–457, 1998.
- [69] Julius O. Smith. Digital Audio Resampling Home Page, 2014.
- [70] Military Specification. MIL-F-8785 (ASG) Flying Qualities of Piloted Airplanes, 1. Technical report, 1954.
- [71] David J. Stensrud. Parameterization Schemes: Keys to Understanding Numerical Weather... Cambridge University Press, 2007.
- [72] John Stimac. Class Notes: Weather an Climate, 2003.
- [73] Byron D. Tapley, Bob E. Schutz, and George H. Born. Statistical Orbit Determination. Elsevier Academic Press, Burlington, MA, 2004.
- [74] J. Taylor. Classical Mechanics. University Science Books, 2005.
- [75] Thies Clima. Thies Clima Ultrasonic Anemometer 3D Datasheet. Technical report, Thies Clima, Gottingen, Germany, 2014.
- [76] M.R. Thorpe, E.G. Banke, and S.D. Smith. Eddy correlation measurements of evaporation and sensible heat flux over Arctic sea ice. Journal of Geophysical Research, 78(18):3573–3584, 1973.
- [77] D. Titterton. Strapdown inertial navigation technology. 2004.

- [78] K. Träumner, Ch. Kottmeier, U. Corsmeier, and A. Wieser. Convective Boundary-Layer Entrainment: Short Review and Progress using Doppler Lidar. Boundary-Layer Meteorology, 141(3):369–391, September 2011.
- [79] U-blox. LEA-6 u-blox 6 GPS Modules. Technical report, U-blox, Thalwil, Switzerland, 2013.
- [80] U-blox. NEO-7P U-blox 7 Precise Point Positioning GNSS module Data Sheet. Technical report, U-blox, Thalwil, Switzerland, 2014.
- [81] Aline van den Kroonenberg, Tim Martin, Marco Buschmann, Jens Bange, and Peter Vörsmann. Measuring the Wind Vector Using the Autonomous Mini Aerial Vehicle M 2 AV. Journal of Atmospheric and Oceanic Technology, 25(11):1969–1982, November 2008.
- [82] G. Wahba. A least squares estimate of satellite attitude. SIAM review, 1965.
- [83] L.K. Wang, S.-C. Hsieh, and E.C.-W. Hsueh. Complete pose determination for low altitude unmanned aerial vehicle using stereo vision. In 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 108–113. IEEE, 2005.
- [84] V. Wulfmeyer, M. Randall, A. Brewer, and R. M. Hardesty. 2-m Doppler lidar transmitter with high frequency stability and low chirp. Optics Letters, 25(17):1228, September 2000.
- [85] Tae Suk Yoo, Sung Kyung Hong, Hyok Min Yoon, and Sungsu Park. Gain-Scheduled Complementary Filter Design for a MEMS Based Attitude and Heading Reference System. Sensors (Basel, Switzerland), 11(4):3816–30, January 2011.

Appendices

Appendix 1

Attitude representation and attitude integration using quaternions

Attitude estimation requires a method for representing the rotation of the body frame relative to the inertial frame. There are several common mathematical representations used for this purpose of which the most common are rotation matrices, quaternions, and Euler angles. Quaternion attitude representations are the least intuitive of these, and some readers may be unfamiliar with quaternion algebra. There are a large number of texts dealing with these subjects, for example Kuipers [41]. This appendix will provide some of the fundamental results of quaternion algebra and quaternion rotation representations applicable to attitude estimation.

A quaternion, \mathbf{q} , is defined as

$$\mathbf{q} = q_0 + \mathbf{i}q_1 + \mathbf{j}q_2 + \mathbf{k}q_3 \quad (1.1)$$

where \mathbf{i} , \mathbf{j} and \mathbf{k} are hyper-complex numbers satisfying ‘Hamilton’s’ rule [27]

$$\mathbf{i}^2 = \mathbf{j}^2 = \mathbf{k}^2 = \mathbf{ijk} = -1 \quad (1.2)$$

Note that quaternions do not satisfy field properties if either a dot product or cross product is used

to define multiplication. Multiplication of quaternions is defined using Hamilton's rule as

$$\begin{aligned}
\mathbf{p} \otimes \mathbf{q} &= (p_0 + \mathbf{i}p_1 + \mathbf{j}p_2 + \mathbf{k}p_3)(q_0 + \mathbf{i}q_1 + \mathbf{j}q_2 + \mathbf{k}q_3) \\
&= p_0q_0 + \mathbf{i}p_1q_0 + \mathbf{j}p_2q_0 + \mathbf{k}p_3q_0 \\
&\quad + \mathbf{i}p_0q_1 + \mathbf{i}^2p_1q_1 + \mathbf{i}\mathbf{j}p_2q_1 + \mathbf{i}\mathbf{k}p_3q_1 \\
&\quad + \mathbf{j}p_0q_2 + \mathbf{j}\mathbf{i}p_1q_2 + \mathbf{j}^2p_2q_2 + \mathbf{j}\mathbf{k}p_3q_2 \\
&\quad + \mathbf{k}p_0q_3 + \mathbf{k}\mathbf{i}p_1q_3 + \mathbf{k}\mathbf{j}p_2q_3 + \mathbf{k}^2p_3q_3 \\
&= p_0q_0 - p_1q_1 - p_2q_2 - p_3q_3 \\
&\quad + \mathbf{i}(p_0q_1 + p_1q_0 + p_2q_3 - p_3q_2) \\
&\quad + \mathbf{j}(p_0q_2 - p_1q_3 + p_2q_0 + p_3q_1) \\
&\quad + \mathbf{k}(p_0q_3 + p_1q_2 - p_2q_1 + p_3q_0)
\end{aligned} \tag{1.3}$$

The complex conjugate of a quaternion is defined as

$$\mathbf{q}^* = q_0 - \mathbf{i}q_1 - \mathbf{j}q_2 - \mathbf{k}q_3 \tag{1.4}$$

and the norm of a quaternion is defined as

$$N(\mathbf{q}) = \sqrt{q_0^2 + q_1^2 + q_2^2 + q_3^2} \tag{1.5}$$

A rotation, for example the rotation of the body frame relative to the inertial frame, may be represented by a unit norm quaternion. If frame \mathbf{B} is rotated an angle α around a vector \vec{u} which is parameterized in frame \mathbf{I} as $[u_x \ u_y \ u_z]^T$, then the quaternion describing this rotation is given by

$$\frac{B}{I}\mathbf{q} = \cos \frac{\alpha}{2} + \mathbf{i} \sin \frac{\alpha}{2} (u_x \mathbf{i} + u_y \mathbf{j} + u_z \mathbf{k}) \tag{1.6}$$

A quaternion is represented by the 4-tuple $[q_0 \ q_1 \ q_2 \ q_3]$. Note, however, that some authors use the 4-tuple $[q_1 \ q_2 \ q_3 \ q_0]$. This difference will cause a corresponding change when quaternion equations are written in matrix formats.

Given the Euler angles for the standard aerospace Euler angle sequence ($\psi = \text{Yaw}$, $\theta = \text{Pitch}$, $\phi = \text{Roll}$) ${}^B_I \mathbf{q}$ may be calculated with the equations

$$q_0 = \cos \frac{\psi}{2} \cos \frac{\theta}{2} \cos \frac{\phi}{2} + \sin \frac{\psi}{2} \sin \frac{\theta}{2} \sin \frac{\phi}{2} \quad (1.7)$$

$$q_1 = \cos \frac{\psi}{2} \cos \frac{\theta}{2} \sin \frac{\phi}{2} - \sin \frac{\psi}{2} \sin \frac{\theta}{2} \cos \frac{\phi}{2} \quad (1.8)$$

$$q_2 = \cos \frac{\psi}{2} \sin \frac{\theta}{2} \cos \frac{\phi}{2} + \sin \frac{\psi}{2} \cos \frac{\theta}{2} \sin \frac{\phi}{2} \quad (1.9)$$

$$q_3 = \sin \frac{\psi}{2} \cos \frac{\theta}{2} \cos \frac{\phi}{2} - \cos \frac{\psi}{2} \sin \frac{\theta}{2} \sin \frac{\phi}{2} \quad (1.10)$$

Formulas and algorithms similarly exist to calculate Euler angles from a quaternion rotation representations, a quaternion rotation representation from a rotation matrix and a rotation matrix from a quaternion rotation representation. The quaternion rotation operation, associated with the quaternion \mathbf{q} and applied to a parameterized vector $[\vec{v}] \in \mathbf{R}^3$, allows the vector's parameterization to be changed from one frame to another rotated frame, as

$$[\vec{v}]_I = {}^B_I \mathbf{q} \otimes [\vec{v}]_B \otimes {}^B_I \mathbf{q}^*. \quad (1.11)$$

The use of the quaternion product operator with a vector operand in this equation has the specific meaning that the vector is treated as a quaternion parameterized as $[0 \ v_x \ v_y \ v_z]$ and where v_x , v_y and v_z are the components of \vec{v} in the B frame.

If the quaternion, ${}^B_I \mathbf{q}$, represents the rotation of the body frame relative to the inertial frame (i.e. the attitude of the body frame) and if the body frame is rotating with rate and direction defined by the rotation vector $[{}^I \vec{\omega}_B]_B$ then the time derivative of ${}^B_I \mathbf{q}$ is given by

$${}^B_I \dot{\mathbf{q}} = \frac{1}{2} {}^B_I \mathbf{q} \otimes [{}^I \vec{\omega}_B]_B \quad (1.12)$$

Here again, $[{}^I \vec{\omega}_B]_B$ is treated as a quaternion as described above. Integration of this equation yields the rotation resulting from a time history of rotation rate.

Appendix 2

Block diagonal structure of the EKF S matrix

Differences between filters may be analyzed using generalized matrices which include elements relating to all states and observations. If state mean and covariance estimate updates using measurements from different sensors are processed separately, then the H , R , S and K matrices used for a particular update are not the “full” generalized matrices. For example an attitude filter may have a generalized S matrix:

$$S \approx \begin{bmatrix} S_{sf} & 0 & 0 & 0 \\ 0 & S_b & 0 & 0 \\ 0 & 0 & S_v & 0 \\ 0 & 0 & 0 & S_p \end{bmatrix}, \quad (2.1)$$

where the sf , b , v and p subscripts denote the blocks related to the accelerometer measurement, magnetometer measurement, velocity measurement and position measurement, respectively. However, when processing an observation from a single sensor a set of matrices representative of the dimension of that single sensor are used. This forces, by construction, the “full” generalized S matrix to have zero blocks, with all blocks relating to any pair of sensors not processed in the same update being zero. If all sensor measurements from different sensors are processed separately then the “full” S matrix is strictly block diagonal as shown above .

GNSS measurements are generally available at a lower rate than gyroscope/accelerometer measurements and, consequently, are typically processed separately. In this case then the upper-

right and lower-left 2x2 blocks of the generalized S matrix in the block form shown above will be all zero.

With respect to analysis of the difference between attitude estimates produced with federated and cascaded filters the (1,2) and (2,1) blocks in the generalized S matrix are also of interest; additional terms introduced if the (3,4) and (4,3) blocks are non-zero will appear equally in both filters correction step equations, and will not contribute any difference between methods. Now, if the gyroscope and accelerometer measurement error are uncorrelated, then it can be shown that the covariance estimate will be the same after applying an update using both sensors' data (together) or after sequentially applying an update using each sensor's data individually. If the updates are processed individually, then, by construction, the equivalent upper-left 2x2 block of the S matrix is block diagonal, as discussed above. If the update is processed as a single update with both sensors' data, however, the covariance estimate will be the same; by implication, the upper-left 2x2 block of the S matrix is equal in both cases and block diagonal.

Appendix 3

The extended Kalman filter algorithm

The extended Kalman filter (EKF) is a well known statistical filter related to the Kalman filter but using linearization of the plant and observation sensor models.

Implementation of an EKF requires a discrete time plant and observation model. Discrete time models may be produced from a continuous time model with a simple first-order approximations, or more precise methods such as a Van Loan approximation. Propagation of a discrete time state is represented as

$$x_k = f(x_{k-1}) + w_{k-1}$$

where w is a process noise term with covariance

$$Q_k = E[w_k w_k^T].$$

Observation sensor measurements are mapped from the state by the measurement model

$$y_k = h(x_k) + v_k$$

where v is a measurement noise term with covariance

$$R_k = E[v_k v_k^T].$$

The EKF operates in a cyclic fashion with a prediction step propagating mean and covariance estimates of the state based on the plant model and process noise and a correction step updating

the state mean and covariance estimates based on the observation sensor measurements. These steps are summarized below.

Prediction Step:

- $\hat{x}_{k|k-1} = f(\hat{x}_{k-1})$
- $P_{k|k-1} = F(\hat{x}_{k-1})P_{k-1}F(\hat{x}_{k-1})^T + Q_{k-1}$
- F is the jacobian of f with respect to x .

Correction Step:

- $\hat{x}_k = \hat{x}_{k|k-1} + K_k[y_k - h(\hat{x}_{k|k-1})]$
- $P_k = [I - K_k H(\hat{x}_{k|k-1})]P_{k|k-1}$
- $K_k = P_{k|k-1} H(\hat{x}_{k|k-1}) S_k^{-1}$
- $S_k = H(\hat{x}_{k|k-1}) P_{k|k-1} H(\hat{x}_{k|k-1})^T + R_k$
- H is the jacobian of h with respect to x .

Note that when S_k and $P_{k|k-1}$ are symmetric (typical) then

- $P_k = P_{k|k-1} - K_k S_k K_k^T$.

Appendix 4

Simulation of sUAS state estimation

Simulation of sUAS state estimation and wind measurement

4.1 Generation of truth data based on an aerodynamic model

The Simulink graphical programming environment, including the aerospace block-set, is used to simulate a sUAS in flight and generate ‘truth data’ for use in further simulation of state estimation filters. A forces-and-moments model of the sUAS using stability and control derivatives produces calculated aerodynamic forces and moments at each simulation time step based on the current sUAS state and the wind model. Wind is modeled using a pre-configured Dryden model block in the aerospace block-set. Control surface actuation is based on an autopilot model with altitude and bank PID control loops. The calculated aerodynamic forces and combined with thrust and gravity and the state is propagated using a 6-degree-of-freedom motion integrator using mass and inertia properties of the airframe. A Matlab script is used to manage the simulation and collect results.

Using these software components data-sets may be produced representing the full state of the sUAS during flight in wind. The aircraft characteristics may be changed through use of different sets of stability and control derivatives and mass and inertial properties. The wind environment may be changed by specifying different mean wind speeds and turbulence levels. The flight condition may be changed by specifying a different autopilot control condition to produce straight and level flight

or circling flight, etc. The exact behavior of the autopilot is not important as the goal is to produce a dataset representative of an sUAS flying in wind with some general behavior - this dataset that includes realistic motions, particularly those due to wind gusts and circling flight paths, can then be used as a truth baseline for further simulation.

The simulation results presented herein were based on simulated flight of a Zagi flying wing airframe. Stability and control derivatives were taken from Beard and McLain [7].

4.2 Simulation of sensor measurements based on truth data

Using sUAS-state data-sets generated as described above a Matlab script is used to generate simulated sensor outputs. Simulation of the sensors includes error models as presented in Equations (4.1), (4.2) and (4.3) as well as quantization error and dynamics models for the gyroscope and accelerometer. Parameters used to generate simulated sensor outputs (with 'standard' error levels) include:

(1) Gyroscope

- Sampling interval 0.01 seconds.
- Zero scale factor, misalignment and acceleration sensitivity errors.
- Natural frequency 10 Hz. Damping ratio 0.707.
- Zero-point drift - sinusoidal with amplitude 1.7 rad/s^2 and frequency 0.01 Hz.
- Noise - angular random walk $1 \text{ deg}/\sqrt{\text{hour}}$.
- 12 bit quantization over $\pm 5 \text{ rad/s}$ range.

(2) Accelerometer

- Sampling interval 0.01 seconds.
- Zero scale factor and misalignment errors.

- Natural frequency 10 Hz. Damping ratio 0.707.
- No zero-point drift.
- Noise - 0.1 g standard deviation.
- 12 bit quantization over ± 5 g range.

(3) Magnetometer

- Sampling interval 0.01 seconds.
- Zero scale factor and misalignment errors.
- Fixed zero-point error from a random Gaussian sample with standard deviation 0.005 Gauss.
- Noise - 0.003 Gauss standard deviation.

(4) GPS

- Sampling interval 0.2 seconds.
- Latency 0.1 seconds.
- Velocity error standard deviation - 0.05 m/s horizontal, 0.1 m/s vertical.
- Position error standard deviation - 0.5 m horizontal, 1 m vertical.


4.3 Simulation of competing sUAS state estimation algorithms

Performance of different sUAS state estimation algorithms was compared through simulation in Matlab using the sensor data-sets described in section 4.2. The estimated sUAS state was compared with the ‘truth’ data-set (section 4.1). Simulation of all state estimation algorithms is straightforward and follows from Equations (4.7), (4.8). and (4.9). Cascaded AHRS/position-velocity algorithms are simulated using two cascaded stages. Equations (4.13) and (4.15) are used to calculate feedback gains for the fixed-step-size gradient descent filter and modified-gradient-descent AHRS filters, respectively while the standard EKF equations presented in Appendix 3 are

used for all EKF filters. Translational acceleration estimates are calculated with either Barton's approximation of Equation (4.31) or with Equation (4.34). All required jacobians may be calculated from Equations (4.5) and (4.6) using Matlab symbolic toolbox.

Appendix 5

Optical reference vector sensor schematics, board layouts, and firmware



5.1 Firmware

The optical reference vector sensor firmware runs on a STMicroelectronics STM32F4DISCOVERY development board. The firmware source code is written in C and was developed in parallel with the CU M4 autopilot code. The device driver modules, timer modules, etc. are generally the same as those used in the CU M4 autopilot code and are not reproduced here. Only the main program file, main.c, is included here which documents the basic functionality of the ORVS functionality.

5.1.1 main.c file listing

```
/**
*****
* @file    ORVSTest1/src/main.c
* @author  D Weibel
* @version V1.0.0
* @date    June 2014
* @brief   Main program body
*****
*
* This is a first test program for the ORVS board.
*   - Collects samples from channel pairs - rate limited by the spi bus
*   - Samples collected for 2 msec period for each pair
*   - Measurement rate regulated to 50Hz
*   - Outputs data to the debug usart
*
* Connections to ORSV board
*   - 5V, 3V and Gnd
*   - PA5=SCK, PA6=MISO, PA7=MOSI
*
*****
*/

// Defined constants
#define NUMSAMPLES 300          // Set above the maximum expected
#define STARTDELAY 1*10000     //180*10000

/* Includes -----*/
#include "main.h"
#include "init.h"
#include "usart.h"
#include "spi.h"
#include <stdio.h>
#include <string.h>
#include <math.h>
```

```

/* Private variables -----*/
/*
 * FAT file system variables
 */
FIL fileID;
FRESULT fileResult;
FATFS FatFs;

// SysTick
//static __IO uint32_t TimingDelay;

/* Private function prototypes -----*/

int main(void)
{
    uint8_t    i=12;
    uint8_t    j=12;
    uint8_t    k;
    uint8_t    msb, lsb;
    uint8_t    channel;
    uint8_t    sample=0;
    uint8_t    stageX, stageY;
    uint8_t    trap;
    uint16_t   nsampleX, nsampleY;
    uint16_t   min1, max1, min2, max2;
    uint16_t   data[4][NUMSAMPLES]; // data[channel, sample]
    float      avg;
    float      sampleAmpl;
    float      channelAmpl[8];
    float      X, Y, sumX, sumY, satX, satY;
    float      Xmeas, Ymeas;
    float      Xstep, Xfrac, Xdiff, Xlow, Xhigh;
    float      Ystep, Yfrac, Ydiff, Ylow, Yhigh;
    char       out[80];
#ifdef SDLogging
    char       dataMsg[10+4*NUMSAMPLES];
#endif

    uint32_t   startTick, loopTick;

    // SysTick configured in Delays.c to 100us and set to highest priority

    //Delay_10usec(10000); // Delay 1 sec

    // Initialize the hardware
    Initialize();

    USART_Puts(DEBUG_USART, (char *)"\nORVS Test 1 startup\n");

    // Initialize the file system and SD card
#ifdef SDLogging

```

```

fileResult = f_mount(&FatFs, "", 1);
if( fileResult != FR_OK ) while(1) {USART_Puts(DEBUG_USART, (char *)"\nf_mount fail\n");};
fileResult = f_open(&fileID, "ORVSt1.BIN", FA_CREATE_ALWAYS | FA_READ | FA_WRITE);
if( fileResult != FR_OK ) while(1) {USART_Puts(DEBUG_USART, (char *)"\nf_open fail\n");};
f_sync(&fileID); // to prevent corruption
#endif

LEDOn(LED_RED);
while (SystemTickCounter < STARTDELAY)
{
}
LEDOff(LED_RED);
LEDOn(LED_GREEN);
LEDOn(LED_BLUE);

/////////////////////////////////////////////////////////////////
// This is the top of the main measurement loop. Loop time is regulated (at the
// bottom of the loop) by loopTick
// The main loop ends when the user pushbutton is pressed.
/////////////////////////////////////////////////////////////////
while (GPIO_ReadInputDataBit(GPIOA, GPIO_Pin_0) == 0)
//while (SystemTickCounter < DATAPERIOD)
{
    loopTick = SystemTickCounter;

    ///////////////////////////////////////////////////////////////////
    // Here we collect samples for X channel pairs. The expected input waveform is at 1Kz
    // and we collect a two period segment

    // Channel 1/2 is sampled. If either are saturated then channel 5/6 is sampled
    ///////////////////////////////////////////////////////////////////

    sample=0;
    max1 = max2 = 0;
    min1 = min2 = 65535;

    // Wait for a SystemTick change to improve timing accuracy
    startTick = SystemTickCounter;
    while(SystemTickCounter == startTick) {}

    // Loop to collect channel 1, 2 samples for 2 milliseconds
    while(SystemTickCounter<startTick+21)
    {

        // Get channel 1
        SPI1_csLow();
        msb = SPI1_send((uint8_t) 0x00); // Start conversion
        msb = SPI1_send((uint8_t) 0); // Get msb
        lsb = SPI1_send((uint8_t) 0); // Get lsb
        SPI1_csHigh();
        data[0][sample] = (((uint16_t)msb)<<8) + (uint16_t)lsb;
        if (data[0][sample] > max1) max1 = data[0][sample];
    }
}

```



```

    SPI1_csHigh();
    data[1][sample] = (((uint16_t)msb)<<8) + (uint16_t)lsb;
    if (data[1][sample] > max2) max2 = data[1][sample];
    if (data[1][sample] < min2) min2 = data[1][sample];

    sample++;
}
nSampleX = sample;

if (min1 < 2 || min2 < 2 || max1 > 65533 || max2 > 65533 || sample<1)
{
    stageX = 0;        // Both stages are saturated :(
    X = 999;          // Set output to 999 indicating instrument is saturated
}
}

if(stageX > 0)
{
    // Either the 2nd or 3rd stage is valid so calculate the X measurement
    for (channel=0;channel<2;channel++)
    {
        avg = 0.0;
        for (sample=0; sample<nSampleX; sample++)
        {
            avg += (float)data[channel][sample];
        }
        avg = avg/(float)nSampleX;
        channelAmpl[channel] = 0.0;
        for (sample=0; sample<nSampleX; sample++)
        {
            sampleAmpl = (float)data[channel][sample] - avg;
            channelAmpl[channel] += fabs(sampleAmpl);
        }
    }

    sumX = (channelAmpl[0]+channelAmpl[1]);
    if (sumX==0)
    {
        X = 998;        // Bad data flag (divide by zero)
    } else {
        X = (channelAmpl[0]-channelAmpl[1]) / sumX;
    }
    satX = ((65533-max1)<(65533-max2)?(65533-max1):(65533-max2));
    satX = (satX<(min1-3)?satX:(min1-3));
    satX = (satX<(min2-3)?satX:(min2-3));
    satX = satX*200./65530.;    // Gives percentage of range left before saturation
}

/* Finally create an X message for the SD card
// Message format   Byte(s)           Content
                    0                   "D"
                    1                   "X"

```

```

                2                stageX                // uint8_t
                3-4              samplesX              // uint16_t
                5-8              satX                  // float
                9-12             rawX                  // float
                13 to 4*samplesX+12 data              // uint16_t, 1,2,1,2...
    */
#ifdef SDLogging
    // Set up message header
    dataMsg[0] = 'D';
    dataMsg[1] = 'X';
    dataMsg[2] = stageX;
    memcpy(&dataMsg[3], &nsampleX, 2);
    memcpy(&dataMsg[5], &satX, 4);
    memcpy(&dataMsg[9], &X, 4);

    for (sample=0;sample<nsampleX;sample++)
    {
        memcpy(&dataMsg[4*sample+13], &data[0][sample], 2);
        memcpy(&dataMsg[4*sample+15], &data[1][sample], 2);
    }
    BuffMsg(dataMsg, 4*nsampleX+13);
#endif

    ////////////////////////////////////////////////////
    // Repeat for channels 3/4, 7/8

    // Channel 3/4 is sampled.  If either are saturated then channel 7/8 is sampled
    ////////////////////////////////////////////////////
    sample=0;
    max1 = max2 = 0;
    min1 = min2 = 65535;

    // Wait for a SystemTick change to improve timing accuracy
    startTick = SystemTickCounter;
    while(SystemTickCounter == startTick) {}

    // Loop to collect channel 3, 4 samples for 2 milliseconds
    while(SystemTickCounter<startTick+21)
    {

        // Get channel 3
        SPI1_csLow();
        msb = SPI1_send((uint8_t) 0x40);    // Start conversion
        msb = SPI1_send((uint8_t) 0);      // Get msb
        lsb = SPI1_send((uint8_t) 0);      // Get lsb
        SPI1_csHigh();
        data[2][sample] = (((uint16_t)msb)<<8) + (uint16_t)lsb;
        if (data[2][sample] > max1) max1 = data[2][sample];
        if (data[2][sample] < min1) min1 = data[2][sample];

        // Get channel 4
        SPI1_csLow();

```

```

msb = SPI1_send((uint8_t) 0x60);    // Start conversion
msb = SPI1_send((uint8_t) 0);      // Get msb
lsb = SPI1_send((uint8_t) 0);      // Get lsb
SPI1_csHigh();
data[3][sample] = (((uint16_t)msb)<<8) + (uint16_t)lsb;
if (data[3][sample] > max2) max2 = data[3][sample];
if (data[3][sample] < min2) min2 = data[3][sample];

sample++;
}
nsampleY = sample;
if (min1 < 2 || min2 < 2 || max1 > 65533 || max2 > 65533 || sample<1)
{
    stageY = 2;          // Flag indicating second stage output should be used
} else {
    stageY = 3;          // Flag indicating third stage output should be used
}

////////////////////////////////////
// Data is collected from channels 7/8 only if 3/4 was saturated

if(stageY == 2)
{
    sample=0;
    max1 = max2 = 0;
    min1 = min2 = 65535;

    // Wait for a SystemTick change to improve timing accuracy
    startTick = SystemTickCounter;
    while(SystemTickCounter == startTick) {}

    // Loop to collect channel 7, 8 samples for 2 milliseconds
    while(SystemTickCounter<startTick+21)
    {

        // Get channel 7
        SPI1_csLow();
        msb = SPI1_send((uint8_t) 0xC0);    // Start conversion
        msb = SPI1_send((uint8_t) 0);      // Get msb
        lsb = SPI1_send((uint8_t) 0);      // Get lsb
        SPI1_csHigh();
        data[2][sample] = (((uint16_t)msb)<<8) + (uint16_t)lsb;
        if (data[2][sample] > max1) max1 = data[2][sample];
        if (data[2][sample] < min1) min1 = data[2][sample];

        // Get channel 8
        SPI1_csLow();
        msb = SPI1_send((uint8_t) 0xE0);    // Start conversion
        msb = SPI1_send((uint8_t) 0);      // Get msb
        lsb = SPI1_send((uint8_t) 0);      // Get lsb
        SPI1_csHigh();
        data[3][sample] = (((uint16_t)msb)<<8) + (uint16_t)lsb;
        if (data[3][sample] > max2) max2 = data[3][sample];

```

```

        if (data[3][sample] < min2) min2 = data[3][sample];

        sample++;
    }
    nsampleY = sample;
    if (min1 < 2 || min2 < 2 || max1 > 65533 || max2 > 65533 || sample<1)
    {
        stageY = 0;    // Both stages are saturated :(
        Y = 999;      // Set output to 999 indicating instrument is saturated
    }
}

if(stageY > 0)
{
    // Either the 2nd or 3rd stage is valid so calculate the Y measurement
    for (channel=2;channel<4;channel++)
    {
        avg = 0.0;
        for (sample=0; sample<nsampleY; sample++)
        {
            avg += (float)data[channel][sample];
        }
        avg = avg/(float)nsampleY;
        channelAmpl[channel] = 0.0;
        for (sample=0; sample<nsampleY; sample++)
        {
            sampleAmpl = (float)data[channel][sample] - avg;
            channelAmpl[channel] += fabs(sampleAmpl);
        }
    }

    sumY = (channelAmpl[2]+channelAmpl[3]);
    if (sumY==0)
    {
        Y = 998;
    } else {
        Y = (channelAmpl[2]-channelAmpl[3]) / sumY;
    }
    satY = ((65533-max1)<(65533-max2)?(65533-max1):(65533-max2));
    satY = (satX<(min1-3)?satX:(min1-3));
    satY = (satX<(min2-3)?satX:(min2-3));
    satY = satY*200./65530.;    // Gives percentage of range left before saturation
}

/* Finally create an Y message for the SD card
// Message format same as X message except header "DY"
*/
#ifdef SDLogging
// Set up message header
dataMsg[0] = 'D';
dataMsg[1] = 'Y';
dataMsg[2] = stageY;
memcpy(&dataMsg[3], &nsampleY, 2);

```

```
memcpy(&dataMsg[5], &satY, 4);
memcpy(&dataMsg[9], &Y, 4);

for (sample=0;sample<nsampleX;sample++)
{
    memcpy(&dataMsg[4*sample+13], &data[2][sample], 2);
    memcpy(&dataMsg[4*sample+15], &data[3][sample], 2);
}
BuffMsg(dataMsg, 4*nsampleY+13);
#endif

//////////////////////////////////////
// Calculate calibrated output angles
//////////////////////////////////////
Xmeas = X;
Ymeas = Y;

// If either raw measurement is out of range set the converted angles to the
// designated invalid value 999
if (Xmeas<xmin || Xmeas>xmax || Ymeas<ymin || Ymeas>ymax)
{
    X = 997;
    Y = 997;
} else {
    //First find interpolation set
    trap = 0;
    while (Xmeas > Xcal[i+1][j] || Xmeas < Xcal[i][j] ||
          Ymeas > Ycal[i][j+1] || Ymeas < Ycal[i][j])
    {
        if (Xmeas > Xcal[i+1][j] && i<24) i=i+1;
        if (Xmeas < Xcal[i][j] && i>0) i=i-1;
        if (Ymeas > Ycal[i][j+1] && j<24) j=j+1;
        if (Ymeas < Ycal[i][j] && j>0) j=j-1;
        trap++;
        if (trap > 25) break;
    }

    // Interpolate
    // Step 1 - calculate X values for high and low Y
    Xstep = angles[i+1] - angles[i];
    Xfrac = Xmeas - Xcal[i][j];
    Xdiff = Xcal[i+1][j] - Xcal[i][j];
    Xlow = angles[i] + Xstep * (Xfrac/Xdiff);
    Xfrac = Xmeas - Xcal[i][j+1];
    Xdiff = Xcal[i+1][j+1] - Xcal[i][j+1];
    Xhigh = angles[i] + Xstep * (Xfrac/Xdiff);
    // Step 2 - calculate Y value for high and low X
    Ystep = angles[j+1] - angles[j];
    Yfrac = Ymeas - Ycal[i][j];
    Ydiff = Ycal[i][j+1] - Ycal[i][j];
    Ylow = angles[j] + Ystep * (Yfrac/Ydiff);
    Yfrac = Ymeas - Ycal[i+1][j];
    Ydiff = Ycal[i+1][j+1] - Ycal[i+1][j];
```

```

    Yhigh = angles[j] + Ystep * (Yfrac/Ydiff);

    if (trap > 25)
    {
        X = 996;          // This 'special' value is output to indicate a problem
        Y = 996;          // in the calibration/interpolation section
        i=12;
        j=12;
    } else {
        // Step 3 - calculate X for interpolated Y
        X = Xhigh*(Yfrac/Ydiff) + Xlow*(1-Yfrac/Ydiff);
        // Step 4 - calculate Y for interpolated X
        Y = Yhigh*(Xfrac/Xdiff) + Ylow*(1-Xfrac/Xdiff);
    }
}

////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////
// Output current readings and perform housekeeping
////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////
#ifdef OUTPUT_TYPE_BINARY
    out[0] = 'D';
    out[1] = 'W';
    out[2] = stageX;
    memcpy(&out[3], &satX, 4);
    memcpy(&out[7], &X, 4);
    memcpy(&out[11], &Y, 4);
    for (k=0;k<15;k++)
    {
        USART_PutChar(DEBUG_USART, out[k]);
    }
#endif

#ifdef OUTPUT_TYPE_RAW
    sprintf(out, "X: %6.3f\tY: %6.3f\tSat: %i\tStage: %i\t%f6.3\n", Xmeas, Ymeas,
        (int)satX, (int)stageX, (float)SystemTickCount/1000.);
    USART_Puts(DEBUG_USART, out);
#endif

#ifdef OUTPUT_TYPE_ANGLE
    sprintf(out, "X: %6.3f\tY: %6.3f\tSat: %i\tStage: %i\t%f6.3\n", X, Y,
        (int)satX, (int)stageX, (float)SystemTickCount/1000.);
    USART_Puts(DEBUG_USART, out);
#endif

#ifdef SDLogging
    // Monitor SD card
    if(SDBufferOverrun==1){
        SDBufferOverrun=0;
        USART_Puts(DEBUG_USART, (char *)"\nSD buffer overrun...\n");
    }
}

```

```

// Write full buffers to SD card
if(SDBuffer1Full==1)
{
    f_write(&fileID, (char *)SDWriteBuffer1, SDBuffer1Len, bytesWritten);
    SDBuffer1Full = FALSE;
    //USART_Puts(DEBUG_USART, (char *)"\nBuffer 1 full...\n");
}
if(SDBuffer2Full==1)
{
    f_write(&fileID, (char *)SDWriteBuffer2, SDBuffer2Len, bytesWritten);
    SDBuffer2Full = FALSE;
    //USART_Puts(DEBUG_USART, (char *)"\nBuffer 2 full...\n");
}
#endif

// Regulate measurement rate
while(loopTick >= SystemTickCounter-200)
{
    LEDToggle(LED_ORANGE);
}

// End of main loop. Terminated by pushing blue button.

LEDOff(LED_GREEN);

// Output last set of raw data
#ifndef OUTPUT_TYPE_BINARY
USART_Puts(DEBUG_USART, (char *)"\nLast data set\n");
for (channel=0;channel<2;channel++)
{
    for (sample=0;sample<nsampleX;sample++)
    {
        sprintf(out, "sample: %i, channel: %i, val: %i\n", sample, channel,
            data[channel][sample]);
        USART_Puts(DEBUG_USART, out);
    }
}
for (channel=2;channel<4;channel++)
{
    for (sample=0;sample<nsampleY;sample++)
    {
        sprintf(out, "sample: %i, channel: %i, val: %i\n", sample, channel,
            data[channel][sample]);
        USART_Puts(DEBUG_USART, out);
    }
}
#endif

// Clean up and close file

```

```

if(SDWhichBuff == 1)
{
    f_write(&fileID, (char *)SDWriteBuffer1, SDBuffer1Len, bytesWritten);
} else {
    f_write(&fileID, (char *)SDWriteBuffer2, SDBuffer2Len, bytesWritten);
}
f_close(&fileID);

USART_Puts(DEBUG_USART, (char *)"\nEND ORVS Test 1\n");

LEDOff(LED_BLUE);
while(1);
}

/*-----
* brief Move data message to the current SD buffer.
* param inMsg: specifies the message buffer.
* retval None
*-----*/
void BuffMsg(char inMsg[], uint16_t length)
{
    int i;

    if (SDWhichBuff == 1)
    {
        // First make sure there is space available
        if(length >= SD_BUFF_SIZE-SDBuffer1Len)
        {
            // Not enough space available. Make sure the other buffer is available
            if(SDBuffer2Full)
            {
                SDBufferOverrun = 1;    // Set flag and do not copy this message
            } else {
                // Here we copy the message to buffer 2
                // First clear remainder of buffer 1
                for (i=SDBuffer1Len;i<SD_BUFF_SIZE;i++)
                {
                    SDWriteBuffer1[i]=0;
                }
                SDWhichBuff = 2;
                SDBuffer1Full = TRUE;
                for (i=0;i<length;i++)
                {
                    SDWriteBuffer2[i] = inMsg[i];
                }
                SDBuffer2Len = i;
            }
        } else {
            // Space available in buffer 1 - copy message

```



```

        for (i=0; i<length; i++)
        {
            SDWriteBuffer1[SDBuffer1Len++] = inMsg[i];
        }
    }

} else {    // This case is if buffer 2 is the current buffer

    // First make sure there is space available
    if(length >= SD_BUFF_SIZE-SDBuffer2Len)
    {
        // Not enough space available. Make sure the other buffer is available
        if(SDBuffer1Full)
        {
            SDBufferOverrun = 1;    // Set flag and do not copy this message
        } else {
            // Here we copy the message to buffer 1
            // First clear remainder of buffer 2
            for (i=SDBuffer2Len;i<SD_BUFF_SIZE;i++)
            {
                SDWriteBuffer2[i]=0;
            }
            SDWhichBuff = 1;
            SDBuffer2Full = TRUE;
            for (i=0;i<length;i++)
            {
                SDWriteBuffer1[i] = inMsg[i];
            }
            SDBuffer1Len = i;
        }
    } else {
        // Space available in buffer 2 - copy message
        for (i=0; i<length; i++)
        {
            SDWriteBuffer2[SDBuffer2Len++] = inMsg[i];
        }
    }
}
}
}

```