THE IMPACT OF GPS VELOCITY BASED FLIGHT CONTROL ON FLIGHT INSTRUMENTATION ARCHITECTURE

Richard P. Kornfeld, R. John Hansman and John J. Deyst

International Center for Air Transportation
Department of Aeronautics & Astronautics
Massachusetts Institute of Technology
Cambridge, MA 02139 USA

June 1999

ICAT-99-5

The Impact of GPS Velocity Based Flight Control on Flight Instrumentation Architecture

by

Richard P. Kornfeld, R. John Hansman and John J. Deyst

Abstract

This thesis explores the use of velocity information obtained by a Global Positioning System (GPS) receiver to close the aircraft's flight control loop. A novel framework to synthesize attitude information from GPS velocity vector measurements is discussed. The framework combines the benefits of high-quality GPS velocity measurements with a novel velocity vector based flight control paradigm to provide a means for the human operator or autopilot to close the aircraft flight control loop. Issues arising from limitations in GPS as well as the presence of a human in the aircraft control loop are addressed.

Results from several flight tests demonstrate the viability of this novel concept and show that GPS velocity based attitude allows for equivalent aircraft control as traditional attitude. Two possible applications of GPS velocity based attitude, an autopilot and a tunnel-in-the-sky trajectory guidance system, are demonstrated in flight. Unlike traditional autopilot and trajectory guidance systems, these applications rely solely on the information obtained from a single-antenna GPS receiver which makes them affordable to the larger General Aviation aircraft community. Finally, the impact of GPS velocity based flight control on the instrumentation architecture of flight vehicles is investigated.

This document is based on the thesis of Richard P. Kornfeld submitted to the Department of Aeronautics and Astronautics at the Massachusetts Institute of Technology in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Aeronautics and Astronautics.

Acknowledgments

Much of this work was supported by a gift from Rockwell-Collins. In particular, the authors would like to thank Dr. Patrick Hwang and Tom Sharpe for always contributing their expertise. The authors also would like to acknowledge the support of the NASA/FAA Joint University Program for Air Transportation and the support of Draper Laboratory. The authors would like to thank Scott Lewis from Northstar and the people from Novatel Technical Support for their help throughout the research.

The research documented in this report would not have been possible without the participation of all the pilots who volunteered to conduct the flight tests. Thanks to all of them.

Thank you Rahman Henderson, Keith Amonlirdviman and Liz Walker for doing an outstanding job in planing, conducting and evaluating some of the flight experiments documented in this report.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction and Overview

The emergence of the Global Positioning System (GPS) as a source of high-quality velocity information to a world-wide user community and the development of other novel sensor technologies offer the potential to increase the integrity of flight instrumentation while at the same time reducing cost. This thesis presents the development and demonstration of a concept which enables GPS-based velocity information to be used to close the aircraft flight control loop. The impact of this novel concept on flight instrumentation architecture is investigated.

In this chapter, the motivation and the objectives of the research documented in this thesis are described. This is followed by an overview of the thesis.

1.1 Motivation

In the past decade, the Global Positioning System (GPS) emerged as a source of high-quality navigation and time information to a world-wide user community. The accuracy of the GPS position and velocity information has thus far only been achieved by inertial navigation systems (INS). With the proliferation of GPS and the expansion of the GPS user community, GPS receiver production has reached a growth where economy of scale principles apply. Consequently, at the time of publication, a standard GPS chip set is available for less than US$200 which is two orders of magnitude less than an INS with comparable performance.

This research documented in this thesis was motivated by the availability of low-cost, high-quality GPS velocity information and addressed the question of the potential impact of this information on the instrumentation architecture of flight vehicles. It was investigated how the availability of GPS velocity information can change existing flight instrumentation architectures in terms of integrity and cost. In addition, the potential for new flight instrumentation concepts based on GPS velocity information was examined.

This thesis focuses on the use of GPS velocity information for aircraft flight control. This is in contrast to the traditional use of GPS velocity information for aircraft guidance. A methodology to synthesize attitude information from GPS velocity measurements has been developed. It combines the benefits of high-quality velocity measurements with a novel flight control paradigm that controls the aircraft velocity vector directly, rather than through attitude as in traditional control schemes. The availability of GPS velocity based attitude information, termed *pseudo-attitude*, creates unique opportunities for new applications.

GPS-based attitude information can greatly *increase the integrity* of cockpit systems. For instance, its use as a backup attitude indicator for General Aviation (GA) aircraft provides the pilot with an additional level of attitude redundancy. Furthermore, GPS-based pseudo-attitude constitutes a source of attitude information that is functionally *independent* from attitude measured by traditional inertial sensor based systems and provides therefore *dissimilar* redundancy. This attitude information can be used in fault detection and isolation schemes as tie-breaker or cross-reference thereby greatly increasing cockpit integrity.

With the availability of GPS-based attitude information, a single-antenna GPS receiver can provide all the information necessary to control and guide aircraft. Classes of aircraft, such as expendable small unmanned aerial vehicles (UAV) which recently began to emerge, can be *instrumented with a single-antenna GPS receiver as the primary sensor.* This has significant weight, size, power and cost advantages compared to traditional instrumentation architectures.

Furthermore, the availability of single-antenna GPS-based position, velocity and flight control information enables the implementation of *autopilot and trajectory guidance systems* solely based on GPS information. Guidance systems such as *tunnel-in the-sky* and *flight director displays* which thus far have relied on expensive sensor hardware can now be implemented using a single-antenna GPS receiver. This has significant system integration and cost advantages and consequently allows the larger General Aviation community to benefit from these systems.

A number of these applications are discussed in this thesis.

1.2 Objectives of the Thesis

The objectives of this thesis are the development and the demonstration of a novel framework within which GPS velocity vector information is used to the close the aircraft flight control and attitude loop. The framework combines the benefits of high-quality GPS velocity measurements with a novel velocity vector based flight control paradigm to provide a means for the human operator or autopilot to close the aircraft flight control loop. The development of the framework takes a human-centered approach to ensure adequate pilot usability. Implementation issues and limitations are addressed, and opportunities identified. A number of applications are implemented and demonstrated in flight.

1.3 Organization of the Thesis

This thesis is divided in three parts. Chapter 2 to 5 deal with the theoretical background of the GPS velocity based flight control concept. This creates the groundwork for the experimental setup and the flight demonstrations discussed in Chapter 6 to 9. Finally, Chapter 10 explores the implications and applications of this novel concept and provides a summary and conclusions. In detail, this thesis is organized as follows:

Chapter 2 discusses the current aircraft control loop structure and presents different flight instrumentation architectures. It then introduces the notion of GPS velocity based flight control and investigates its impact on the flight instrumentation architecture.

Chapter 3 explains the concept of velocity based flight control in more detail. The notion of velocity based pseudo-attitude is introduced and its synthesis from aircraft velocity and acceleration information is discussed. A novel pseudo-attitude display is presented and display update rate and latency are discussed. The results of a preliminary simulator study on the effectiveness of velocity based attitude and the required display update rate are presented.

Chapter 4 discusses the different aspects of the Global Positioning System (GPS) which are pertinent to the generation of GPS-based velocity and acceleration information. The principle of operation and the observable of GPS are briefly explained. The GPS receiver architecture and operation, and the algorithms to generate velocity and acceleration information from the GPS observables are discussed. The bandwidth and related trade-offs as well as error sources of GPS velocity and acceleration are examined. Finally, GPS integrity, availability and continuity issues relevant to the generation of velocity and acceleration information are highlighted.

In Chapter 5 the loop closure around GPS velocity based pseudo-attitude information is discussed. This chapter relies on the concepts and insights presented in Chapter 3 and 4. A linearized analysis is used to investigate the open- and closed-loop behavior of pseudo-attitude based flight control. An autopilot design is discussed as a case example.

Chapter 6 introduces the objectives of the flight tests and discusses the flight test setup. The implementation of a flight test system is described. The hardware as well as the software aspects of the instrumentation are addressed. Initial testing efforts are outlined and an overview of the flight tests is given.

Chapter 7 discusses the experimental evaluation of the pseudo-attitude system. The flight test objectives and the flight test protocol are outlined. Objective and subjective results are presented and discussed.

Chapter 8 presents the flight demonstration of a pseudo-attitude based flight director / autopilot system. The flight test objectives, the setup and the flight test protocol are outlined, and objective results are presented and discussed.

Chapter 9 presents the flight demonstration and the evaluation of pseudo-attitude based tunnel-in-the-sky trajectory guidance systems. Two perspective flight path displays are flight tested and compared to traditional ILS guidance scheme. Flight test objectives, setup and flight test protocol are outlined. Objective and subjective results are then presented and discussed.

Chapter 10 summarizes the research work documented in this thesis and examines the implications and applications of GPS-based flight controls.

Chapter 2

The Impact of GPS Velocity Based Flight Control on Flight Instrumentation Architecture

This chapter introduces the concept of GPS velocity based flight control and proposes a novel flight instrumentation architecture that is based on this concept. The new architecture is compared to a number of current flight instrumentation architectures. It is shown that the new architecture has the potential to drastically reduce the cost of flight instrumentation.

The chapter starts with a description of the classical aircraft control loops in Section 2.1 and gives an overview over current flight instrumentation architectures in Section 2.2. Section 2.3 introduces the notion of GPS velocity vector based flight control and discusses the novel flight instrumentation architecture. It also presents an instantiation of the new architecture which was used throughout most of the research presented in this thesis.

2.1 Classical Aircraft Control Loops

Classical aircraft control schemes rely on a multi-loop feedback design in which the different loops are nested within each other. In a typical flight control configuration the loops are the guidance loop, the flight control or attitude command loop, and the stability or control augmentation loop, as shown in Figure 2.1. The different loops and the task they perform are described in the following:

- *Guidance Loop*: Navigation, in the context of this thesis, refers to the process of establishing the position and velocity state of the aircraft. Guidance refers to the process of using this information to command the vehicle to follow a pre-defined trajectory. The guidance loop, thus, generates guidance commands by differencing the desired and measured aircraft position and velocity states and feeding them to the next inner loop, the flight control or attitude command loop so as to reduce the state deviation and ensuring that the aircraft flies along the desired track.

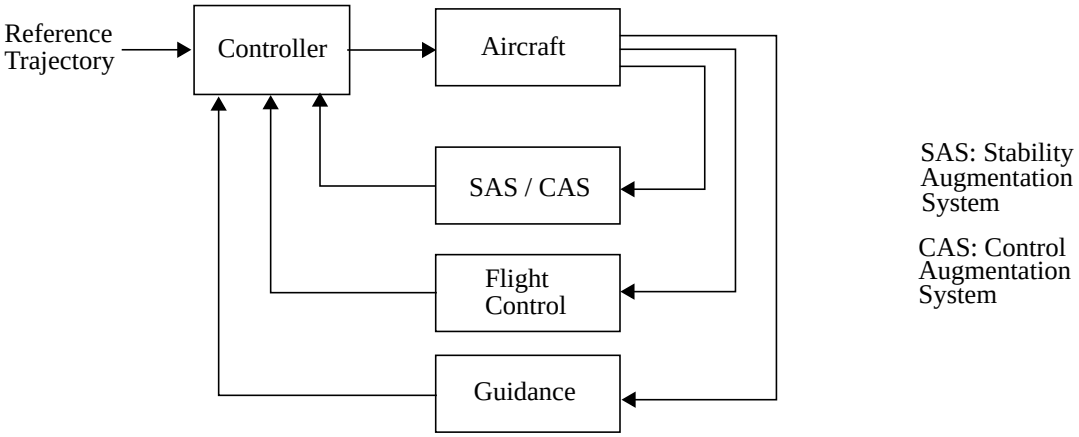

Figure 2.1: Classical Flight Control Loops

- *Flight Control Loop or Attitude Command Loop*: The flight control or attitude command loop is used to change the aircraft state in order to follow the guidance commands. This loop, thus, generates flight control commands and feeds them to the aerodynamic control surfaces and the engine controls causing the aircraft to achieve the commanded aircraft state. Autopilots needed to provide 'pilot relief' typically operate in this loop.

- *Stability Augmentation System (SAS) or Control Augmentation System (CAS) Loop*: The stability augmentation system is the inner most loop and is used to suppress the effects of unwanted inherent aircraft modes such as the dutch roll in the lateral or the short period in the longitudinal direction. It thereby facilitates the design of the outer loops and insures that the outer loops function properly. The modes are typically excited by aerodynamic control deflections and gust disturbances. By feeding back appropriate control variables in the SAS loop their effects can be damped out and their response decay time decreased. The control augmentation system (CAS) loop improves the transient response properties of the aircraft and provides the pilot with a particular type of response to the control inputs. This loop often enhances the inherent deficient aircraft modes as well.

The next two sections briefly present typical lateral and longitudinal aircraft control schemes and the feedback variables used to close the individual loops.

2.1.1 Lateral Feedback Loop Closure

Figure 2.2(a) shows a typical lateral feedback structure. The lateral guidance loop is commonly closed using horizontal position and velocity information, denoted d and $\dot{d}$ (or, equivalently, heading information, ψ). The measured position and velocity state is differenced from the desired state given by the reference trajectory. This deviation, using appropriate guidance laws, generates bank angle commands ϕ_c which are fed to the flight control loop. This loop commands aileron deflections δ_a to achieve the desired roll angle. The banked aircraft experiences a sideward acceleration that changes the velocity vector in the direction commanded by the guidance laws. The blocks containing K denote the respective gains and compensators.

The lateral flight control loop is, thus, closed using the roll angle ϕ as feedback variable to follow the commanded bank angle and regulate against disturbances. This configuration controls the velocity and acceleration vector indirectly by commanding and controlling the aircraft roll angle which in turn generates the aircraft acceleration necessary to change the direction of the velocity vector.

A yaw damper is often employed as a SAS to improve the dutch roll behavior. A yaw rate feedback with washout circuitry K_r which generates commands δ_r to the rudder is normally sufficient to dampen this mode. A roll CAS feeding back roll rate p is sometimes used as an additional inner loop to improve the roll response.

2.1.2 Longitudinal Feedback Loop Closure

Figure 2.2(b) shows a typical longitudinal feedback closure. The guidance variables controlled are altitude h, air speed u, and vertical speed $\dot{h}$ or flight path angle γ. The deviation from the desired aircraft states, using the appropriate guidance laws, results in pitch attitude and airspeed commands, θ_c and u_c, fed into the flight control loop. This loop generated elevator δ_e and throttle δ_{th} inputs to achieve the commanded pitch attitude and airspeed using knowledge of the aircraft dynamics. This feedback structure controls the flight path state indirectly through the control of pitch attitude.

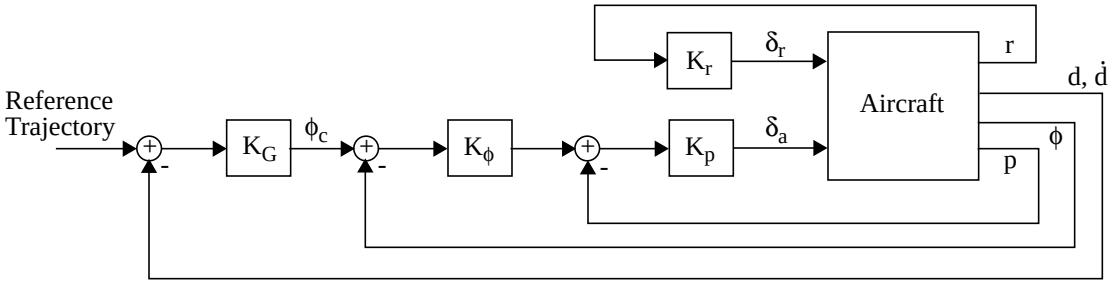

(a)

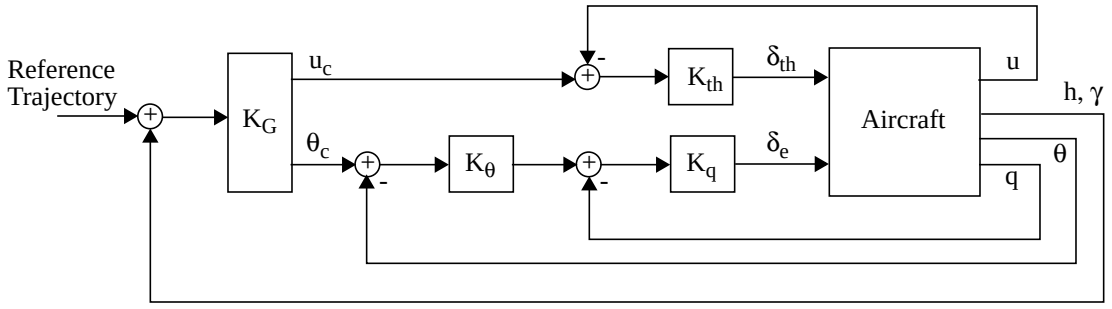

(b)

Figure 2.2: (a) Lateral Flight Control Loops (b) Longitudinal Flight Control Loops

A pitch damper, using pitch rate q feedback is typically added if the aircraft short period mode is not well damped. Its elevator commands are added as a high frequency component to the elevator command of the pitch attitude loop.

2.2 Current Instrumentation Architectures

This section gives an overview over traditional instrumentation architectures used to close the loops outlined in the previous section. With regard to this thesis, the scope is limited mainly to inertial sensors and radio navigation instrumentation. Air data sensors are not considered because the thesis primarily focuses on the attitude command or flight control loop closure where air data is of minor relevance.

First, traditional Attitude and Heading Reference System (AHRS) and Inertial Navigation System (INS) based instrumentation architectures are discussed. Next, an INS/ GPS based architecture is considered which uses the synergy of inertial and single-

antenna GPS information. Finally, a more recent multi-antenna GPS-based instrumentation architecture is presented that relies on carrier phase measurements to close the flight control loop.

2.2.1 AHRS and INS Based Instrumentation Architecture

Figure 2.3 shows an AHRS based instrumentation architecture. It relies on measurements of inertial quantities such as turn rate and acceleration to calculate all the necessary feedback variables. At the heart of the Attitude and Heading Reference System is an Inertial Measurement Unit (IMU).† It consists of at least three gyros and accelerometers, mounted typically on three orthogonal axes, and thus senses accelerations and turn rates in three dimensions. The accelerations and turn rates are available as outputs for feedbacks in SAS and CAS loops.

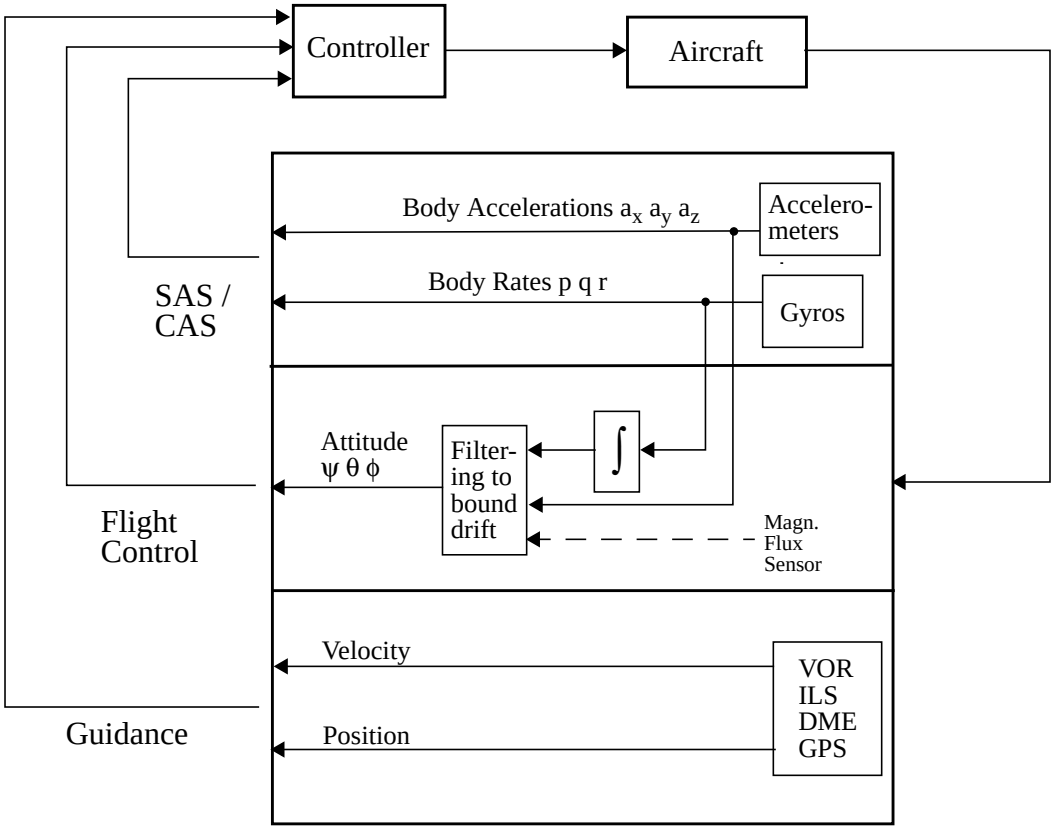

Figure 2.3: AHRS Based Instrumentation Architecture

† Only the more recent strap-down systems are considered here.

The turn rates are integrated once to yield traditional aircraft attitude which is typically expressed in Euler angles (heading angle ψ, pitch angle θ and roll angle ϕ). Due to bias and drift inherent in the gyro sensor measurements, the attitude as a result of the integration process will drift over time. Acceleration measurements can be used to bound the attitude drift. In the absence of considerable aircraft dynamics or averaged over a longer time interval, the accelerometers measure the gravity vector and, thus, act like a mechanical pendulum to indicate the vertical direction. The acceleration measurements, averaged over time, are then used in an electronic erection loop to prevent the roll and pitch attitude to drift over time. The averaging time constant is subject to a trade-off: too short of a time constant will cause aircraft accelerations to be sensed as gravitational acceleration thereby nulling any indicated aircraft pitch and roll indications. On the other hand, too long of a time constant will not bound the attitude drift effectively enough. For typical AHRS implementation the time constant is in the order of few minutes. A drawback of this mechanization is apparent if the aircraft maneuvering time exceeds the average time constant (such as in an extended steady turn). In that case, the actual aircraft acceleration will be interpreted as gravitational acceleration and any indicated aircraft attitude will be reset over time.

While the roll and pitch angle drift can be bounded using accelerometer measurements, the azimuth or heading angle drift can not, and thus necessitates the availability of additional sensor information such as magnetic compass or magnetic flux sensor measurements. These measurements will then be incorporated in the attitude filtering to bound the heading drift. This is indicated with the dashed line in Figure 2.3.

The attitude drift is an inherent limitation of AHRS instrumentation and drives the performance requirements of the rate gyros used. For AHRS in aircraft applications, high-end tactical grade gyros (~0.1 deg/hr), such as fiber optic gyros, are typically necessary in order to achieve the required attitude accuracy (Tazartes 1995, Schmidt 1997). As a consequence of the high gyro quality required, the cost of AHRS traditionally range from $30k to $50k at the time of publishing. Its use is, therefore, normally limited to commercial and business aircraft.[†]

AHRS based flight control instrumentations have the advantages of being entirely self-contained and of having high availability, high sensor bandwidth and low sensor noise characteristics.

In an AHRS based instrumentation architecture the guidance loop is typically closed using radio-navigation aids or Doppler radar information. Radio-navigation systems include systems such as DME, VOR, ILS and GPS.

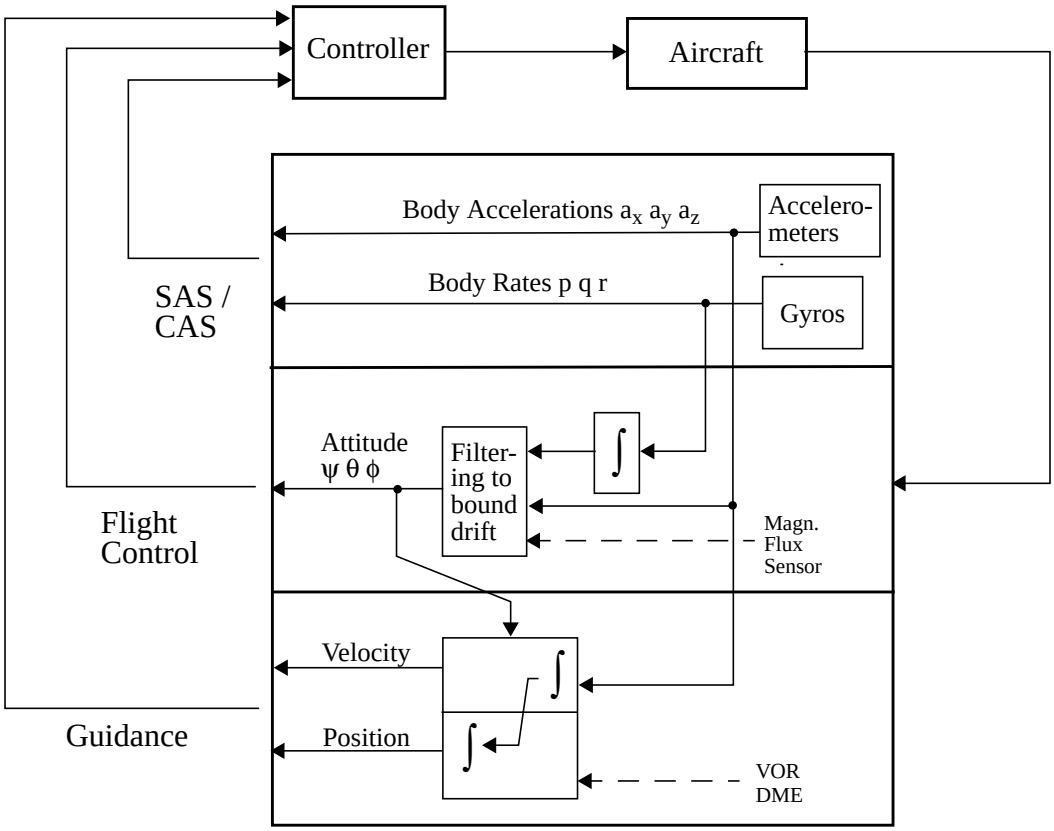

Figure 2.4: INS Based Instrumentation Architecture

† For low-end aircraft, such as for the large fleet of General Aviation (GA) aircraft, the AHRS functions are performed by a gimballed vertical gyro for pitch and roll and a directional gyro for heading information. Traditionally, these instruments are mechanical, and driven by a vacuum pump or electrically. They are thought to be prone to failures, but their cost is a fraction of the cost of traditional AHRS. Moreover, due to their entirely mechanical nature, these instruments provide no electronic output of the attitude information.

In an Inertial Navigation System (INS) based instrumentation architecture, the guidance loop is closed using position and velocity information calculated from the accelerometer measurements. Figure 2.4 shows an INS based instrumentation architecture. The measurements are transformed from the aircraft body axes into a suitable reference frame using the calculated attitude and are integrated once for velocity and a second time for position estimates. This information may be blended with position and velocity measurements obtained from additional radio-navigation aids.

The limitation of inertial systems due to gyro and accelerometer drift rates become apparent by considering the number of integrations necessary to obtain position and velocity from accelerations and turn rates. Each integration increases the rate at which the resulting quantity drifts over time. Acceptable position and velocity drift rates require, therefore, gyro and accelerometer of high performance. Typically, for an INS system of 1.0 nmi/hr accuracy, navigational grade accelerometer and gyros with biases in the order of 10^{-5} g and 10^{-2} deg/hr, respectively, are necessary (Phillips 1996). Ring laser gyros are commonly used for this application. This results in high cost for INS based instrumentation architectures. The cost of inertial navigation systems currently exceeds $50K and architectures based on INS are limited to the higher end of the aircraft spectrum. Advantages of INS based architectures include the availability of independent and self-contained position and velocity information, high sensor bandwidth and low sensor noise characteristics.

2.2.2 INS/GPS Based Instrumentation Architecture

With the advent of the Global Positioning System (GPS), instrumentation concepts have been developed which synergistically combine the properties of inertial sensors and GPS measurements. Inertial sensors have high bandwidth and low noise and, thus, good high frequency behavior, but their inherent biases give raise to drift in attitude, position and velocity. GPS, on the other hand, typically provides noisy position and velocity measurements at limited bandwidth, but the measurements are absolute and therefore not

affected by any drift problem. The good low frequency behavior of GPS, thus, complements the good high frequency behavior of inertial sensors in an optimal way. Figure 2.5 shows a possible INS/GPS based instrumentation architecture.

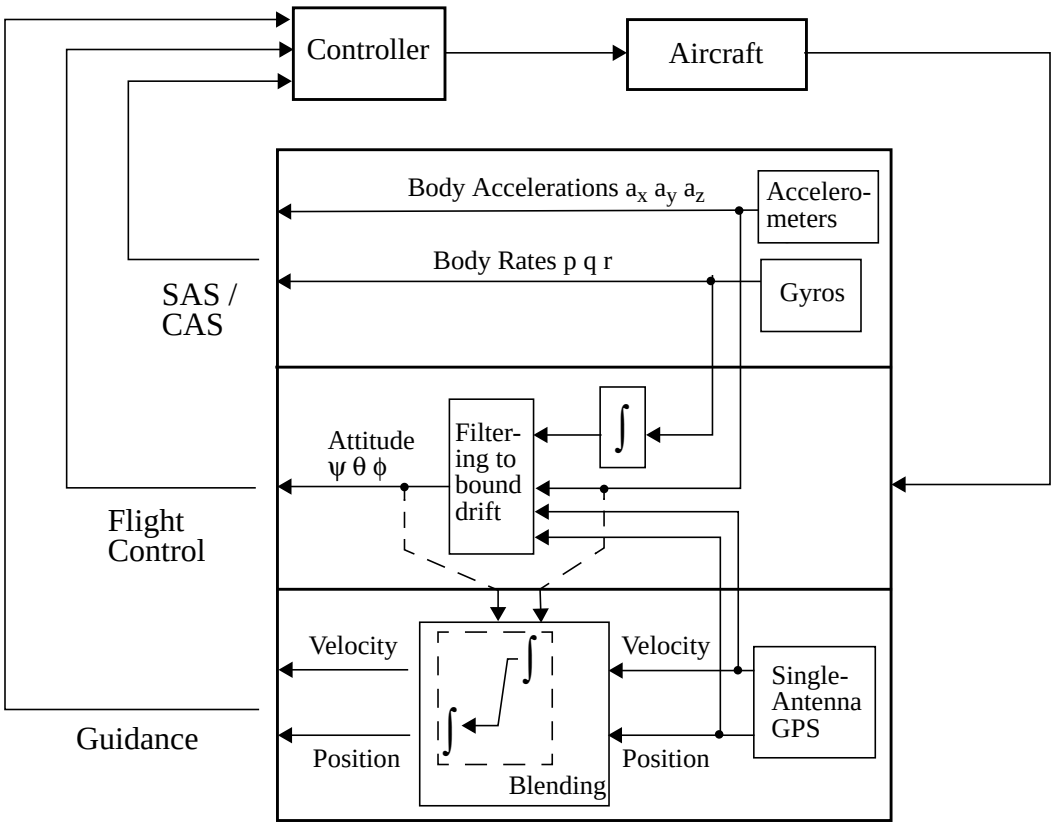

Figure 2.5: INS/GPS Based Instrumentation Architecture

As part of the synergism, GPS information calibrates inertial sensor errors and reduces the drift in inertial based attitude, position and velocity information. This is done with the use of an error model that relates position and velocity errors to gyro and acceleration errors. At the same time, the inertial based (attitude, position and velocity) information is available at significant higher rate than GPS information. Inertial information may also be used to help the GPS receiver during satellite acquisition.

In the architecture discussed in the previous section, the GPS sensor is used purely as a navigation aid and its information is used to close the guidance loop. Here, GPS position and velocity information is used to correct the aircraft attitude and is blended with inertially derived position and velocity information, that is, the GPS position and velocity measurements are used for the flight control loop and guidance loop. The integration of GPS and inertial information as shown in Figure 2.5 is commonly referred to as 'loosely coupled' GPS/INS architecture (Phillips 1996). Alternate integration concepts range from separate INS and GPS systems, with GPS information resetting the INS solution periodically, to tightly coupled INS/GPS systems where raw GPS (pseudo-range and delta-range) and inertial measurements are combined in a single filter.

The use of GPS measurements to calibrate the inertial sensor errors allows for a reduction in the sensor quality required to achieve comparable performance as with an AHRS. Typically, low-end tactical grade gyros and accelerometers, such as fiber optic or micromachined tuning fork quartz gyros and micromachined vibrating quartz accelerometers, with biases of the order of 10 deg/hr and 10^{-3} g are used resulting in lower costs (Boeing 1997). INS/GPS tactical grade units range from $8k to $20k and are commonly employed in the mid and high-end aircraft segment.

As an additional advantage, an integrated INS/GPS system provides the user with attitude information during GPS outages or jamming. The calibration of the inertial sensor errors hereby reduces the rate at which the attitude information drifts.

2.2.3 Multi-Antenna GPS-Based Instrumentation Architecture

In recent years multi-antenna GPS-based attitude sensors have been developed (Cohen 1996). They rely on interferometric principles to determine the vehicle attitude. By measuring the difference in GPS carrier phase between a pair of antennae, the receiver determines the range difference between the pair of antennae and the satellite. Range differences obtained using carrier phase measurements from multiple satellites with three or more antennae with known baselines allow the receiver to compute three-axis attitude of a vehicle.

The GPS receiver initially only measures the fractional part of the differential phase. The integer part of the range difference, corresponding to multiple of the GPS carrier wavelength, must be determined by independent means before the differential phase measurement can be interpreted as a differential range measurement. This problem is commonly referred to as integer ambiguity resolution and numerous algorithm to solve for this ambiguity have been implemented.

Multi-antenna GPS-based attitude determination is a direct measurement of the vehicle attitude and hence is not affected by drift problems. Since the principal observables are carrier phase difference measurements, it is not susceptible to Selective Availability. However, its accuracy is proportional to the inverse of the antenna baseline lengths. Thus, larger baselines reduce the attitude error and hence the vehicle dimensions constrain the achievable accuracy.

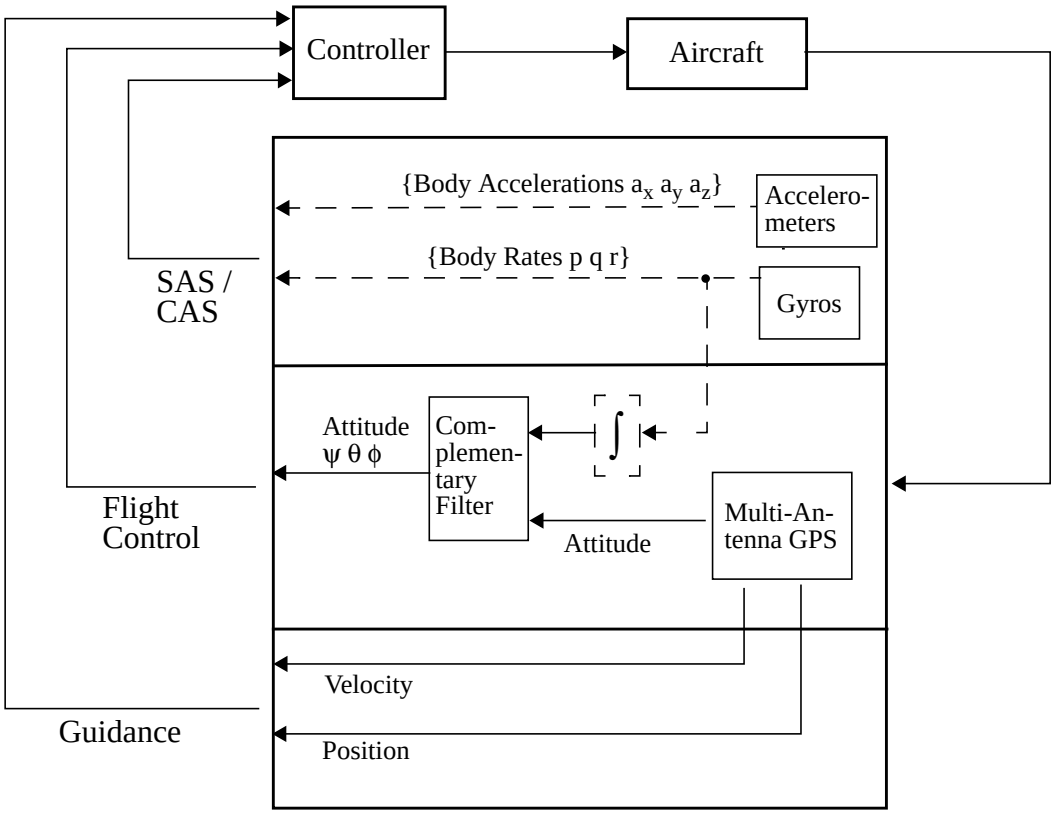

Figure 2.6: Multi-Antenna GPS-Based Instrumentation Architecture

Figure 2.6 shows a typical multi-antenna GPS-based flight instrumentation architecture. The GPS attitude sensor provides roll, pitch and yaw information for the flight control loop closure. At the same time, the GPS receiver can be configured to provide position and velocity information that is used to close the navigation and guidance loop. This information, unlike the GPS derived attitude, is obtained from ranging measurements to four or more satellites using a single antenna. If necessary, additional gyros and accelerometers can be used to dampen unwanted aircraft motions or to improve the aircraft response in the SAS/CAS loop. This is indicated through the dashed lines in Figure 2.6.

It is interesting to note that in the multi-antenna GPS-based architecture, inertial sensors are used to *augment* GPS attitude information, whereas in the integrated INS/GPS based architecture the converse is true. That is, GPS augments the INS based attitude information.

The inertial sensor performance necessary for SAS/CAS loop closures is significantly lower than the inertial sensor performance necessary to close the flight control loop. The fact that primarily high-frequency components of accelerometer and gyro outputs are fed back in this loop, makes their biases and drift rates less significant and allows the use of low-cost automotive grade inertial sensors for this task. Gyros and accelerometers with a typical performance of 180 deg/hr and 1 mg, respectively, may be utilized (Gebre 1998, Schmidt 1997). The cost of such sensors are less than $100 (for large volumes) at the time of publication and are projected to decrease significantly in the future.

A synergy exists between GPS attitude and inertial sensor based attitude, similar to the one found in the combined integration of INS/GPS systems discussed in Section 2.2.2. By combining the drift free, low bandwidth, GPS attitude information with high bandwidth, drift-affected, inertial sensor based attitude information in a complementary filter, the advantages of both sensors are exploited. This is shown in Figure 2.6. GPS attitude is used to calibrate the rate gyro biases on-line and, at the same time, a higher attitude update rate is available using the calibrated inertial attitude output. The availability of actual GPS attitude measurements (as opposed to GPS position and velocity measurements) allows for

an improved gyro bias calibration and significantly reduces the performance requirements of the gyros. Also, properly calibrated inertial sensors allow for continued operation during temporary GPS outages.

Hayward (1997) and Gebre (1998) implemented an GPS/Inertial AHRS for General Aviation (GA) applications utilizing three antennae and automotive grade gyros. The antennae were configured in a isosceles triangle with baselines of 50 cm and 36 cm. They demonstrated a attitude accuracy of better than 0.2 deg.

Though the availability of multi-antenna receivers on the market is still limited, this concept has the potential to serve a larger GA community in the future. Some of the disadvantages associated with multi-antenna GPS-based systems are the extensive antenna installations and baseline calibrations, the aircraft specific certification, the currently limited GPS integrity and the need for ambiguity resolution after loosing lock and subsequent reacquisition.

2.3 GPS Velocity Vector Based Flight Control - A New Flight Control Paradigm

This thesis presents a new paradigm for closing the flight control loop. It has the distinct advantage that the control variables necessary to close the loop are completely observable from single-antenna GPS measurements. This creates the opportunity for a flight instrumentation architecture that can be primarily based on a single-antenna GPS receiver and has, thus, the potential to greatly reduce instrumentation complexity. The new paradigm is based on sensing and controlling the inertial velocity vector directly, rather than controlling it through aircraft attitude, as in conventional control schemes.

Traditionally, the pitch and roll attitude are controlled to achieve a desired velocity vector and, hence, flight path change. In the lateral direction, an aircraft roll angle is established in order to generate an acceleration force that, in turn, changes the velocity vector. Similarly, in traditional pitch attitude based longitudinal flight control schemes, pitch and/or thrust control is used to achieve a desired flight path angle and speed.

In the proposed flight control paradigm, the flight path vector and its rate of change are sensed and controlled directly. In order for this approach to be successful, however, the aircraft has to be well behaved ‘around the velocity vector’. That is, unwanted aircraft modes have to be satisfactorily damped and the aircraft response to control inputs has to be adequate.

The velocity vector, flight path angle and the acceleration vector are completely observable from single-antenna GPS measurements. The GPS receiver measures high-quality carrier Doppler frequency shifts that are used to compute velocity information. Acceleration can be inferred by backdifferencing or by Kalman filtering the velocity information.

A useful *representation* of velocity vector based flight control variables is *flight path angle* in the longitudinal direction and *pseudo-roll* angle in the lateral direction. Flight path angle is the angle between the inertial velocity vector and the local level plane, and is used as a surrogate of pitch angle. Pseudo-roll angle represents the roll angle around the velocity vector axis and is a substitute for traditional roll angle. Pseudo-roll corresponds to the observed lateral rate of change of the velocity vector and is determined from the acceleration vector perpendicular to the velocity vector. The combined use of these attitude or flight control variables is novel and is referred to as *pseudo-attitude* to distinguish them from traditional attitude. Chapter 3 explains the derivation of pseudo-attitude in much greater detail. It will be shown that for coordinated flight the pseudo-roll angle closely corresponds to the traditional roll angle and that, therefore, similar control strategies can be employed as for traditional roll angle.

The new flight control paradigm enables an instrumentation architecture that is primarily based on a single-antenna GPS receiver. This architecture is discussed in Section 2.3.1. Section 2.3.2 briefly presents a demonstration example of this new architecture, namely a pseudo-attitude system for General Aviation aircraft. The rest of this thesis is devoted to the derivation, synthesis and flight test of pseudo-attitude and the demonstration of some of its applications.

2.3.1 Single-Antenna GPS-Based Instrumentation Architecture

Figure 2.7 shows a possible single-antenna GPS-based instrumentation architecture. A single-antenna GPS receiver is used as a primary means to obtain position and velocity information which serves to close the guidance loop. Using the pseudo-attitude concept, velocity information is also used to close the flight control or attitude command loop. This is accomplished by inferring acceleration and subsequently synthesizing pseudo-attitude to close the loop.

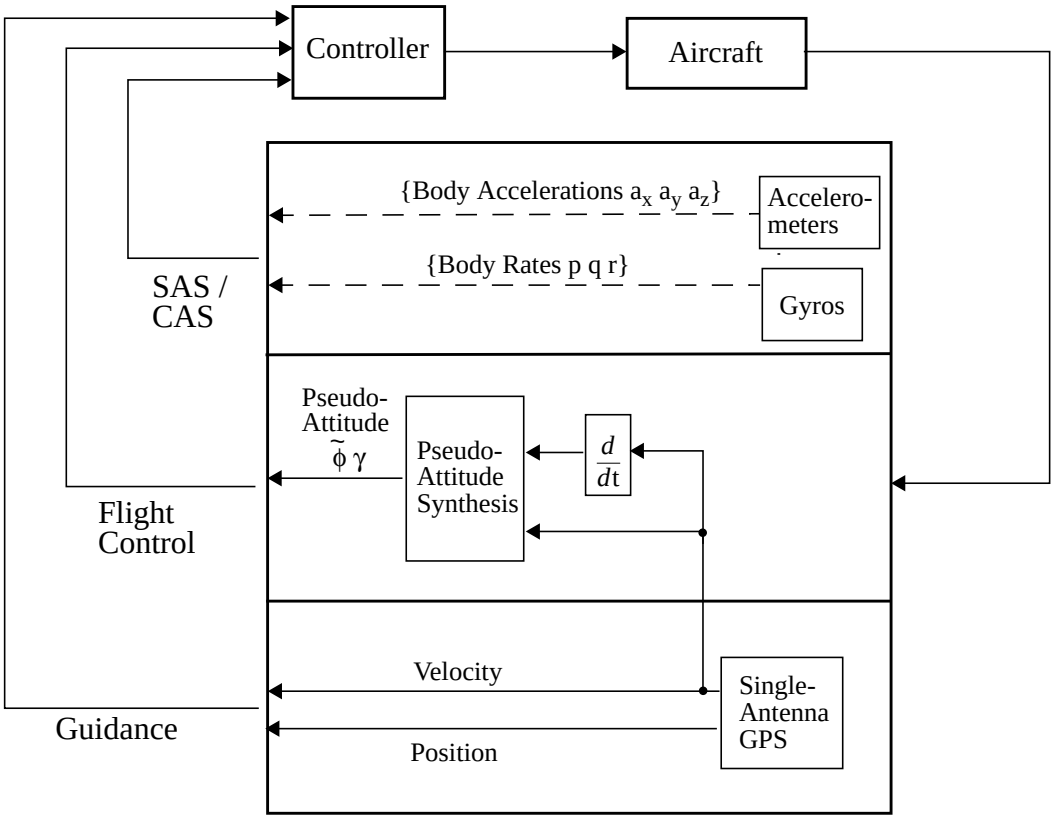

Figure 2.7: Single-Antenna GPS-Based Instrumentation Architecture

If necessary, additional inertial sensors can be employed to dampen unwanted aircraft modes. In Figure 2.7, a SAS/CAS loop, if required, is schematically shown by the dashed arrows. Since primarily high-frequency components of accelerometer and gyro outputs are fed back in this loop, their biases and drift rates are less significant. Consequently, low-cost gyros and accelerometers of automotive grade can be used for these tasks. Their cost are typically less than $100 at the time of publication.

Similar to the multi-antenna GPS-based architecture, this architecture uses low-cost inertial sensors only to *augment* the GPS-based pseudo-attitude, as part of a SAS/CAS loop or to increase the bandwidth. This is in contrast to traditional INS/GPS based architectures where GPS is used to correct the primary INS derived attitude solution.

The single-antenna GPS-based instrumentation architecture relies on a sensor that, in recent years, has found broad acceptance as a navigation aid and is, thus, readily available at affordable cost. At the time of publication, GPS receivers with update rates as high as 10 Hz are available for $2-3k†. Hence, this architecture may be implemented at significant lower cost than traditional instrumentation architectures.

A distinct advantage of pseudo-attitude is the fact that it constitutes an absolute measurement of the aircraft state and, hence, provides drift free attitude information. An additional advantage is the minimal required installation. Unlike multi-antenna GPS attitude, where a number of antennae have to be installed and their baselines have to be known or estimated, this architecture relies on the installation of single antenna, ideally close to the center of gravity. In addition, no ambiguity resolution is necessary to operate the system. Finally, no initial alignment as for inertial sensors is required.

However, the use of an outer-loop variable for inner loop control, that is, the differentiation of velocity to obtain acceleration information is tied to a trade-off involving noise and bandwidth of acceleration information, and thus sets an inherent limit on the achievable performance. Furthermore, GPS availability and integrity issues have to be addressed. This is of particular importance since this architecture utilizes GPS information not only in the guidance but also in the flight control or attitude command loop. For example, GPS outages or jamming can lead to the loss of attitude control. These issues will be addressed in later Chapters.

---

† Cost of the receiver hardware. Certification, if necessary, is not included.

2.3.2 Demonstration Example: Pseudo-Attitude System for General Aviation Aircraft

This section presents a simple example of the single-antenna GPS-based instrumentation architecture, namely a pseudo-attitude system for small General Aviation (GA) aircraft. This system was used as prototype to demonstrate the concept of GPS-based pseudo-attitude because of its relative simplicity in implementing and testing it.

Figure 2.8 shows the pseudo-attitude system as part of the attitude command loop which, in this case, is closed by the pilot. The need for inner-loop stabilization of small GA aircraft is typically greatly diminished by the inherent design of these aircraft. In addition, the pilot is assumed to fly the aircraft in a coordinated manner by compensating any experienced sideforce by appropriate rudder inputs. The navigation and guidance loop, not shown in Figure 2.8, is assumed to be closed by the pilot using available GPS position and velocity information.

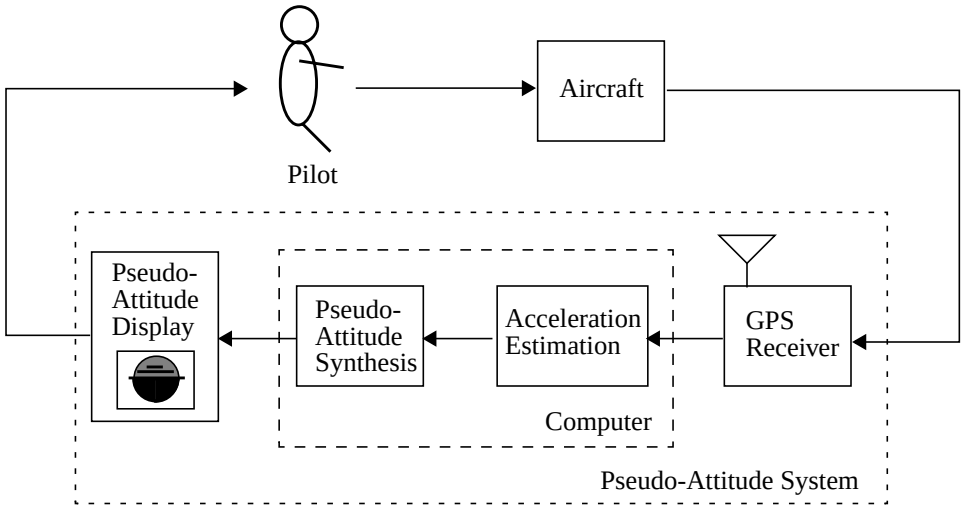

Figure 2.8: Pseudo-Attitude Based Flight Control Loop

The pseudo-attitude system consists of a GPS receiver providing three-dimensional velocity information, a computer executing the pseudo-attitude synthesis algorithm, and a display showing the pseudo-attitude information. In the implementation described in this thesis, a Novatel 3151R GPS receiver, operated in single-point mode, was used as the

primary velocity source. Aircraft velocity and acceleration, however, are necessary to synthesize pseudo-attitude. The velocity output is therefore fed into a Kalman filter which estimates acceleration. The velocity and acceleration information are the input to the pseudo-attitude synthesizing algorithm which, together with the Kalman filter, is implemented in an onboard computer. The calculated pseudo-attitude is then displayed on an electronic pseudo-attitude display.

Chapter 3

Velocity Vector Based Flight Control

This chapter describes the synthesis and display of velocity vector based attitude information. In particular, it introduces the notion of *pseudo-attitude* which is a useful representation of attitude information and is completely observable from velocity and acceleration information. Pseudo-attitude, thus, enables a GPS receiver providing high quality Doppler shift based velocity and acceleration measurements to be used for flight control loop closures.

Pseudo-attitude consists of flight path angle and pseudo-roll angle, defined as a rotation about the velocity vector axis. Its synthesis is based on a simple point mass aircraft model that does not require particular knowledge about the aircraft. It will be shown that under coordinated flight conditions pseudo-roll corresponds closely to traditional roll angle. Furthermore, pseudo-attitude allows for functionally similar flight control loop closures in the longitudinal and lateral directions as traditional attitude. This is discussed further in Chapter 5.

Section 3.1 discusses the synthesis of velocity based pseudo-attitude information. In addition, the synthesis of traditional attitude information is considered in order to explore the differences between pseudo-attitude and traditional attitude. The treatment in this section assumes perfect velocity and acceleration measurements. The realistic sensor characteristics of GPS are treated in Chapter 4 and their implications on loop closures are discussed in Chapter 5. Section 3.2 introduces a possible pseudo-attitude display that allows the human to close the attitude or flight control loop around velocity based attitude information. Section 3.3 summarizes this chapter.

3.1 Derivation of Velocity Vector Based Attitude information

The aircraft *flight control problem* has the objective of determining the appropriate control surface deflections to achieve a desired attitude and, ultimately, to execute a desired flight path change. Determining the aircraft attitude when the aircraft trajectory is given may, thus, be called the *inverse problem* of aircraft general motion and has been subject of some research (Kato 1986, Hauser 1997). Kato (1986) shows that the inverse problem of the airplane general motion has no unique solution and indicates that if additional constraints, such as coordinated flight are included, the aircraft attitude is unique to the given flight path and may, at least in principle, be determined by solving three simultaneous implicit nonlinear equations. No solution for the case of coordinated flight is offered, however.

This section develops an approach to synthesize attitude information from the aircraft trajectory, that is, from the velocity along the aircraft flight path. The synthesis relies on the assumption of coordinated (or nearly coordinated) flight. This implies that the aircraft flies in a manner so as to eliminate any sideforce experienced by pilots and passengers. Flight coordination is typically achieved by banking the airplane to obtain a desired turn rate and by the use of rudder input to cancel any residual sideforce. The assumption of coordinated flight is valid for most flight conditions encountered by conventional aircraft and, thus, does not constitute a significant limitation to this concept. Furthermore, it will be shown in later chapters that the synthesized attitude information is also useful in steady uncoordinated flight, such as in the presence of a non-zero steady sideslip.

The attitude information synthesized from the aircraft trajectory has been termed *pseudo-attitude* to distinguish it from traditional attitude consisting of pitch and roll angles (Kornfeld 1998a, b). In contrast to traditional attitude, which is referenced to the aircraft body axes, pseudo-attitude is referenced to the aircraft velocity vector and consists of flight path angle γ with respect to the (local horizontal) ground plane, substituting for traditional aircraft pitch angle, and a pseudo-roll angle $\tilde{\phi}$ about the aircraft velocity vector axis, substituting for traditional roll angle. The pseudo-roll angle is defined as the effective

bank angle which corresponds to the observed lateral rate of change of the velocity vector. Pseudo-attitude, unlike traditional attitude, provides a direct indication of the flight path. Figure 3.1 illustrates the definition of pseudo-attitude.

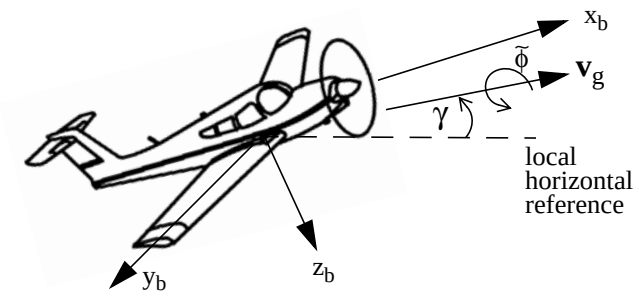

Figure 3.1: Illustration of Pseudo-Attitude

The synthesis of pseudo-attitude is treated in Section 3.1.4. In order to provide a basis for the pseudo-attitude algorithm, the synthesis of attitude in wind axes and of traditional attitude in body axes, under the assumption of coordinated flight, is discussed first in Section 3.1.2 and Section 3.1.3, respectively. The insights gained from the synthesis of these attitude variables help explain the approach taken to synthesize pseudo-attitude and help identify the commonalities and differences between them.

To simplify the discussion in this chapter, the following assumptions are made:

- The aircraft is assumed to be in coordinated flight.

- The atmosphere is assumed to be in (nearly) uniform motion relative to the Earth. Thus if **w** is the wind velocity, then

$$\frac{d}{dt}\mathbf{w} \approx 0 \tag{3.1}$$

  This assumption will in fact not be valid all the time, since the atmosphere is usually in non-uniform motion in time or space. However, by choosing the wind vector **w** as the space and time average of the atmospheric motion over an appropriate space-time domain, an average uniform motion can be assumed

(Etkin 1972). The local or temporal deviations from the mean atmospheric motions are turbulence or gusts, and will be treated as disturbance inputs to the attitude determination system.

- The earth is assumed to be flat and fixed in space, thereby neglecting the effects of earth curvature and rotation. This assumption is valid for flight velocities encountered by most conventional aircraft (typically subsonic or low supersonic speeds < Mach 3) (Etkin 1972).

- The aircraft is assumed to be a rigid body having a plane of symmetry.

The assumption of coordinated flight and uniform wind motion are relaxed in later sections and their effect on pseudo-attitude is discussed in Section 3.1.5 and in Chapter 5. The next section, presents the coordinate frames used throughout this chapter. Finally, Section 3.1.6 briefly describes other velocity based control variables.

3.1.1 Coordinate Frames

The following coordinate frames are used:

- The NED frame $\mathbf{F}_{NED}$ is an earth-fixed, local level coordinate system which has its origin instantaneously located at the current position of the aircraft center of gravity (c.g.), and its axes aligned with the directions of North, East and the local vertical (Down). Due to the flat-earth assumption, the NED frame is, in effect, treated as an inertial reference frame in which accelerations and angular rates are measured. The velocity of the aircraft c.g. relative to $\mathbf{F}_{NED}$, i.e. with respect to the ground, is denoted $\mathbf{v}_g$ and expressed in NED coordinates as $\mathbf{v}_g = (v_{gN}, v_{gE}, v_{gD})$.

- The body frame $\mathbf{F}_B$ is a body-fixed, orthogonal axes set which has its origin at the aircraft c.g., and its axes aligned along the roll (x_b), pitch (y_b) and yaw (z_b) axes of the aircraft. That is, the x_b-axis extends forward out the vehicle's nose, the y_b-axes extends out the right wing, and the z_b-axis extends out the bottom of the vehicle. The x_b-z_b plane is usually a plane of geometric symmetry. Since the aircraft

rotational inertia matrix is constant in body axes, the rotational dynamics equations are typically expressed in this frame. The frame $\mathbf{F}_B$ has angular velocity $\omega = (p,q,r)$ relative to $\mathbf{F}_{NED}$.†

- The atmosphere-fixed reference frame $\mathbf{F}_A$ is moving with respect to the NED frame $\mathbf{F}_{NED}$ at wind velocity **w**. Its origin is at the aircraft c.g. and its axes are parallel to the NED axes. Because of the assumption of uniform wind motion, this frame is an inertial frame. The velocity of the aircraft c.g. relative to $\mathbf{F}_A$, is referred to as the *relative wind*, and is the relevant velocity for aerodynamic forces in atmospheric flight. It is denoted $\mathbf{v}_a$ and expressed in NED coordinates as $\mathbf{v}_a = (v_{aN},v_{aE},v_{aD})$.

- The wind axes frame $\mathbf{F}_W$ is an orthogonal axes set which has its origin fixed to aircraft c.g., and the x_w-axis pointing into the relative wind, i.e x_w is directed along the velocity vector $\mathbf{v}_a$. The z_w-axis lies in the plane of symmetry of the aircraft (i.e. in the x_b-z_b plane), and y_w is orthogonal to x_w and z_w to form a right hand sided coordinate frame. The frame $\mathbf{F}_W$ has angular velocity $\omega_w = (p_w,q_w,r_w)$ relative to $\mathbf{F}_{NED}$.

The orientation of the body frame $\mathbf{F}_B$ with respect to the NED frame $\mathbf{F}_{NED}$ is given by three angles, namely yaw ψ, pitch θ and roll ϕ angle. These are the Euler angles of the body axes. The rotation matrix $\mathbf{R}(\psi,\theta,\phi)$ which transforms any vector from $\mathbf{F}_{NED}$ to $\mathbf{F}_B$ is

† The stability axes system $\mathbf{F}_S$ is a special body axes reference frame used primarily in the study of small disturbances from a steady reference flight condition. The x_s-axis is chosen to lie on the projection of $\mathbf{v}_a$ onto the aircraft plane of symmetry, the z_s-axis is in the body x_b-z_b plane and the y_s-axis is orthogonal to x_s and z_s to form a right hand sided coordinate frame. For a symmetric flight condition (no sideslip β) $\mathbf{F}_S$ coincides with the wind axes $\mathbf{F}_W$ in the reference condition, but departs from it, moving with the body, during the disturbance. The stability axis frame will be used in context of the linearized analysis of pseudo-attitude in Chapter 5.

then composed of three consecutive rotations about the respective yaw, pitch and roll axes, in that order, as shown in Figure 3.2(a), and given by

$$\begin{aligned}\mathbf{R}(\psi,\theta,\phi) &= \mathbf{R}(\psi)\cdot\mathbf{R}(\theta)\cdot\mathbf{R}(\phi)\\ &= \begin{bmatrix}1 & 0 & 0\\ 0 & \cos\phi & \sin\phi\\ 0 & -\sin\phi & \cos\phi\end{bmatrix}\cdot\begin{bmatrix}\cos\theta & 0 & -\sin\theta\\ 0 & 1 & 0\\ \sin\theta & 0 & \cos\theta\end{bmatrix}\cdot\begin{bmatrix}\cos\psi & \sin\psi & 0\\ -\sin\psi & \cos\psi & 0\\ 0 & 0 & 1\end{bmatrix}\end{aligned} \tag{3.2}$$

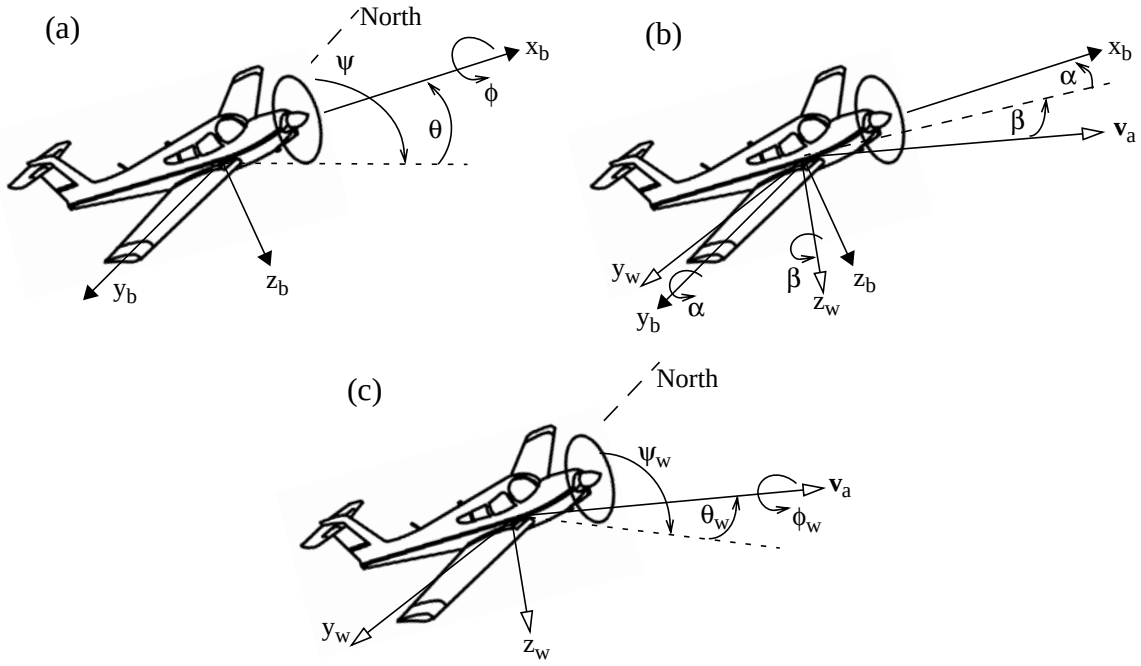

Figure 3.2: Definition of (a) Euler Angles in Body Axes, (b) Aerodynamic Angles (c) Euler Angles in Wind Axes

Similarly, the rotation $\mathbf{R}(\psi_w,\theta_w,\phi_w)$ that defines the orientation of the wind axes frame $\mathbf{F}_W$ relative to $\mathbf{F}_{NED}$ is given by the Euler angles (ψ_w,θ_w,ϕ_w) of the wind axes system. The definition of the Euler angles of the wind axes is shown in Figure 3.2(c).

The body frame $\mathbf{F}_B$ and the wind axes frame $\mathbf{F}_W$ are related through the angle of attack α and sideslip angle β.† The angle of attack is the angle between the body x_b-axis and the projection of the relative wind $\mathbf{v}_a$ onto the aircraft plane of symmetry. The sideslip angle is

† To simplify the discussion, it is assumed here that the body x_b-axis is parallel to the basic aerodynamic reference direction, i.e the mean aerodynamic chord or the zero lift line. In this case, α, as used here, will be the same as that commonly used in aerodynamic theory and in wind tunnel testing. Otherwise, it differs by a constant offset

the angle between $\mathbf{v}_a$ and the aircraft plane of symmetry. These are the aerodynamic angles and, together with the relative wind $\mathbf{v}_a$, are of fundamental importance in determining the aerodynamic forces that act on the aircraft. The rotations that carry $\mathbf{F}_W$ into $\mathbf{F}_B$ are shown in Figure 3.2(b). The rotations are about the z_w-axis by the sideslip angle β and about the y_b-axis by α, in that order. In terms of Euler angles this sequence of rotation can be expressed as $\mathbf{R}(-\beta,\alpha,0)$ (Etkin 1972). Thus, the relations between Euler angles of the body and the wind frames are (Kato 1986):

$$\mathbf{R}(\psi, \theta, \phi) = \mathbf{R}(-\beta, \alpha, 0) \cdot \mathbf{R}(\psi_w, \theta_w, \phi_w) \tag{3.3}$$

From this equation, the Euler angles of the body axes can be obtained in terms of Euler angles of the wind axes and the aerodynamic angles α, β, and vice versa. Figure 3.3, adapted from (Stengel 1977), summarizes the relations between the different reference frames and shows the rotations necessary to carry one frame into another.

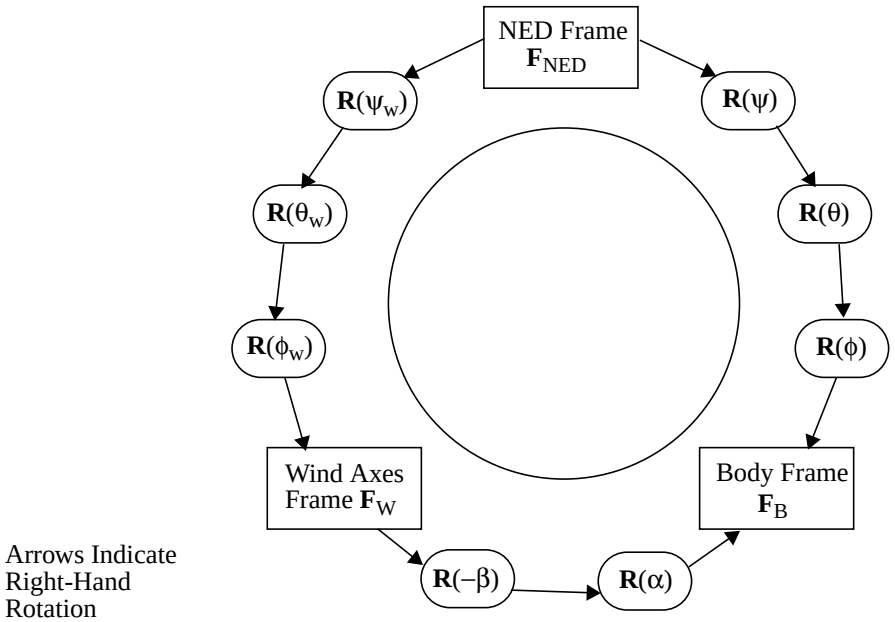

Figure 3.3: Reference Frame Transformations

For coordinated flight the sideslip β is zero and Eq. (3.3) is modified to

$$\mathbf{R}(\psi, \theta, \phi) = \mathbf{R}(0, \alpha, 0) \cdot \mathbf{R}(\psi_w, \theta_w, \phi_w) \tag{3.4}$$

In this case, the y_b axis coincides with the y_w axis, and the wind axes frame relates to the body axes frame through a single rotation about the y-axes by the angle of attack α.

In the following treatment it is assumed that the trajectory, or equivalently, the velocity $\mathbf{v}_g$ and acceleration $\mathbf{a}_g$ with respect to the ground are given (e.g. obtained from measurements of a GPS receiver). For the synthesis of attitude in wind and body axes, it is further assumed that the wind vector **w** and the angle of attack α are known, respectively.

3.1.2 Wind Axes Attitude Synthesis

In this section, the synthesis of attitude in wind axes, (ψ_w,θ_w,ϕ_w), is discussed. Given $\mathbf{v}_g = (v_{gN},v_{gE},v_{gD})$, the velocity of the aircraft center of gravity with respect to the ground, and the wind vector **w**, the velocity of the aircraft center of gravity with respect to the wind frame, $\mathbf{v}_a = (v_{aN},v_{aE},v_{aD})$, is determined by

$$\mathbf{v}_a = \mathbf{v}_g - \mathbf{w} \tag{3.5}$$

The direction of $\mathbf{v}_a$ is aligned with the wind x_w-axis. The heading angle in wind axes or, equivalently, the track angle with respect to the surrounding air, ψ_w, is given by

$$\psi_w = \operatorname{atan}\left(\frac{v_{aE}}{v_{aN}}\right) \tag{3.6}$$

where the +/-π ambiguity has to be resolved using the signs of the velocity components. The pitch angle in wind axes or, equivalently, flight path angle with respect to the surrounding air, θ_w, is given by

$$\theta_w = \operatorname{atan}\left(\frac{-v_{aD}}{\sqrt{v_{aN}^2 + v_{aE}^2}}\right) \tag{3.7}$$

where a positive θ_w indicates a climb.

For the determination of the roll angle in wind axes, ϕ_w, the following considerations are necessary. The acceleration $\mathbf{a}_a$ of the aircraft center of gravity with respect to the atmosphere-fixed frame $\mathbf{F}_A$ is obtained by differentiating $\mathbf{v}_a$ with respect to time in $\mathbf{F}_A$. Differentiating Eq. (3.5) and invoking the assumption in Eq. (3.1) yields

$$\mathbf{a}_a = \left.\frac{d\mathbf{v}_a}{dt}\right|_A \approx \left.\frac{d\mathbf{v}_g}{dt}\right|_{NED} = \mathbf{a}_g \tag{3.8}$$

that is, $\mathbf{a}_a$ is equivalent to the given aircraft acceleration with respect to the ground, $\mathbf{a}_g$. Separating the acceleration $\mathbf{a}_a$ into components which are tangential and normal to the flight path vector $\mathbf{v}_a$ yields

$$\mathbf{a}_a = \mathbf{a}_a^t + \mathbf{a}_a^n \tag{3.9}$$

where the superscripts “t” and “n” denote the tangential and normal components, respectively. They are given by

$$\mathbf{a}_a^t = \frac{\mathbf{a}_a \cdot \mathbf{v}_a}{|\mathbf{v}_a|^2}\mathbf{v}_a \tag{3.10}$$

$$\mathbf{a}_a^n = \mathbf{a}_a - \mathbf{a}_a^t \tag{3.11}$$

Similarly, the gravitational acceleration $\mathbf{g}$ is divided into tangential and normal components by

$$\mathbf{g}^t = \frac{\mathbf{g} \cdot \mathbf{v}_a}{|\mathbf{v}_a|^2}\mathbf{v}_a \tag{3.12}$$

$$\mathbf{g}^n = \mathbf{g} - \mathbf{g}^t \tag{3.13}$$

and $\mathbf{g}$ is expressed in the coordinates of the NED frame as

$$\mathbf{g} = \begin{bmatrix} 0 \\ 0 \\ g_o \end{bmatrix} \tag{3.14}$$

where $g_o = 9.81\ m/s^2$ and is assumed to be constant. Then, because coordinated flight is assumed, the normal component of the aircraft acceleration is equal to the sum of the lift acceleration vector **l** and the normal component of gravity, i.e.

$$\mathbf{a}_a^n = \mathbf{l} + \mathbf{g}^n \tag{3.15}$$

Figure 3.4 shows the relevant accelerations with the velocity vector $\mathbf{v}_a$ pointing into the page. The lift acceleration vector **l** is, by definition, perpendicular to the flight path vector $\mathbf{v}_a$ and can in turn be expressed as the difference of the normal component of the aircraft acceleration and the normal component of gravity, i.e.

$$\mathbf{l} = \mathbf{a}_a^n - \mathbf{g}^n \tag{3.16}$$

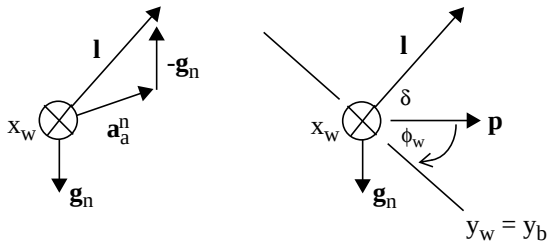

Figure 3.4: Relevant Forces for the Synthesis of Attitude in Wind Axes

In aircraft flight, the direction of the lift vector **l** is controlled by the aircraft roll angle and, consequently, carries information about the aircraft bank angle. In order to extract the roll information, a horizontal reference vector **p** is defined as

$$\mathbf{p} = \mathbf{g} \times \mathbf{v}_a = \mathbf{g}^n \times \mathbf{v}_a \tag{3.17}$$

The angle δ between **l** and **p** is determined by

$$\delta = \operatorname{acos}\left(\frac{\mathbf{l} \cdot \mathbf{p}}{|\mathbf{l}| \cdot |\mathbf{p}|} \right) \tag{3.18}$$

The complementary angle to δ, i.e. the angle between **p** and the wind y_w-axis, which is perpendicular to **l** and $\mathbf{v}_a$, is the *roll angle in wind axes*, ϕ_w, and is given by

$$\phi_w = \text{asin}\left(\frac{\mathbf{l} \cdot \mathbf{p}}{|\mathbf{l}| \cdot |\mathbf{p}|}\right) \tag{3.19}$$

Appendix A offers a proof of the roll synthesis that serves as a complementary illustration to the geometric argument given above. The proof uses the force equation of the aircraft equations of motion to show that Eq. (3.8) - Eq. (3.19) yield the roll angle in wind axes.

3.1.3 Traditional Attitude Synthesis

If the angle of attack is non-zero, the aircraft attitude in body axes is determined from the attitude in wind axes by performing an additional rotation about the common y-axis by the angle of attack, as discussed in Section 3.1.1. This rotation does not introduce any sideforce and, thus, the resulting attitude in body axes still corresponds to a coordinated flight condition.

The evaluation of the attitude rotation equation given in Eq. (3.4), i.e.

$$\mathbf{R}(\psi, \theta, \phi) = \mathbf{R}(0, \alpha, 0) \cdot \mathbf{R}(\psi_w, \theta_w, \phi_w) \tag{3.20}$$

yields the following useful expressions

$$\sin\theta = \cos\alpha \cdot \sin\theta_w + \sin\alpha \cdot \cos\theta_w \cdot \cos\phi_w \tag{3.21}$$

$$\sin\phi = (\cos\theta_w \cdot \sin\phi_w)/(\cos\theta) \tag{3.22}$$

$$\sin(\psi_w - \psi) = (-\sin\alpha \cdot \sin\phi)/\cos\theta_w \tag{3.23}$$

from which the attitude in body axes, (ψ,θ,ϕ), can easily be calculated if α is known. For the case of $\alpha = 0$, the above expressions simplify to

$$\sin\theta = \sin\theta_w, \quad \sin\phi = \sin\phi_w, \quad \sin(\psi_w - \psi) = 0 \tag{3.24}$$

as expected. For the special case of level flight ($\phi_w = 0$), Eq. (3.21) yields the familiar relation

$$\sin\theta = \cos\alpha \cdot \sin\theta_w + \sin\alpha \cdot \cos\theta_w = \sin(\theta_w + \alpha) \tag{3.25}$$

i.e. for level flight pitch angle is the sum of flight path angle with respect to the surrounding air and the angle of attack.

In many cockpit architectures, the angle of attack is not available as a measurement and may have to be estimated instead, using the aircraft lift curves, the flight altitude and the aircraft weight. A simple approach is outlined in (Case 1996). Similarly, in the case where the wind vector is unknown it may be inferred from in-flight measurements of the velocity vector $\mathbf{v}_g$ during a constant airspeed turn. This approach is outlined in (Hollister 1990) in the context of Loran measurements. Often, however, accurate lift curves or the aircraft weight are not available, and sufficiently accurate estimates of α or the wind vector cannot be obtained. These shortcomings may limit the applicability of the body axes attitude synthesis in practice.

Difference Between Attitude in Wind Axes and Body Axes

In order to compare body axis attitude with pseudo-attitude described in the next subsection, it is useful to first quantify the difference between body axis attitude and wind axes attitude. Using a first order Taylor approximation, it can be shown from Eq. (3.21) - Eq. (3.22) that the difference between ϕ and ϕ_w is given by

$$\phi - \phi_w \approx \tan\theta_w \cdot \sin\phi_w \cdot \alpha \tag{3.26}$$

and does not exceed 1.9 deg in magnitude for the flight conditions limited by

$$-15° < \theta_w < 15°, \quad -45° < \phi_w < 45°, \quad -10° < \alpha < 10° \tag{3.27}$$

Thus, the roll angle in body axes is approximately equal to the roll in wind axes, i.e.

$$\phi \approx \phi_w \tag{3.28}$$

in the typical operating range of most civil aircraft. The difference between the aircraft pitch angle in body and wind axes is to first order

$$\theta - \theta_w \approx \cos\phi_w \cdot \alpha \tag{3.29}$$

and is largest for level flight. In this case, the difference equals to the angle of attack, as expected from Eq. (3.25). Finally the difference between ψ and ψ_w is given in Eq. (3.23).

3.1.4 Pseudo-Attitude Synthesis

Pseudo-attitude is a surrogate for traditional aircraft attitude and is synthesized from the inertial velocity and acceleration vectors $\mathbf{v}_g$ and $\mathbf{a}_g$, respectively. Pseudo-attitude is therefore entirely observable from GPS velocity and acceleration measurements and its synthesis does not rely on the additional knowledge of angle of attack or wind information. High quality GPS velocity information is available from Doppler frequency shift measurements of the GPS carrier tracking loop as discussed in more detail in the next chapter.

Pseudo-attitude is referenced to the aircraft velocity vector and consists of flight path angle in the longitudinal direction and pseudo-roll angle around the aircraft velocity vector axis in the lateral direction. Flight path angle γ is the angle between the velocity vector and the local level plane, and is used as a surrogate for pitch angle. Pseudo-roll angle $\tilde{\phi}$ is defined as the effective bank angle which corresponds to the observed lateral rate of change of the velocity vector and is a substitute for traditional roll angle. Pseudo-attitude, unlike traditional attitude, provides a direct indication of the flight path. It will be shown, however, that for coordinated flight the pseudo-roll angle closely corresponds to traditional roll angle. Figure 3.1 shows the velocity vector axis, $\mathbf{v}_g$, used for the definition of pseudo-attitude, and the body axes, (x_b, y_b, z_b), used for the definition of the traditional attitude.

Flight path angle is defined as the angle between $\mathbf{v}_g$ and the local level ground plane and is given by

$$\gamma = \text{atan}\left(\frac{-v_{gD}}{\sqrt{v_{gN}^2 + v_{gE}^2}}\right) \tag{3.30}$$

where a positive γ indicates a climb.

The pseudo-roll $\tilde{\phi}$ is determined from the known aircraft acceleration $\mathbf{a}_g$ and the gravitational acceleration $\mathbf{g}$ as shown in Figure 3.5. The force diagram in this figure closely resembles the force diagram in Figure 3.4 which shows an aircraft flying a coordinated turn. The only difference is that in the diagram of Figure 3.5 the inertial velocity vector axis, i.e. the axis aligned with $\mathbf{v}_g$, is used to resolve the forces instead of the aircraft wind x_w-axis, which is the axis aligned with the velocity vector $\mathbf{v}_a$ relative to the air.

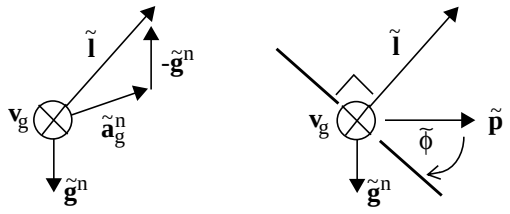

Figure 3.5: Determination of Pseudo-Roll

Consequently, the pseudo-roll angle is determined in a manner similar to the roll in wind axes, described above, but with $\mathbf{v}_a$ replaced by $\mathbf{v}_g$. First, a pseudo-lift acceleration vector $\tilde{\mathbf{l}}$ is defined as the vector difference of $\tilde{\mathbf{a}}_g^n$ and $\tilde{\mathbf{g}}^n$, the components of $\mathbf{a}_g$ and $\mathbf{g}$ perpendicular to the aircraft velocity vector $\mathbf{v}_g$, respectively. That is

$$\tilde{\mathbf{l}} = \tilde{\mathbf{a}}_g^n - \tilde{\mathbf{g}}^n \tag{3.31}$$

where similar to Eq. (3.10) - Eq. (3.13)

$$\tilde{\mathbf{a}}_g^t = \frac{\mathbf{a}_g \cdot \mathbf{v}_g}{|\mathbf{v}_g|^2} \cdot \mathbf{v}_g \tag{3.32}$$

$$\tilde{\mathbf{a}}_g^n = \mathbf{a}_g - \tilde{\mathbf{a}}_g^t \tag{3.33}$$

$$\tilde{\mathbf{g}}^t = \frac{\mathbf{g} \cdot \mathbf{v}_g}{|\mathbf{v}_g|^2} \cdot \mathbf{v}_g \tag{3.34}$$

$$\tilde{\mathbf{g}}^n = \mathbf{g} - \tilde{\mathbf{g}}^t \tag{3.35}$$

The tilde (~) is used to distinguish these quantities from the ones used in the synthesis of attitude in wind axis. The pseudo-roll $\tilde{\phi}$ is then determined as the complement of the angle between the pseudo-lift vector and a local horizontal reference $\tilde{\mathbf{p}}$

$$\tilde{\phi} = \text{asin}\left(\frac{\tilde{\mathbf{l}} \cdot \tilde{\mathbf{p}}}{|\tilde{\mathbf{l}}| \cdot |\tilde{\mathbf{p}}|}\right) \tag{3.36}$$

where the local horizontal reference is defined by

$$\tilde{\mathbf{p}} = \mathbf{g} \times \mathbf{v}_g = \tilde{\mathbf{g}}^n \times \mathbf{v}_g \tag{3.37}$$

Ground track heading χ complements the flight path angle and pseudo-roll angle to form a complete set of velocity referenced attitude variables $(\chi, \gamma, \tilde{\phi})$. χ is defined as

$$\chi = \text{atan}\left(\frac{V_{gE}}{V_{gN}}\right) \tag{3.38}$$

where the +/-π ambiguity has to be resolved using the signs of the velocity components. The properties of attitude variables $(\chi, \gamma, \tilde{\phi})$ are examined in the next subsection.

3.1.5 Properties of Pseudo-Attitude

Pseudo-attitude, unlike traditional attitude, provides a direct indication of the aircraft trajectory. Flight path angle and ground track heading give a direct indication of the velocity vector direction and aircraft trajectory. By way of contrast, pitch and heading angle give the direction the aircraft is pointing and provide only an indirect indication of flight path. Lambregts (1979), investigating a velocity vector control steering mode, states that “the inertial flight path angle best characterizes the airplane maneuver in the vertical plane... . In the horizontal plane, ground track is the most representative variable for control of the aircraft along an earth referenced path... .”

In the lateral direction, pseudo-roll angle is a direct indication of the lateral rate of change of the velocity vector or, equivalently, of the lateral aircraft acceleration. For most flight conditions, however, pseudo-roll corresponds closely to traditional roll. This is discussed in the following section.

Difference Between Pseudo-Roll and Roll Angle in Wind Axes

It was previously shown that roll angle in wind axes is approximately equivalent to roll in body axes with a difference of less than two degrees for the flight envelope specified in Eq. (3.27). The difference between roll in wind axes and pseudo-roll, in turn, is due to the different axes chosen to resolve the gravity and aircraft accelerations. For the former, it is the axis aligned with $\mathbf{v}_a$, whereas for the latter the axis aligned with $\mathbf{v}_g$ is used. The difference of the two axes is due to the wind vector $\mathbf{w}$ only, as shown in Eq. (3.5).

To simplify the discussion, the representative case of an aircraft flying a coordinated level turn under the presence of a horizontal wind vector $\mathbf{w}$ is considered. Then, using a second order Taylor approximation, the difference between pseudo-roll $\tilde{\phi}$ and roll in wind axes ϕ_w is given by

$$\phi_w - \tilde{\phi} \approx \frac{1}{2} \cdot \frac{\tan\phi_w}{1 + (\tan\phi_w)^2} \cdot \rho^2 \tag{3.39}$$

where ρ designates the angle between the vectors $\mathbf{v}_a$ and $\mathbf{v}_g$, that is

$$\rho = \operatorname{acos}\left(\frac{\mathbf{v}_a \cdot \mathbf{v}_g}{|\mathbf{v}_a| \cdot |\mathbf{v}_g|}\right) \tag{3.40}$$

For a relatively severe case of an aircraft flying at an airspeed of 100 knots and turning with a roll in wind axes of 45 deg, a 40 knot crosswind induces an instantaneous difference of no more than 2 degrees. Furthermore, taking the difference of roll in wind axes and traditional roll given in Eq. (3.26) into account, it can be concluded that pseudo-roll closely corresponds to traditional roll angle in this flight condition.

So far, the treatment of pseudo-roll assumed a constant wind vector and coordinated flight. These assumptions are not always valid. The effects of atmospheric non-uniformities and of uncoordinated flight on pseudo-roll are therefore discussed next.

Effects of Uncoordinated Flight on Pseudo-Roll

Even though the assumption of coordinated flight is valid most of the time, there are flight conditions, such as during a severe slip or yaw maneuver, or during stall, where this assumption does not hold true. In these instances pseudo-roll angle may differ considerably from the aircraft bank angle. The effects of the uncoordinated flight conditions on pseudo-roll are discussed here:

- Sidesliping and Yawing: Sidesliping may be caused, among other factors, by severe yawing, excessive rudder input or by an undamped dutch roll motion. It induces a sideward acceleration with respect to the velocity vector that, when measured directly by GPS or inferred from GPS trajectory data, is indistinguishable from an acceleration originating from a banked aircraft (i.e a deflected lift vector) during a coordinated turn.† Consequently, the pseudo-roll synthesis interprets the sensed sideforce as a pseudo-roll angle when, in fact, the aircraft is level.

† An accelerometer mounted along the aircraft y_b-axis, however, measures specific side force and allows to distinguish between coordinated and uncoordinated flight.

Unintentional sideslips are typically of transient nature. The decay times are a function of particular aircraft characteristics. For instance, in case of the dutch roll mode, the decay time is determined from dutch roll damping. Transient sideslip effects will be illustrated in greater detail in Chapter 5 in the context of a linearized aircraft control loop analysis. Fortunately, as will be shown, the effects due to sideslip are typically benign and can be mitigated, if necessary, using stability augmentation feedback systems (SAS).

In the case of a steady sideslip maneuver, a constant offset between pseudo-roll and traditional roll can be observed. For instance, for an aircraft flying with wings level and a steady sideslip, a constant pseudo-roll offset exists due to a steady side acceleration. Conversely, when a roll angle is applied to compensate for the sideforce induced by the sideslip and maintain a straight flight path, no pseudo-roll angle is indicated while a real roll angle exists. The latter may occur when flying a twin-engine aircraft with a failed engine. Generally in these conditions, due to the constant nature of the offset, pseudo-roll angle still conveys the information to control the aircraft flight path. That is, a constant offset in pseudo-roll may be mentally removed by the pilot. For the example case of a twin-engine aircraft with a failed engine, pseudo-roll indication could be used to maintain a straight ground track since a non-zero pseudo-roll is indicative of a change in the velocity vector direction. Aircraft control under the presence of steady sideslip conditions, and the effectiveness of the pseudo-attitude display to track a straight and level ground track under sideslip conditions, will be demonstrated in Chapter 7 as part of the experimental evaluation.

- Non-linear Flight Regime, Stall, Spin: The use of pseudo-attitude is most appropriate in linear or quasi-linear flight regimes where the angle between the aircraft body x_b-axis and the velocity vector is small. In stall or post-stall flight conditions, this angle increases dramatically. Typically, while stalling, the aircraft has a large positive pitch angle and is loosing speed rapidly and, subsequently, altitude at an increasing rate. This causes the aircraft velocity vector to point

towards the ground with an increasingly negative flight path angle. In this situation, the lift acceleration decreases rapidly and the main aircraft acceleration is along the velocity vector. In these conditions, the relation between measured sideforce and roll information is lost and pseudo-roll information “degenerates”. The latter is true for spin conditions as well.

In addition, the recovery from a stall or spin is achieved by establishing the appropriate aircraft attitude. The use of velocity vector based pseudo-attitude information does, thus, not provide sufficient information to adequately regain aircraft control in these situations. A complete non-linear simulation of the behavior of pseudo-attitude during stall and spin is beyond the scope of this thesis.

Effects of Atmospheric Non-Uniformities on Pseudo-Roll

In reality, the atmosphere is in non-uniform motion in time or space. Gusts or turbulence, which are local or temporal deviations from the mean atmospheric motion, as well as windshears may affect the aircraft trajectory and subsequently have an effect on pseudo-attitude:

- Turbulence: The disturbance of primary concern is side gust velocity. This disturbance input appears as a sideslip gust input whose effects are twofold. First, the nearly ‘instantaneous’ nature of the gust input induces a momentary side acceleration which, by the same argument as for sideslip, translates into a corresponding instantaneous change in pseudo-roll. Second, the sideslip gust input may excite aircraft dutch roll motion. The associated sideslip oscillations may, in severe cases, induce a corresponding pseudo-roll oscillation. In Chapter 5 the effects of side gusts are illustrated in the context of a linearized aircraft control loop analysis. Fortunately, as will be shown, for moderate turbulence the effects are minor.

- Wind Shear: Horizontal wind shear, such as experienced by climbing or descending through boundary layer motion next to the ground, subjects the aircraft to sustained side acceleration. Similar to sideslip, the side acceleration, when

measured directly by GPS or inferred from GPS trajectory data, cannot be discriminated from an acceleration originating from a banked aircraft during a coordinated turn. Accordingly, a non-zero pseudo-roll angle will result where, in reality, the aircraft is not necessarily banked. However, in these instances pseudo-roll indication could be used to maintain a straight ground track. Because the pseudo-roll angle correlates with the rate of change of the velocity vector, the pseudo-roll indication induces a self-correcting pilot control input. The pilot will tend to bank the airplane to correct for any indicated pseudo-roll angle and thereby maintain a straight ground track.

3.1.6 Other Velocity Based Control Variables

With a pseudo-attitude reference frame fixed, additional velocity vector based control variables, such as pseudo-pitch rate $\tilde{q}$ and pseudo-yaw rate $\tilde{r}$, can be defined. They are calculated from GPS velocity and acceleration information as follows:

$$\tilde{q} = \frac{\tilde{\mathbf{a}}_g^n \cdot \mathbf{i}_L}{|\mathbf{v}_g|} \qquad \text{where} \qquad \mathbf{i}_L = \frac{\tilde{\mathbf{l}}}{|\tilde{\mathbf{l}}|} \tag{3.41}$$

$$\tilde{r} = \frac{\tilde{\mathbf{a}}_g^n \cdot \mathbf{i}_C}{|\mathbf{v}_g|} \qquad \text{where} \qquad \mathbf{i}_C = \frac{\mathbf{v}_g \times \tilde{\mathbf{l}}}{|\mathbf{v}_g \times \tilde{\mathbf{l}}|} \tag{3.42}$$

and $\tilde{\mathbf{l}}$ is defined in Eq. (3.31). For the case of no wind, Eq. (3.41) and Eq. (3.42) yield pitch and yaw rate in wind axes, q_w and r_w. The use of these variables for inner-loop aircraft control is limited due to the bandwidth constraints of GPS acceleration, as shown in the next chapter. This topic is therefore not pursued further.

3.2 The Display of Velocity Vector Based Attitude Information

This section discusses the display of velocity vector based attitude information to close the pilot's control loop. Section 3.2.1 gives a brief overview of current display formats containing velocity vector information and of research conducted previously in this field. Section 3.2.2 presents a novel pseudo-attitude display that is fed by velocity vector based

pseudo-attitude, described in the previous section. This display was used throughout the research described in this thesis. Section 3.2.3 discusses display related issues such as display update rate and latency. Finally, Section 3.2.4 reviews the results of a preliminary simulator study of the pseudo-attitude display.

3.2.1 Introduction

The display of velocity vector information on a primary flight display has found broad acceptance in the military aircraft fleet. Typically, velocity vector information is displayed in the form of a flight path or velocity vector symbol. Together with additional essential information, the velocity vector symbol is projected on the aircraft windscreen using a head-up display (HUD), thus allowing the pilot to see important aircraft data overlaid with the actual visual scene. Figure 3.6 shows a typical head-up display projection. Shown are the flight path vector symbol, traditional attitude, heading, velocity and altitude information. The flight path vector symbol moves freely over the head-up display area indicating the instantaneous direction of travel relative to the aircraft attitude symbol, that is, relative to the aircraft's current azimuth and elevation. The aircraft's inertial acceleration is inferred indirectly by the rate of change of the flight path vector symbol's position on the screen

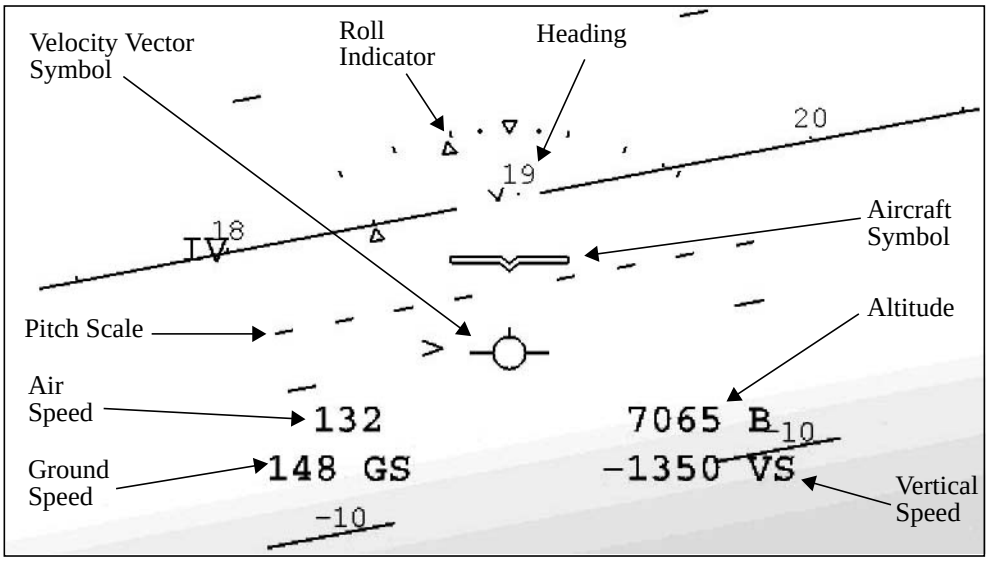

Figure 3.6: Head-up Display (HUD)

Currently, only few commercial aircraft are equipped with HUDs. Most of today's commercial aircraft, however, provide horizontal velocity vector information on navigation displays that can operate in the ground track mode. The information displayed includes ground track heading and ground speed along with land marks and way points, and allows the aircraft to be controlled along an earth referenced path. A typical navigation display is the Electronic Horizontal Situation Indicator (EHSI) shown in Figure 3.7.

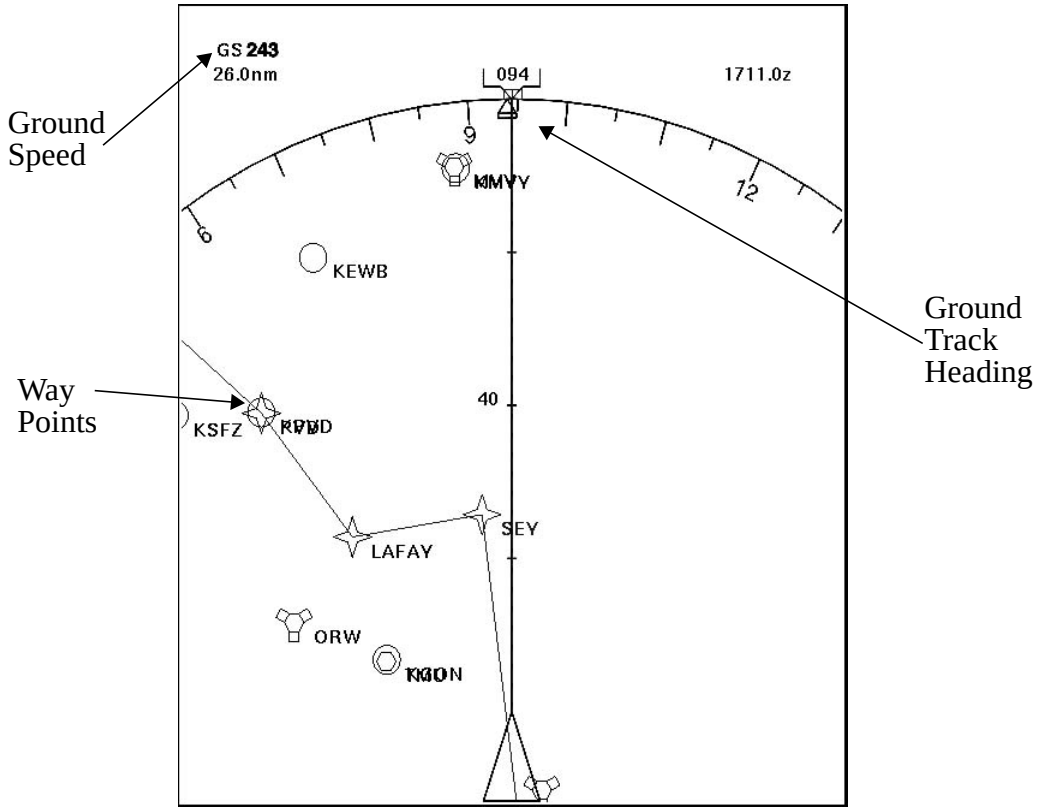

Figure 3.7: Electronic Horizontal Situation Indicator (EHSI)

Steinmetz (1986) suggested a velocity vector aligned attitude indicator. On the display, the velocity vector symbol is fixed in the center, indicating the direction of travel, while the aircraft attitude symbol is moving freely, indicating the direction in which the aircraft is pointing relative to the velocity vector. Figure 3.8 shows the velocity vector aligned attitude indicator for the case of an aircraft having a left crab angle due to a crosswind. The vertical motion of the horizon is coupled to the flight path angle and the pitch angle is given by the position of the aircraft attitude symbol with respect to the horizon. This is in

contrast to conventional attitude indicators which are centered around the direction the aircraft is pointing and in which the vertical motion of the horizon is coupled to the pitch angle. The pilot has then to extract information about a flight path angle change from the relative motion between the horizon and the flight path vector symbol.

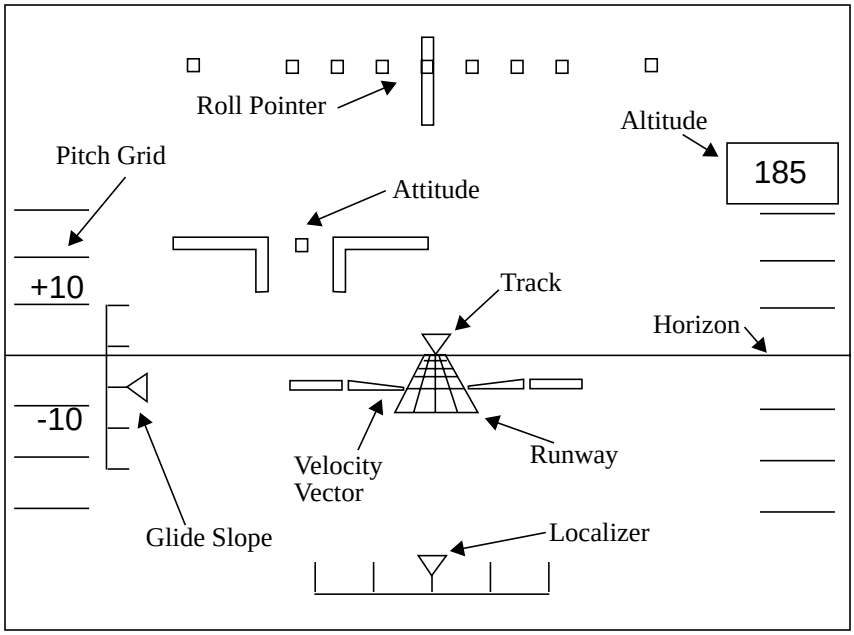

Figure 3.8: Velocity Vector Aligned Attitude Indicator (Steinmetz 1986)[a]

a. In the presence of left-to-right crosswind.

Steinmetz (1980) describes a simulator study and actual flight tests comparing pilot performance and pilot opinion for a velocity vector aligned attitude indicator and a conventional attitude indicator. The velocity vector aligned attitude indicator was tested in conjunction with a velocity vector control wheel steering mode developed by Lambregts (1979). The velocity vector information to drive the display was obtained from an onboard inertial navigation system (INS). He reports that although the statistical results of the objective performance measures for the two displays are inconclusive, the pilots indicated a clear preference for the velocity vector aligned display format.

More recently, Theunissen (1997) compared velocity vector aligned perspective flight path displays with attitude aligned perspective flight path displays in a simulator study. The perspective flight path display showed the desired flight trajectory in the form of a

tunnel-in-the-sky that was overlaid on an attitude indicator referenced to either the velocity vector symbol or aircraft attitude symbol, with the other one moving freely. He reports that the pilots were able to stay well within tunnel for both displays and that the statistical analysis did not reveal a significant difference between them. Similar to the results reported by Steinmetz, the pilot subjects in their subjective rating unanimously preferred the velocity vector aligned frame of reference.

3.2.2 Pseudo-Attitude Display

The displays presented previously show velocity vector information in conjunction with traditional aircraft attitude information as conveyed by an artificial horizon. In this section a novel pseudo-attitude display is presented. This display shows pseudo-attitude which is *entirely* based on velocity vector information. The pseudo-attitude display acts as a substitute for the traditional attitude indicator.

Pseudo-attitude is displayed in a manner quite similar to traditional attitude. The distinguishing feature of the pseudo-attitude display is that its aircraft symbol is referenced to the inertial velocity vector. Figure 3.9 shows the comparison between traditional and pseudo-attitude displays. The left display shows the traditional artificial horizon which has its aircraft symbol referenced to the direction the aircraft is pointing, that is, referenced to the aircraft fuselage centerline. In the case shown in Figure 3.9, the aircraft is in level flight and has a pitch angle of 5 deg and a roll angle of 20 deg.

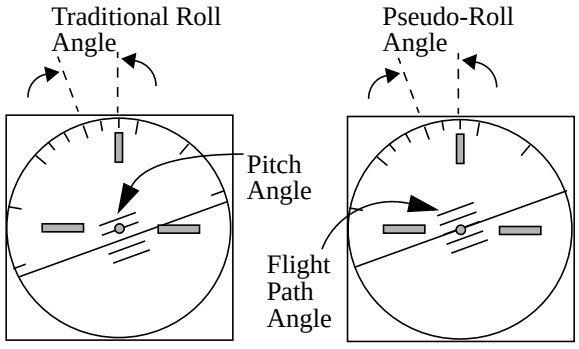

Figure 3.9: Attitude and Pseudo-Attitude Display

The right display shows the pseudo-attitude of the aircraft in the same flight condition. The roll representation indicates the pseudo-roll angle which in the case shown is very close to the traditional roll angle. The pitch information is replaced with the flight path angle and the horizon is, thus, coupled to the flight path angle. As shown in Figure 3.9, because the aircraft is in a level turn the aircraft symbol lies on the horizon. For straight and level flight, the difference between pitch angle and flight path angle indication is due to the angle of attack.

A potential issue of the pseudo-attitude display may be its similarity to the traditional attitude display. This may cause pilots, who are used to the traditional attitude indicator, some confusion as to whether pitch or flight path angle is being displayed.

3.2.3 Display Update Rate and Latency

Two important display parameters for control loop closure are display update rate and display latency. The display update rate is related to the sampling frequency of the control loop. The sampling frequency must be sufficient to ensure stability and controllability of the system. The display update rate may have to meet additional perceptual requirements of the pilot. Display latency is the time delay between the aircraft response and the corresponding response of the cockpit display. It has an impact on the achievable bandwidth as well as the handling qualities of the system. Display update rate and latency are discussed next in more detail.

Control Loop Sampling Frequency and Display Update Rate

The choice of a sampling frequency in a digital control system, in general, involves an important trade-off. A higher sampling frequency allows for fast reaction times and good control performance at the expense of higher demands on the digital control electronics. A lower sampling frequency, on the other hand, implies more relaxed control systems demands but may incur lower control performance. Therefore, the sampling frequency and display update rate are typically chosen to be the lowest frequency which satisfies stability, controllability and perceptual requirements.

It is generally difficult to determine stability boundaries of a controlled system as a function of the sampling interval because of the transcendental nature in which the eigenvalues are dependent on the sampling interval (Mansour 1990). With the exception of simple dynamic systems where stability boundaries can be determined in closed form, the stability of a sampled system must be evaluated for a particular sampling frequency.

A theoretical lower bound on the sampling frequency is given by Shannon's sampling theorem which requires the sampling frequency to be twice the closed-loop bandwidth. This criterion, however does not ensure that adequate controllability is attained. Rather, empirical rules have been developed for the choice of the sampling frequency that gives acceptable controllability. Ackermann (1983) and Mansour (1990) suggest a as a 'rule of thumb' a sampling frequency of 10 times the largest eigenfrequency of the continuous system. Similarly, Powell and Katz (1975) suggest a sampling frequency of 10 to 20 times the closed-loop bandwidth. For the example of a Cessna 182 which has its largest eigenfrequency in the 3 - 4 rad/sec (0.5 -0.65 Hz) range, Ackermann's rule suggests a display update rate of approximately 6 Hz.

An additional requirement on the display update rate is given by the human perception of the display. To convey the illusion of continuous motion, the display update rate must exceed approximately 10 Hz (Theunissen 1997). Higher update rates yield a more smoothly animated picture. However, while the minimum update rates follow from stability and control considerations, the perceptual requirements may be met using interpolation or extrapolation at the required frequency, but do not necessarily require higher observation rates.

Display Latency

Display latency is the time delay between aircraft response and the corresponding response of the cockpit display. It may be composed of time delays introduced by different components and functions of the aircraft control loop.

Sampling delay is introduced by the very nature of a sampled system†. Next, the measurement and processing of data may cause additional delay. For example, the velocity measurements of the Novatel 3151R GPS receiver is based on the average of two successive phase measurements obtained from the carrier phase-locked loop. Thus, for a receiver update rate of the 10 Hz, the latency introduced is 50 ms.

Also, the filtering of data introduces latency. For instance, perceptual requirements can dictate a certain amount of filtering to obtain a smoothly animated picture. This is the case for the display of GPS-based pseudo-attitude. Filtering has to be applied to reduce the noise in the acceleration information that is inferred from velocity measurements by differentiation.

Time delays reduce closed-loop system stability by reducing available phase margin. Consequently, achievable gain and bandwidth are reduced as well. Latency in the control system affects the flying qualities by requiring the pilot to compensate for it, thereby increasing task demand load and degrading task performance. Extensive research on the relation of latency and flying qualities has been conducted. Some of the research is documented in (McRuer 1974), (King 1993) and (Stengel 1980).

There exists a trade-off between stability and perceptual requirements for the display of GPS-based pseudo-attitude. The former require a minimum of filtering to reduce the latency, whereas the latter call for enough filtering to obtain a smooth display indication. To gain more insight into this trade-off, control-loop analysis, simulations and flight tests were conducted. These are discussed in later chapters.

3.2.4 Preliminary Simulator Study of the Pseudo-Attitude Display and the Required Display Update Rate

A preliminary simulator study of the pseudo-attitude display was conducted on the MIT Advanced Cockpit Simulator to evaluate its efficacy and to determine the minimum display update rate necessary for pilots to maintain adequate aircraft control.‡

† A zero-order-hold (ZOH), using a second-order approximation, can be expressed as a pure time delay with a dead time of half the sampling interval (Stengel 1980).

The simulator was based on the dynamic model of a Cessna 182 (C182). The latter resembles in dimension and weight a Piper Arrow aircraft† which was the aircraft type used for flight tests described in subsequent chapters. The investigation included the objective measurement of pilot performance and the subjective assessment of pilot preference and was based on 8 pilot subjects.

The results of the study suggested a similar objective performance and subjective pilot preference for both the traditional attitude and the pseudo-attitude display. In regard to the update rate, the objective results suggested a minimum display update rate of 2 Hz. Subjective results, however, showed a clear pilot preference for a 6 Hz display update rate. The latter is also in good correspondence with Ackermann's rule of thumb for the choice of the sampling rate for adequate controllability, as stated in Section 3.2.3. The rule calls for a sampling rate of 10 times the highest frequency of the aircraft, which in the case of the C182 is approximately in the 3 - 4 rad/sec range (0.5 - 0.65 Hz).

3.3 Chapter Summary

This chapter presented the synthesis and display of velocity based attitude information. In particular, it discussed the synthesis of pseudo-attitude from velocity and acceleration information. The synthesis of traditional attitude and attitude in wind axes was outlined and the differences and commonalities to pseudo-attitude explored.

In addition, this chapter introduced a novel pseudo-attitude display that allows the human to close the attitude loop around velocity based attitude information. Issues such as display update rate and display latency were investigated.

---

‡ The simulator study was performed by Henderson (1996, 1997) in close collaboration with the author of this thesis. This research work was part of an undergraduate summer project jointly supervised by the author of this thesis and Dr. R. John Hansman.The study is included in abbreviated form in Appendix B.

† The aerodynamic model of the Piper Arrow is unavailable in the public domain.

Chapter 4

GPS-Based Velocity and Acceleration

This chapter discusses the generation of velocity and acceleration information using the Global Positioning System (GPS). This information is used to synthesize pseudo-attitude, as described in the previous chapter. In particular, the emphasis is put on the generation of velocity and acceleration information using a single-antenna GPS receiver in stand-alone (i.e. non-differential) mode. Differential GPS operation is considered only when relevant to the current discussion.

GPS was conceived and implemented in the late seventies and eighties and achieved its operational status in the first half of the nineties. Originally designed as a military navigation system, GPS is used today in a vast variety of applications, one of which is the subject of this thesis. In the past decade a vast amount of research has been conducted on a broad variety of GPS related aspects and current research efforts in GPS continue unabated. Much of the material discussed in this chapter is therefore based on research work performed by others over the years. An attempt is made to give credit where appropriate.

It is beyond the scope of this chapter to give a detailed overview of the Global Positioning System. For a detailed discussion the reader is referred to Parkinson (1996a) or Kaplan (1996). The treatment in this chapter assumes that the reader has a background on the basic principles of GPS.

The objectives of this chapter are twofold. The first objective is to give the reader a basic understanding of the mechanisms of GPS-based velocity and acceleration information generation. This includes a description of the GPS observables and the method of measuring and processing them to yield velocity and acceleration information. The second objective is to highlight the issues, trade-offs and limitations in the signal

processing architecture, in the receiver design, and in the GPS system as a whole, which *pertain* to the synthesis of GPS velocity based flight controls. These two objectives are not treated separately, but are interwoven throughout the chapter.

Section 4.1 briefly discusses the relevant principles of operation and the GPS observables. Section 4.2 concentrates on the GPS receiver architecture and identifies the receiver functions which are relevant to the measurement of GPS carrier Doppler shift information. In Section 4.3, the generation of velocity and acceleration information from carrier Doppler shift measurements is discussed. Section 4.4 investigates the bandwidth of GPS velocity and acceleration and discusses the fundamental trade-off between bandwidth and noise reduction. Section 4.5 presents a number of GPS error sources and investigates their impact on GPS velocity and acceleration measurements and, subsequently, on the synthesis of pseudo-attitude. Section 4.6 identifies the limitations of GPS due to integrity, availability and continuity issues. Their implications on the use of pseudo-attitude will be examined more fully in Chapter 10. Finally, Section 4.7 provides a chapter summary.

It should be noted that details of a particular GPS receiver architecture are generally of proprietary nature and are thus not available in the open literature. Hence, the material presented in Section 4.2 uses a generic receiver model and discusses its operation on a conceptual level. References to the architecture of the Novatel 3151R GPS receiver -the receiver used throughout the flight tests documented in this thesis- are made whenever the appropriate information is available.

4.1 Principle of Operation and GPS Observables

GPS is a satellite based navigation system designed to provide a properly equipped user with position and velocity information anywhere on the globe. The GPS space segment consists of 24 satellites (21 primary and 3 spare satellites) arranged in six orbital planes. Each orbital plane contains 4 satellites in 12-hour orbits. This constellation provides a user located anywhere on the world with the visibility of four or more satellites

at any time. This is necessary for continuous positioning capability, as will be shown below. Since the satellites are not in geosynchronous orbits the satellite geometry, as seen by an observer at a fixed location on earth, is continuously changing.

GPS satellites transmit signals using two carrier frequencies, L1 (1575.42 MHz) and L2 (1227.6 MHz). The carrier frequencies are modulated by various spread spectrum signals which contain information necessary to determine position and velocity. The basic principles of position and velocity determination using these signals are explained next. The treatment of position determination is included because the determination of velocity requires the knowledge of the user position beforehand.

4.1.1 Principle of Position Determination

The basic function of GPS positioning is the determination of the user position $\mathbf{r}_u$ from GPS signals. The GPS receiver accomplishes this by using the propagation delay of GPS signals from an array of satellites to the receiver's GPS antenna, i.e. the time interval for the signals to travel from the satellites to the receiver. The time references of the different GPS satellites are precisely synchronized and known (GPS systems time). If differences between satellite clocks occur they can be quantified exactly and modulated onto the carrier as navigation messages. The GPS receiver time, on the other hand, relies on an autonomous clock within the receiver which, due to inherent clock uncertainties, will be offset from the GPS system time by an unknown bias b_u. The observed propagation delay t_{pd} of the GPS signal from each satellite is, thus, offset from the actual signal travel time by the same bias b_u.

The observed propagation delay t_{pd} from each satellite scaled by the speed of light in vacuum c, corresponds to a range measurement and is commonly referred to as *pseudo-range* ρ, i.e.

$$\rho = c \cdot t_{pd} \tag{4.1}$$

The term *pseudo-range* is utilized because the range measurement contains the geometric range from the satellite to the receiver as well as a range component due to the receiver clock offset. The pseudo-range observation ρ_i between a user and satellite *i* can be related to the user position and clock offset by

$$\rho_i = \left|\mathbf{r}_i - \mathbf{r}_u\right| + c \cdot b_u + \varepsilon_i \tag{4.2}$$

where $\mathbf{r}_i$ is the three dimensional satellite position at signal transmit time, $\mathbf{r}_u$ is the receiver position at the receive time, and ε_i is the composite of various error sources including atmospheric delays, Selective Availability (SA) and receiver noise.

Eq. (4.2) shows that the GPS positioning problem is four dimensional consisting of the simultaneous determination of the three dimensional user position $\mathbf{r}_u$ and, as the fourth dimension, the receiver clock offset b_u or, equivalently, GPS system time. This problem can be solved, in principle, from pseudo-range measurement observations from at least four satellites. The actual solution will be presented in Section 4.3. The underlying coordinate system currently used by GPS is the 1984 World Geodetic System (WGS-84), but the user position can be expressed in any desired coordinate system.

4.1.2 Principle of Velocity Determination

Three dimensional user velocity $\mathbf{v}_u$ can be determined from the observed frequency shift of the received GPS carrier signal. The observed carrier frequency differs from the nominal L1 or L2 carrier frequency due to Doppler shifts caused by the relative motion of the satellite with respect to the user, as well as a receiver clock frequency bias f_u.†

The Doppler shift caused by the relative motion of satellite *i* and user is given by the projection of the relative velocity onto the line of sight, scaled by the ratio of the transmitted carrier frequency L_1 to the speed of light c, i.e.

$$\Delta f_{D_i} = -\left(\frac{\mathbf{v}_i - \mathbf{v}_u}{c} \cdot \mathbf{1}_i\right) \cdot L_1 \tag{4.3}$$

† Satellite clock frequency errors are known and transmitted as navigation messages modulated onto the carrier.

where $\mathbf{v}_i$ is the velocity of satellite *i*, and $\mathbf{1}_i$ is the line of sight vector from the user to satellite *i*. The satellite velocity vector $\mathbf{v}_i$ can be computed in the receiver using ephemeris information modulated as a navigation message onto the carrier signal. The observed carrier frequency shift Δf_i can be related to the user velocity and unknown receiver clock frequency bias f_u as

$$\Delta f_i = \Delta f_{D_i} + f_u + \eta_i \tag{4.4}$$

where η_i is the composite of various errors in the frequency shift observation including the effects of Selective Availability and receiver noise. Eq. (4.4) can be converted to a *pseudo-range rate* observation by scaling it by the ratio c/L_1 and substituting Eq. (4.3) into Eq. (4.4), i.e.

$$\begin{aligned} \frac{c}{L_1} \cdot \Delta f_i &= \frac{c}{L_1} \cdot (\Delta f_{D_i} + f_u + \eta_i) \\ &= \delta_i = (\mathbf{v}_i - \mathbf{v}_u) \cdot \mathbf{1}_i + \delta_u + \zeta_i \end{aligned} \tag{4.5}$$

where δ_i is referred to as *pseudo-range rate* or *delta range* from the user to satellite *i*, δ_u is the receiver clock frequency bias, and ζ_i is the error in the observation, all in m/s.

Similar to the GPS positioning problem presented in Section 4.1.1, the user velocity determination is a four dimensional problem. It must be solved for the three user velocity components and the receiver clock frequency bias simultaneously. Observations from four or more satellites are therefore necessary to solve the problem, at least in principle. Because the line of sight vector has to be known, the velocity is generally computed in conjunction with or just after the position of the user is determined. The actual velocity information generation is presented in Section 4.3.

4.2 GPS Receiver Architecture and Measurement Generation

This section discusses the architecture and operation of a generic GPS receiver and highlights the issues which are relevant to the generation of velocity and acceleration information and to the synthesis of pseudo-attitude. Section 4.2.1 discusses the GPS

receiver architecture and operation in general, and the carrier phase-locked loop (PLL) in particular. Section 4.2.2 then describes how delta range measurements are generated from PLL outputs. An overview of the GPS signal structure, which is inherently connected to the GPS receiver architecture, is given in Appendix C.1.

4.2.1 GPS Receiver Architecture and Operation

For the following treatment the receiver is assumed to be tracking the satellites in view with the tracking loops operating in a coherently correlating manner.

Top Level Receiver Architecture

A high-level block diagram of a generic GPS receiver is shown in Figure 4.1 (Zhuang 1996, Ward 1996). The blocks relevant to the generation of velocity and acceleration information are emphasized in this figure.

GPS RF signals of all satellites in view are received by an antenna with nearly hemispherical gain coverage. The received signals are amplified by a low-noise pre-amplifier which sets the noise figure of the receiver. The signals are then passband filtered to minimize out-of-band noise, and subsequently down-converted in two stages to intermediate frequency (IF) and baseband. The signal conditioning preserves the signal modulations and Doppler shifts, but lowers the carrier frequency to a baseband frequency where an analog-to-digital conversion takes place.

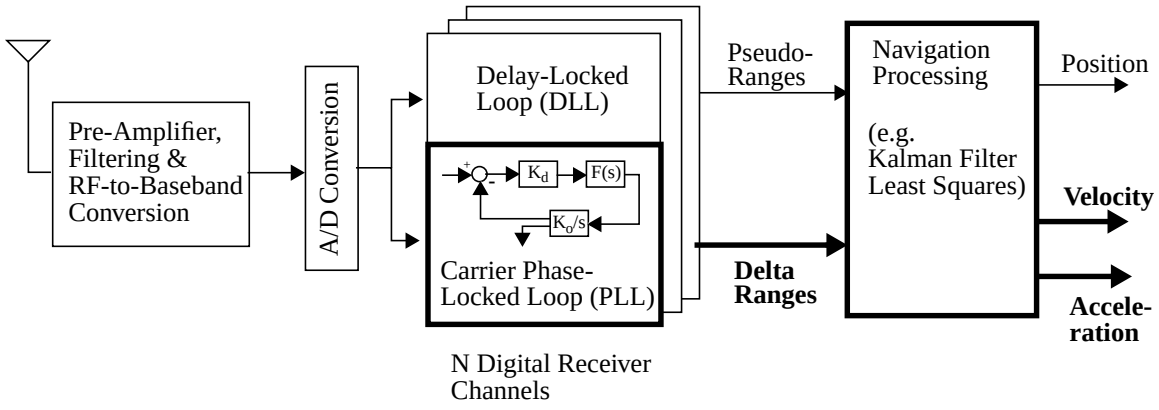

Figure 4.1: Generic Receiver Architecture

The digitized baseband signals are processed by each of the N digital receiver channels. The number of receiver channels corresponds to the number of satellites which the receiver can track. The functions of the receiver channels are commonly performed by a digital baseband processor (DBP). Each channel contains a delay lock loop (DLL) that tracks the incoming C/A code of a particular satellite and generates pseudo-range measurements, and a phase-locked loop (PLL) which tracks the incoming carrier phase and measures Doppler shifts of that satellite. In addition, the receiver channels demodulate the navigation data from the baseband signals.

The DLL measures pseudo-ranges to a particular satellite by tracking the incoming Coarse/Acquisition (C/A) code and aligning it with a receiver generated replica of that satellite's code in a correlator. If the replica code tracks the incoming code exactly, maximum correlation is achieved. Any misalignment of the replica code phase with respect to the incoming C/A code phase decreases the correlator output and produces a difference signal so that the amount and the direction of the phase change can be corrected for by the DLL. These difference signals are converted to pseudo-ranges as explained in (Ward 1996). Since pseudo-range measurements only play a secondary role† in the determination of velocity and acceleration measurements, the DLL will not be considered in more detail here.

The carrier phase-locked loop generates delta range measurements by tracking the incoming carrier phase and associated Doppler shift. The PLL design is inherently connected to the achievable performance of the delta range measurements. The following sections explain the function of the PLL and the generation of delta range measurements in greater detail.

The set of pseudo-range and delta range measurements generated by the receiver channels is fed into the navigation processor where a navigation solution is computed, typically at an output rate between 1 and 10 Hz. The navigation processor calculates estimates of three dimensional position, velocity and sometimes acceleration as well as the

† Pseudo-range measurements determine the user position and thus the line-of-sight vector from the user to the satellite. This vector is used in Eq. (4.5).

receiver clock biases, as outlined in Section 4.3 in more detail. In stand alone GPS receivers, the navigation process is commonly performed in a single high-speed processor that also performs the other baseband functions.

The Novatel 3151R GPS receiver, used for the flight tests described in this thesis, has 12 digital channels and can provide a navigation solution for position and velocity at an update rate of 10 Hz.

Carrier Phase-Locked Loop

The carrier phase-locked loop (PLL) tracks the incoming carrier phase signal and extract velocity information in form of measured carrier Doppler shifts. Figure 4.2 shows a generic linearized phase-locked loop typically used in GPS receivers. The carrier PLL tracks the incoming signal Θ_i by eliminating any phase difference Θ_e, between the incoming carrier signal and the carrier replica generated by a local voltage controlled oscillator (VCO), in a closed loop manner. Any Doppler shift due to satellite or user motion will cause the phase of the incoming signal to be advanced or delayed with respect to the phase of the local carrier. The difference is measured by a carrier phase discriminator and is fed into the VCO to bring the local carrier replica in phase alignment with the incoming carrier phase. To reduce the effects of thermal noise, the phase difference is low-pass filtered and the filter output is fed into the VCO. The phase and Doppler frequency shift of the incoming signal can then be obtained from those of the local VCO output. A more detailed description of the PLL operation is given in Appendix C.2 and in (Zhuang 1996).

The PLL design determines the noise characteristics (thermal noise error) of the measurements and the ability of the receiver to track GPS signals under dynamic conditions (dynamic performance) (Ward 1996). Both are fundamental for the generation of pseudo-attitude from GPS velocity and acceleration measurements.

The amount of thermal noise error in the phase measurements is directly related to the amount of filtering necessary to obtain smooth velocity and acceleration information and to the latency introduced thereby.† The dynamic performance, on the other hand, has to be

sufficient to allow the receiver to track GPS signals even under dynamic conditions such as in steep aircraft turns or in fast maneuvers. Dynamic stress errors, i.e. errors in tracking the incoming carrier signal, as well as the occurrence of complete loss of phase-lock have thereby to be minimized.

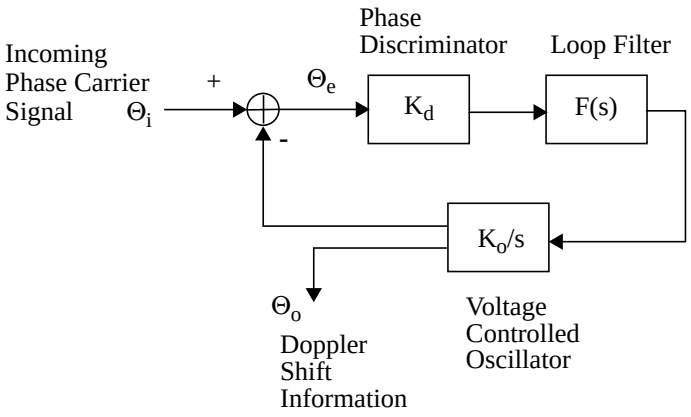

Figure 4.2: Generic Linearized Third-Order PLL

Noise error and dynamic performance are concurrent characteristics and result in a design trade-off. This trade-off is briefly outlined here with a more detailed description given in (Ward 1996). The carrier phase-locked loop is completely characterized by the choice of the carrier loop filter (order and bandwidth), the carrier loop discriminator, and the predetection integration time (correlation time)†.

For the carrier Doppler phase measurements to have low noise, the carrier loop filter bandwidth should be narrow, the discriminator should be a phase discriminator, and the predetection integration time should be long. On the other hand, to minimize the dynamic stress error and to increase the dynamic range where the signal is in lock, the carrier loop

† In Section 3.2.3, the trade-off between between stability of the control loop and perceptual requirements for the display of GPS-based pseudo-attitude was outlined. The former require a minimum of filtering to reduce the latency, whereas the latter call for enough filtering to obtain a smooth display indication.

† i.e. the phase signal is integrated (averaged) before it is fed into the carrier phase discriminator. See Appendix C.2.

filter bandwidth should be wide, the discriminator should be a differential phase discriminator (i.e. frequency-locked loop), and the predetection integration time should be short (Ward 1996).

In order to not loose phase-lock in dynamic flight conditions, the PLL must be at least of *third-order*. This is explained using the linearized PLL in Figure 4.2. For clarity, the treatment is performed in the continuous time (or s-) domain. However, the results are easily transferable to the discrete-time domain. In the linearized model, the voltage controlled oscillator (VCO) is modeled as an integrator. Using a second-order low-pass filter of the form

$$F(s) = \frac{2\omega_o s^2 + 2\omega_o^2 s + \omega_o^3}{s^3} \tag{4.6}$$

as a loop filter F(s) results, in conjunction with the integrator of the VCO, in a third-order PLL (Ward 1996). For the simplified case of unity gains, i.e. $K_d = K_o = 1$, the transfer function H(s) and the error transfer function $H_e(s)$ are given by

$$H(s) = \frac{\Theta_o}{\Theta_i} = \frac{2\omega_o s^2 + 2\omega_o^2 s + \omega_o^3}{s^3 + 2\omega_o s^2 + 2\omega_o^2 s + \omega_o^3} \tag{4.7}$$

$$H_e(s) = \frac{\Theta_e}{\Theta_i} = \frac{s^3}{s^3 + 2\omega_o s^2 + 2\omega_o^2 s + \omega_o^3} \tag{4.8}$$

Using the final value theorem, the steady-state phase error $\Theta_e(t \to \infty)$ can be determined for different phase inputs. Table 4.1 shows the phase error of the third-order PLL for phase inputs Θ_i corresponding to line-of-sight velocity, acceleration and jerk steps. R denotes the line-of-sight range to the satellite, Θ_i is the phase input, ω_i is the corresponding frequency input, ω_o is the loop filter natural frequency, and $k = L_1/c$ is a scaling constant.

As can be seen, the third-order loop tracks constant line-of-sight velocity and acceleration with no error, but is sensitive to jerk stress. Thus, in order to track GPS signals during vehicle acceleration a third-order or higher-order PLL must be used. A line-

Line-of-Sight Dynamics	**Frequency Input ω_i**	**Phase Input Θ_i**	**Phase Error $\Theta_e(t \to \infty)$**
Range Velocity Step (dR/dt)	$k \cdot \frac{dR/dt}{s}$	$k \cdot \frac{dR/dt}{s^2}$	0
Range Acceleration Step (d^2R/dt^2)	$k \cdot \frac{d^2R/dt^2}{s^2}$	$k \cdot \frac{d^2R/dt^2}{s^3}$	0
Range Jerk Step (d^3R/dt^3)	$k \cdot \frac{d^3R/dt^3}{s^3}$	$k \cdot \frac{d^3R/dt^3}{s^4}$	$k \cdot \frac{d^3R/dt^3}{\omega_o^3}$

Table 4.1: Phase Error due to Line-of-Sight Dynamics[a]

a. R denotes the line-of-sight range to the satellite, Θ_i is the phase input, ω_i is the corresponding frequency input, ω_o is the loop filter natural frequency, and $k = 2\pi L_1/c$ is a scaling constant.

of-sight jerk input, that is, a varying line-of-sight acceleration, induces a *dynamic stress error*. The steady-state error due to a jerk step input is inversely proportional to ω_o^3, and therefore a function of the tracking loop bandwidth, and directly proportional to the line-of-sight jerk input d^3R/dt^3. The peak dynamic stress error may be slightly larger than the steady-state error if the loop filter response to a step function has overshoot.

The Novatel GPS 3151R receiver has third-order PLL, a carrier tracking bandwidth of 15 Hz and a predetection integration time of 100 ms. The jerk is limited to 4.5 g/s in order to not loose lock (Novatel 1995, 1996).

There is a trade-off between PLL loop order and loop stability (Zhuang 1996). Higher order loops have better dynamic tracking performance, i.e. smaller tracking error and faster convergence, but are less stable than lower order loops. Thus, care has to be taken that the loop parameters selected ensure loop stability.

In practice, a well-designed GPS receiver will incorporate a flexible carrier phase-locked loop design which adapts its PLL characteristics according to the anticipated dynamics. This is possible since most of the functions determining the loop characteristics are programmed in the receiver signal processing software. The Novatel 3151R, for example, allows the carrier tracking loop bandwidth to be adjusted from 2.5 Hz for stationary use up to 15 Hz for high dynamics.

4.2.2 Delta Range Measurement Generation

Delta range measurements are obtained from the change in carrier phase during a specified time interval. In principle, delta range measurements could be derived from differential measurement of the *code* tracking loop. The measurements, however, are nearly *three orders of magnitude* less noisy if taken from the *carrier* phase-locked loop (Ward 1996).

The carrier Doppler phase measurements are extracted by the receiver baseband process from the carrier tracking loop using a carrier accumulator. The carrier accumulator consists of integer cycle count N_{CA} and the fractional cycle count Φ_{CA} of the carrier Doppler phase measurements. The delta range δ_i is simply the change in phase in the carrier accumulator during a specified time interval T_u averaged over this interval and scaled with the wavelength λ_1 of the L1 carrier frequency. That is

$$\delta_i = [(N_{CA_i}(t) + \Phi_{CA_i}(t)) - (N_{CA_i}(t - T_u) + \Phi_{CA_i}(t - T_u))] \cdot \frac{\lambda_{L1}}{T_u} \tag{4.9}$$

Since delta range is an average velocity based on the time difference T_u between successive carrier phase measurements and not an instantaneous Doppler velocity measurement, it has a latency of $T_u/2$ seconds. For the case of the Novatel 3151R receiver providing 10 Hz data the latency is 50 ms.

4.3 Velocity and Acceleration Generation

The previous sections described the fundamental principles of GPS position and velocity determination, the GPS observables, namely pseudo-range and delta range, and how GPS receivers extract them from the incoming GPS signals. This section explains how velocity and acceleration information is generated from GPS observables.

The physical measurements of pseudo-range and delta range from at least four satellites and the measurement equations Eq. (4.2) and Eq. (4.5) described in the previous sections are in principle sufficient to generate a single-point solution of the position, velocity and receiver clock bias states (Axelrad 1996). In this approach, the measurement

equations are linearized about some nominal value, e.g. about the current best estimate. A least squares solution to the linearized measurement equations and a set of four or more pseudo-range and delta range measurements, taken at the same time, gives the corrections to the current best estimate, thus, yielding an improved state estimation. This approach is used for the Novatel 3151R GPS receiver and is discussed in Appendix C.3.

A disadvantage of the single-point solution approach is that it does not include any of the known user dynamics in the estimation of the states and it does not carry any information from one measurement cycle to the next.

A Kalman filter overcomes this shortcoming by optimally incorporating the a priori knowledge of the user dynamics in the estimation of the desired states (Axelrad 1996). The Kalman filter consists of two models. As in the single-point solution, it uses a linearized measurement model to describe the connection of the measurements to the states to be estimated. In addition, it uses a linear (or linearized) process model to describe the dynamics of the states and, thus, incorporates a priori knowledge of the system.

Section 4.3.1 discusses the measurement and process model for a discrete extended Kalman filter (EKF) that provides position, velocity and acceleration information from pseudo-range and delta range measurements. The velocity and acceleration information are then available as the input to the pseudo-attitude synthesis.

Many commercially available GPS receivers, including the Novatel 3151R receiver, do not contain acceleration states in their navigation solution. In these instances, acceleration can be obtained by two approaches. In the first approach, the filter discussed in Section 4.3.1 may be implemented on an external CPU and fed by raw pseudo-range and delta range output of the GPS receiver to produce a navigation solution which includes the desired GPS acceleration states. Although simple in principle, the actual implementation may be quite elaborate due to the large amount of additional computation necessary for tasks such as the calculation of satellite orbits and error corrections transmitted in the navigation message.

In Section 4.3.2, a second, simpler approach to obtain acceleration is discussed. It infers acceleration information directly from the GPS receiver velocity output and, thus bypasses the large amount of computation arising in the first approach. Because of the noise present in the receiver measurements, simple back-differencing of the velocity output is not feasible. Instead, it relies on an external Kalman filter to estimate acceleration. This approach was implemented for the experimental work documented in this thesis in Chapters 6 - 9.

4.3.1 Receiver Internal Kalman Filter

In stand-alone GPS applications, generally, a discrete, extended Kalman Filter (EKF) is used. Its characteristics are given by a process model and measurement model. These models are described next.

Process Model

The process model describes the dynamics of the states to be estimated by the filter and summarizes what is *a priori* known about the process. In the current context, the process model contains the dynamic model of the GPS receiver platform motion and a model of the receiver clock errors.

Generally, for vehicles with high dynamics, the model describing the receiver platform motion includes states for position, velocity and acceleration in three dimensions (Axelrad 1996). This model causes the EKF to assume that jerk -the time derivative of acceleration- is a white noise process. While this assumption holds for most applications where acceleration information has to be estimated, it does not in the context of the synthesis of velocity vector based attitude information. As discussed in Section 3.1.4, the estimated accelerations relate directly to the aircraft roll angle. Consequently, the derivative of acceleration, jerk, corresponds to roll rate which is typically correlated over time. It is

therefore appropriate to include jerk as a separate state in the estimation and not to treat it as white noise (Kornfeld 1998a, b). Figure 4.3 shows the resulting process model. The

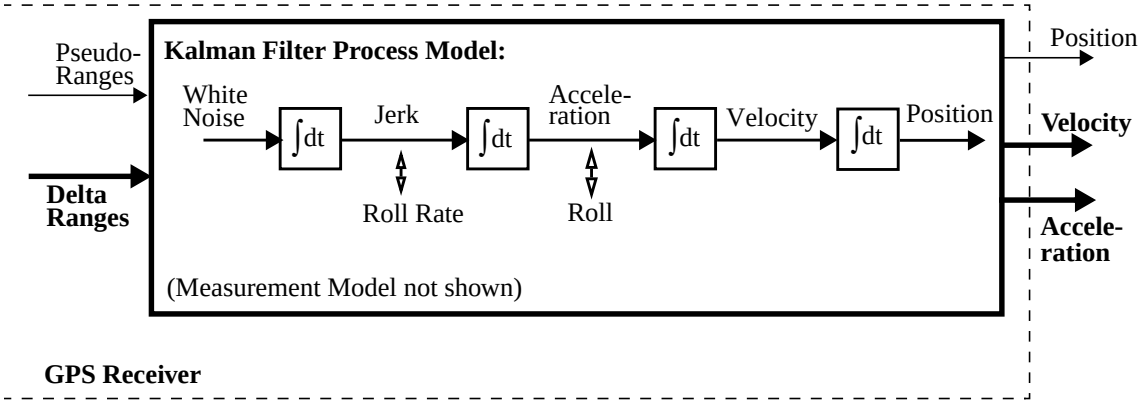

Figure 4.3: Kalman Filter Process Model

discrete-time dynamic model for the receiver platform motion can be expressed as

$$\mathbf{x}_p(k) = \Phi_p(\Delta t)\,\mathbf{x}_p(k-1) + \mathbf{w}_p(k-1) \tag{4.10}$$

The state vector $\mathbf{x}_p$ and the state transition matrix Φ_p are defined as

$$\mathbf{x}_p = [\mathbf{r}_u \;\; \mathbf{v}_u \;\; \mathbf{a}_u \;\; \mathbf{j}_u]^T \tag{4.11}$$

$$\Phi_p(\Delta t) = \begin{bmatrix} \mathbf{I} & \mathbf{I}\,\Delta t & \mathbf{I}\,\Delta t^2/2 & \mathbf{I}\,\Delta t^3/6 \\ 0 & \mathbf{I} & \mathbf{I}\,\Delta t & \mathbf{I}\,\Delta t^2/2 \\ 0 & 0 & \mathbf{I} & \mathbf{I}\,\Delta t \\ 0 & 0 & 0 & \mathbf{I} \end{bmatrix} \tag{4.12}$$

where $\mathbf{r}_u$, $\mathbf{v}_u$, $\mathbf{a}_u$ and $\mathbf{j}_u$ are the three dimensional user position, velocity, acceleration and jerk vectors, in that order. $\mathbf{I}$ is a 3x3 identity matrix and Δt is the filter update time interval. $\mathbf{w}_p(k)$ is a vector that models the process noise, i.e. the effect of random dynamics between

measurement updates; its elements are white noise sequences. The corresponding process noise covariance matrix $\mathbf{Q}_p$ is determined from (Gelb 1974) as

$$\mathbf{Q}_p = \begin{bmatrix} \mathbf{S}\,\Delta t^7/252 & \mathbf{S}\,\Delta t^6/72 & \mathbf{S}\,\Delta t^5/30 & \mathbf{S}\,\Delta t^4/24 \\ \mathbf{S}\,\Delta t^6/72 & \mathbf{S}\,\Delta t^5/20 & \mathbf{S}\,\Delta t^4/8 & \mathbf{S}\,\Delta t^3/6 \\ \mathbf{S}\,\Delta t^5/30 & \mathbf{S}\,\Delta t^4/8 & \mathbf{S}\,\Delta t^3/3 & \mathbf{S}\,\Delta t^2/2 \\ \mathbf{S}\,\Delta t^4/24 & \mathbf{S}\,\Delta t^3/6 & \mathbf{S}\,\Delta t^2/2 & \mathbf{S}\,\Delta t \end{bmatrix} \tag{4.13}$$

$$\mathbf{S} = \begin{bmatrix} S_N & 0 & 0 \\ 0 & S_E & 0 \\ 0 & 0 & S_D \end{bmatrix}$$

where S_N, S_E, S_D are the white noise spectral amplitudes for the North, East and Down directions, respectively. The process noise covariance $\mathbf{Q}_p$ represents the uncertainty in the dynamic model. The determination of $\mathbf{S}$ is typically based on an estimate of the expected vehicle dynamics.

The model of the receiver clock errors consists of two states, the user clock bias b_u and the drift f_u introduced in Section 4.1. They represent the phase and frequency errors in the crystal oscillator in the receiver. In the model commonly employed both the frequency and phase are expected to random walk over a short period of time. Its discrete-time process is given by (Brown 1997) and is included here for completeness:

$$\mathbf{x}_c(k) = \Phi_c(\Delta t)\,\mathbf{x}_c(k-1) + \mathbf{w}_c(k-1) \tag{4.14}$$

where the state vector $\mathbf{x}_c$ and the state transition matrix Φ_c are defined as

$$\mathbf{x}_c = [c b_u \;\; \delta_u]^T \tag{4.15}$$

$$\Phi_c(\Delta t) = \begin{bmatrix} 1 & \Delta t \\ 0 & 1 \end{bmatrix} \tag{4.16}$$

The vector $\mathbf{w}_c(k)$ models the process noise and its elements are white noise sequences. The corresponding process noise covariance matrix $\mathbf{Q}_c$ is determined from the conventional Allen variance parameters that are often used to describe clock drift. Its derivation can be found in (Brown 1997).

The complete process model is obtained by combining the dynamic model for the receiver platform motion and the receiver clock model

$$\mathbf{x}(k) = \begin{bmatrix} \mathbf{x}_p(k) \\ \mathbf{x}_c(k) \end{bmatrix} = \begin{bmatrix} \Phi_p(\Delta t) & 0 \\ 0 & \Phi_c(\Delta t) \end{bmatrix} \begin{bmatrix} \mathbf{x}_p(k-1) \\ \mathbf{x}_c(k-1) \end{bmatrix} + \begin{bmatrix} \mathbf{w}_p(k-1) \\ \mathbf{w}_c(k-1) \end{bmatrix} \tag{4.17}$$

and the error covariance $\mathbf{Q}$ of the complete model is given by

$$\mathbf{Q} = \begin{bmatrix} \mathbf{Q}_p & 0 \\ 0 & \mathbf{Q}_c \end{bmatrix} \tag{4.18}$$

Measurement Model

The EKF relies on linear process and measurement models. While the process model discussed above is itself linear, the measurement equations of the pseudo-range and delta range observables are not, as seen in Section 4.1. To obtain a linear measurement model, Eq. (4.2) and Eq. (4.5) must be linearized and in relation to the states to be estimated. The linearization is performed about the current best estimate of the state

$$\hat{\mathbf{x}} = [\hat{\mathbf{r}}_u \;\; \hat{\mathbf{v}}_u \;\; \hat{\mathbf{a}}_u \;\; \hat{\mathbf{j}}_u \;\; c\hat{b}_u \;\; \hat{\delta}_u]^T \tag{4.19}$$

Hence, given an a priori best estimate of the states, predicted pseudo-ranges $\hat{\rho}_i$ and delta ranges $\hat{\delta}_i$ can be calculated as

$$\begin{aligned} \hat{\rho}_i &= \left|\mathbf{r}_i - \hat{\mathbf{r}}_u\right| + c \cdot \hat{b}_u + \hat{\varepsilon}_i \\ \hat{\delta}_i &= (\mathbf{v}_i - \hat{\mathbf{v}}_u) \cdot \hat{\mathbf{1}}_i + \hat{\delta}_u + \hat{\zeta}_i \end{aligned} \tag{4.20}$$

where $\mathbf{r}_i$ is the satellite position, $\hat{\mathbf{1}}_i$ is the estimated line of sight vector from the user to the satellite *i*, and $\hat{\varepsilon}_i$, $\hat{\zeta}_i$ are the estimates of the range error and range error rate, respectively.

The difference between the actual measurements for satellites 1..n, and the predicted values can then be modeled as linearly related to the error in the states (Axelrad 1996). That is,

$$\begin{bmatrix} \rho_1 \\ \vdots \\ \rho_n \\ \delta_1 \\ \vdots \\ \delta_n \end{bmatrix} - \begin{bmatrix} \hat{\rho}_1 \\ \vdots \\ \hat{\rho}_n \\ \hat{\delta}_1 \\ \vdots \\ \hat{\delta}_n \end{bmatrix} = \begin{bmatrix} -\mathbf{1}_1^T & 0 & 0 & 0 & 1 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ -\mathbf{1}_n^T & 0 & 0 & 0 & 1 & 0 \\ 0 & -\mathbf{1}_1^T & 0 & 0 & 0 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & -\mathbf{1}_n^T & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta\mathbf{r}_u \\ \Delta\mathbf{v}_u \\ \Delta\mathbf{a}_u \\ \Delta\mathbf{j}_u \\ c\,\Delta b_u \\ \Delta\delta_u \end{bmatrix} + \begin{bmatrix} \Delta\varepsilon_1 \\ \vdots \\ \Delta\varepsilon_n \\ \Delta\zeta_1 \\ \vdots \\ \Delta\zeta_n \end{bmatrix} \tag{4.21}$$

or in more compact form

$$\Delta\mathbf{z} = \mathbf{z} - \hat{\mathbf{z}} = \mathbf{z} - \mathbf{h}(\hat{\mathbf{x}}) = \mathbf{G}\Delta\mathbf{x} + \Delta\varepsilon \tag{4.22}$$

where $\mathbf{z}$ is the actual measurement vector containing pseudo-ranges and delta ranges and $\hat{\mathbf{z}} = \mathbf{h}(\hat{\mathbf{x}})$ is the predicted measurement vector computed from Eq. (4.20) for all satellites in view. The first matrix on the right hand side is the is measurement connection matrix $\mathbf{G}$, and is frequently referred to as the geometry matrix because it contains the line of sight vectors from the user to the satellites. The vector $\Delta\varepsilon$ models the noise in the measurements and its elements are assumed to be white noise sequences. The corresponding measurement noise covariance matrix $\mathbf{R}$ is a measure of the receiver noise and is receiver dependent.

Extended Kalman Filter

With the four matrixes Φ, $\mathbf{Q}$, $\mathbf{G}$ and $\mathbf{R}$ known, the EKF is in principle specified. Initial estimates of the state $\hat{\mathbf{x}}_o$ and the state error covariance matrix $\mathbf{P}_o$ have to be established for its initialization. An EKF formulation can be found in (Axelrad 1996) and is given in Appendix C.4.

4.3.2 External Kalman Filter

The primary purpose of the external Kalman filter is to provide an estimate of the acceleration $\mathbf{a}_g$, based on the receiver velocity output $\mathbf{v}_g$. This approach is useful in instances where the receiver does not provide acceleration information as part of its navigation solution, because it estimates acceleration information in a simple and rapid manner. The approach discussed here was also pursued for the experimental work documented in this thesis, since acceleration information was not available from the Novatel 3151R GPS receiver.

GPS velocity information is commonly provided in a North, East and Down coordinate frame.† Separate, but identical Kalman filters are therefore necessary to estimate accelerations in the North, East and Down directions. Figure 4.4 shows the external Kalman filter configuration schematically with the GPS receiver block representing the Novatel 3151R.‡ In the following, the filter is presented in the context of the East direction.

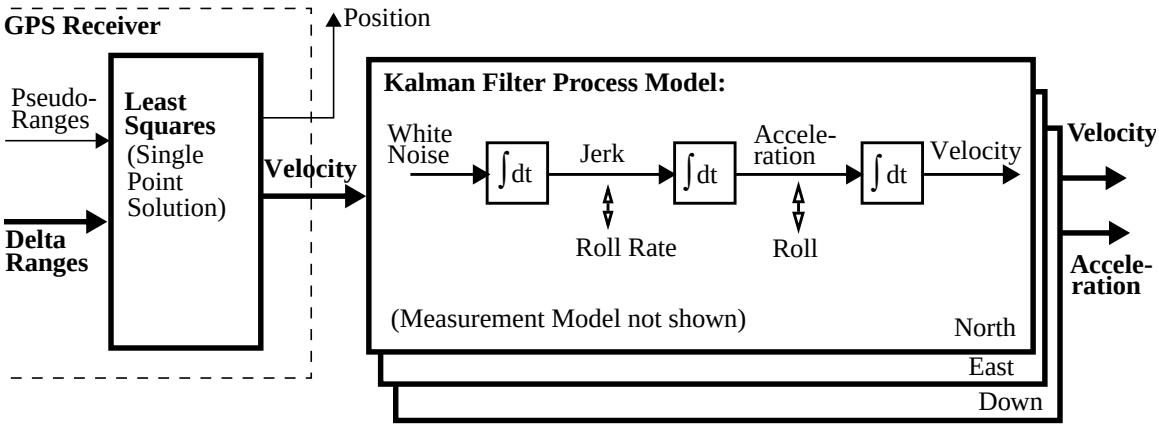

Figure 4.4: External Kalman Filter Process Model

† The velocity information may be expressed in North, East and Down components or, equivalently, in horizontal speed, ground track heading and vertical speed.

‡ In Figure 4.4, the velocity information is shown to be filtered by the Kalman filter as well. This reflects the actual implementation used in the flight tests described in later chapters. However, the velocity filtering is not required for the purpose of generating flight path angle information, because the noise level affecting the GPS receiver velocity output is more than two orders of magnitude lower than typical aircraft speeds.

Process Model

As for the internal filter discussed above, it is appropriate to include jerk as a separate state in the estimation and not to treat it as white noise. Consequently, a triple integrator plant is used as the process model including velocity, acceleration and jerk states. The discrete-time model of the triple integrator for the East direction is given by

$$\mathbf{x}_E(k) = \Phi_E(\Delta t)\,\mathbf{x}_E(k-1) + \mathbf{w}_E(k-1) \tag{4.23}$$

where the state vector $\mathbf{x}_E(k)$ for the east direction and the state transition matrix Φ_E are defined as

$$\mathbf{x}_E(k) = \begin{bmatrix} v_E(k) \\ a_E(k) \\ j_E(k) \end{bmatrix} = \begin{bmatrix} \text{East Velocity} \\ \text{East Acceleration} \\ \text{East Jerk} \end{bmatrix} \tag{4.24}$$

$$\Phi_E = \begin{bmatrix} 1 & \Delta t & \Delta t^2/2 \\ 0 & 1 & \Delta t \\ 0 & 0 & 1 \end{bmatrix} \tag{4.25}$$

Δt is the filter update time interval. $\mathbf{w}_E$ is a vector which models the process noise and its elements are white noise sequences. The corresponding process noise covariance matrix $\mathbf{Q}_E$ is determined (Gelb 1974) as

$$\mathbf{Q}_E = \begin{bmatrix} \Delta t^5/20 & \Delta t^4/8 & \Delta t^3/6 \\ \Delta t^4/8 & \Delta t^3/3 & \Delta t^2/2 \\ \Delta t^3/6 & \Delta t^2/2 & \Delta t \end{bmatrix} \cdot S_E \tag{4.26}$$

where S_E is the white noise spectral amplitude. The process noise covariance $\mathbf{Q}_E$ represents the uncertainty in the dynamic model. The determination of S_E is, typically, based on an estimate of the expected vehicle dynamics.

Measurement Model

The measurement input for the filter is the East velocity output of the GPS receiver. The measurement model for the East direction is given by

$$y_E(k) = \mathbf{H}_E \mathbf{x}_E(k) + u_E(k) \tag{4.27}$$

where, for the east direction, $y_E(k) = v_E(k)$. The measurement connection matrix $\mathbf{H}_E = [1\ 0\ 0]$ represents the fact that only velocity measurements are available. $u_E(k)$ models the measurement noise and the corresponding measurement noise variance R_E (a scalar in this case) depends on the particular GPS receiver used.

Kalman Filter Implementation

With Φ_E, $\mathbf{Q}_E$, $\mathbf{H}_E$ and R_E known, the Kalman Filter is specified. The filter equations are similar to the ones for the EKF and are given in Appendix C.4. S_E and R_E have to be carefully selected or tuned since they have a significant impact on the convergence, bandwidth and accuracy of the filter solutions, as shown in Section 4.4.

In some instances the navigation solution of the GPS receiver itself is based on a Kalman filter. The use of an external Kalman filter to estimate acceleration will then result in a cascaded Kalman filter configuration which may yield slightly suboptimal solutions as compared to a single Kalman filter solution (Levy 1996).

The velocity output of the Novatel 3151R GPS receiver is based on a single-point solution, and the approach outlined in this chapter does, therefore, not lead to a cascaded filter configuration. Moreover, as will be seen in subsequent chapters, this choice achieved sufficient performance to realize the desired objectives.

4.4 GPS Velocity and Acceleration Bandwidth

The generation of velocity and acceleration information is a two-step process, as discussed in Section 4.2 and Section 4.3. First, carrier Doppler phase information is extracted from the incoming GPS signal using a phase-locked loop. The PLL acts as a low pass filter, as can be seen from Eq. (4.7), tracking only signal dynamics in its bandwidth.

Next, all the delta range measurements are processed to generate velocity and acceleration information. This may be accomplished in two different ways, as shown in the previous section. In a first approach, the delta range measurements are combined in a Kalman filter to generate velocity and acceleration information (internal Kalman filter). Alternatively, the delta range measurements are first processed in a least squares solution to obtain instantaneous velocity which is subsequently fed into an external Kalman filter to estimate acceleration information. In both approaches, the Kalman filter assumes a dynamic model of the GPS receiver platform motion which determines the bandwidth of the filter. This section focuses on the external Kalman filter approach since it was implemented for the experimental work documented in this thesis.

The bandwidth of a filter is inherently connected to amount of lag the filter applies to its output signal and, thus, has a profound impact on the performance of a control system closed around it. In addition, the bandwidth of a filter determines the amount of noise at its output when subjected to a noisy input signal. In this context, one defines a *noise equivalent bandwidth* that characterizes the amount of noise a filter passes.†

There exists an obvious trade-off between good dynamic response and good noise performance of a filter. A higher bandwidth allows for better dynamic tracking at the expense of more noise in the output signal, and vice versa.‡ In the following, the noise

† It is defined as the bandwidth of a filter with unity gain which is fed by white noise and produces the same noise power at the output as the actual filter.

‡ It is important to distinguish between the *bandwidth* and the *output rate* of the receiver. The former determines the frequency range in which the receiver measures platform dynamics. The latter is the frequency at which the receiver makes the measurements available. While for control purposes (discussed in Chapter 5) both quantities are imporant, in the current context only the bandwidth is considered.

versus lag or bandwidth trade-off for both the PLL and the Kalman filter are briefly presented. The section concludes with some design considerations.

Phase-Locked Loop Noise vs. Bandwidth Trade-Off

Eq. (4.7) gives the transfer function of a third-order PLL with natural frequency ω_o. The corresponding PLL bandwidth, which scales with ω_o, defines the range in which the PLL tracks vehicle motion. In case of the Novatel 3151R, the receiver has a tracking bandwidth of 15 Hz. Also, Table 4.1 indicates that the dynamic tracking error due to jerk input is inversely proportional to ω_o^3. A higher natural frequency and tracking bandwidth reduces the dynamic stress errors and allows the measurements of vehicle dynamics at higher frequencies.

On the other hand, larger bandwidth leads to more noise in the delta range measurements and hence in the velocity estimate. It can be shown that the phase jitter standard deviation σ_θ of the PLL output due to thermal noise is given by

$$\sigma_\theta = \sqrt{\frac{B_n}{S/N_o} \cdot \left(1 + \frac{1}{2T \cdot S/N_o}\right)} \qquad [\text{rad}] \tag{4.28}$$

where B_n is the carrier loop noise equivalent bandwidth, S/N_o is the signal to noise ratio of the PLL input, and T is the predetection integration time (van Dierendonck 1996, Ward 1996). For the third-order loop of Eq. (4.7), B_n is shown to be $\omega_o / 1.2$ (van Dierendonck 1996). Eq. (4.28) indicates that the output phase noise variance scale linearly with B_n and hence with the PLL bandwidth. Consequently, a larger PLL bandwidth increases the noise in the velocity estimates. Figure 4.5 shows the output phase noise as a function of the signal to noise ratio for a number of different PLL noise bandwidths. The predetection integration time is 100 ms for all the plots.

As can be seen, acceptable dynamic response and noise performance of the PLL have to be found in a design trade-off.

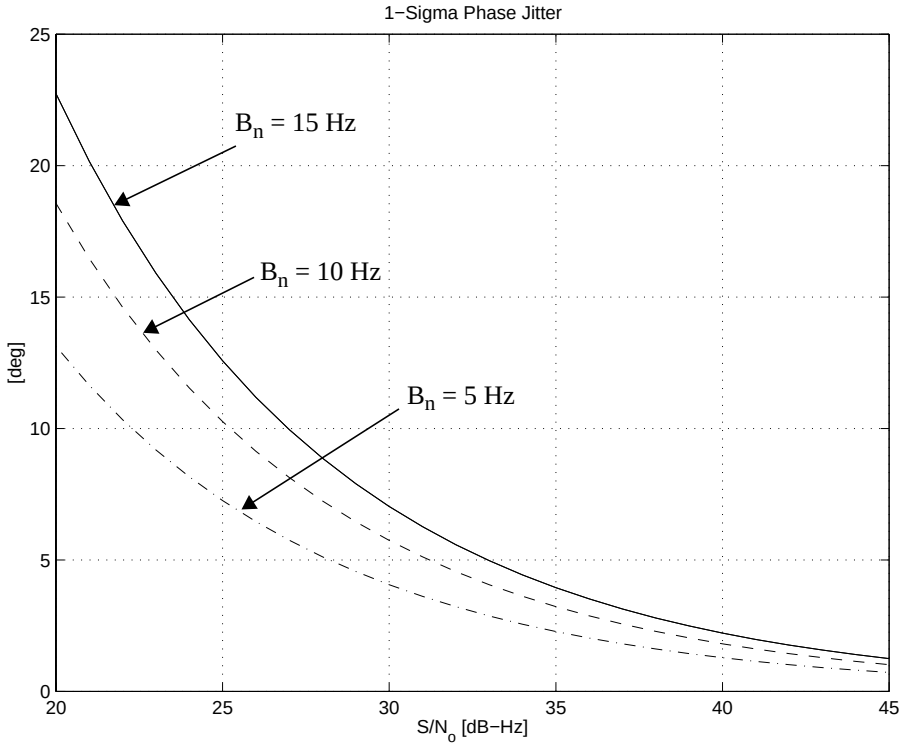

Figure 4.5: Phase Jitter for Different PLL Noise Bandwidths

Kalman Filter Noise vs. Bandwidth Trade-Off

In the following, the bandwidth of the external Kalman filter of Section 4.3.2 is examined.† The filter has a triple integrator process model and is used to estimate acceleration from velocity inputs. The filter was used throughout the research work documented in this thesis and was implemented in the flight test system discussed in Chapter 6.

The Kalman filter bandwidth is a function of the assumed receiver platform dynamics, expressed in the process noise covariance matrix **Q** (or in the white noise spectral amplitude S), and the measurement noise, expressed in the measurement variance R (a scalar in this case). The relation between the Kalman filter estimator bandwidth and quantities **Q** and R is best understood by examining the estimator transfer function as

† Only the bandwidth of the acceleration estimation is considered here. The filtering of the velocity information is not required for the purpose of generating flight path angle information, because the noise level affecting the GPS receiver velocity output is more than two orders of magnitude lower than typical aircraft speeds.

determined by the steady-state (positive definite) solution to the algebraic Riccati equation (ARE) (Friedland 1986). For reasons of simplicity the transfer function is determined in the continuous-time domain. Using the derivation in Appendix D and the continuous-time equivalent of the discrete-time triple integrator model given in Eq. (4.23) - Eq. (4.25), the velocity to acceleration transfer function is determined as

$$\frac{A(s)}{V(s)} = \frac{2\Omega^2 s^2 + \Omega^3 s}{s^3 + 2\Omega s^2 + 2\Omega^2 s + \Omega^3} \tag{4.29}$$

where A(s) and V(s) denote the acceleration and velocity in the s-domain. The natural frequency Ω is given by

$$\Omega = \left(\frac{S}{R\Delta t}\right)^{1/6} \tag{4.30}$$

where Δt denotes the filter update time interval of the discrete time filter. For the implementation discussed in this thesis, the update rate is 10 Hz and, thus, $\Delta t = 0.1$ seconds.

Eq. (4.30) indicates that the filter bandwidth is a function of the ratio of the white noise spectral amplitude and measurement noise variance. The equation also illustrates the noise versus bandwidth trade-off. A larger dynamic uncertainty increases Ω and subsequently the filter bandwidth. This effect causes the filter to track vehicle dynamics more readily, but comes at the expense of larger errors in the state estimate. This can be seen from the error covariance propagation equation in the Kalman filter, where the state error covariance **P** is *increased* by **Q** (Eq. (C.13) in Appendix C.4). On the other hand, with a smaller amount of dynamic uncertainty, or equivalently, with a 'relative' large amount of measurement noise, Ω and the filter bandwidth can decrease. In this case, estimation errors are more attributable to the measurement noise, rather than to the vehicle dynamics, and the state estimate is updated with more weight on the previously projected state estimate than on the new measurement (see state update equation Eq. (C.10) in Appendix C.4) which effectively causes the filter to narrow its bandwidth.

The measurement noise variance R is a function of the PLL phase noise variance given in Eq. (4.28). Static tests were performed to identify the noise performance of the Novatel 3151R receiver. Velocity data for the GPS receiver at rest was acquired. The velocity data in the Down direction is shown in Figure 4.6. Any correlated error due to Selective

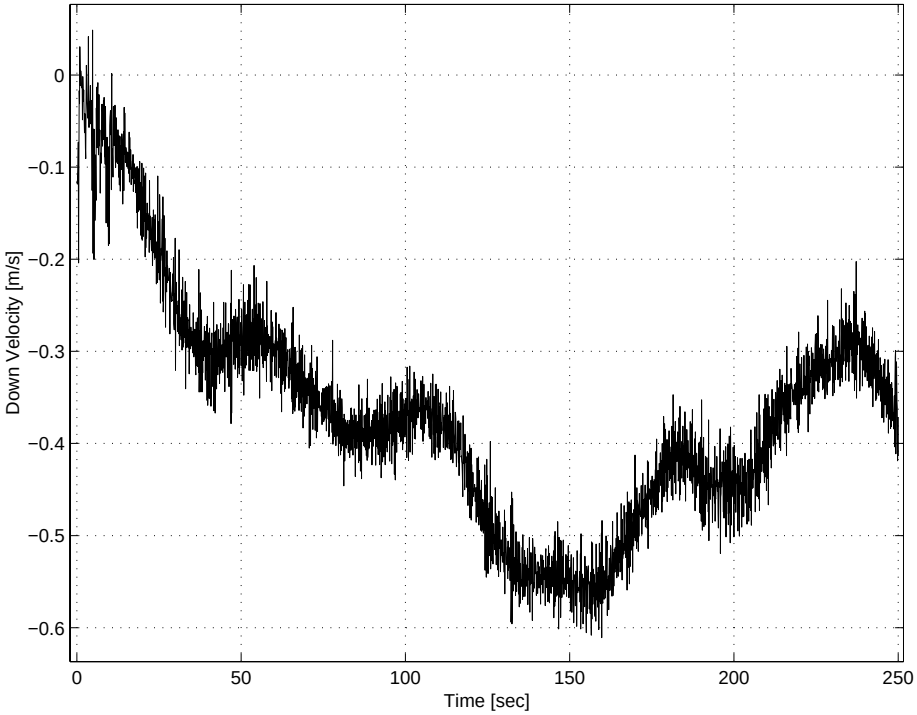

Figure 4.6: GPS Down Velocity Data Obtained Under Static Conditions

Availability was removed using polynomial splines, and an autocorrelation analysis showed that the remaining data was nearly uncorrelated. The resulting approximate noise variances were $R_N = R_E = (0.01)^2$ m^2/s^2 for the North and East directions, and $R_D = (0.1)^2$ m^2/s^2 in the Down direction for a 10 Hz update rate. The poorer noise performance of the Down direction is due to the poorer measurement geometry in the vertical direction.

The process noise covariance matrix **Q** depends strongly on the receiver platform dynamics. In most cases, the white noise spectral amplitude S initially has to be estimated and later, during tests, ‘tuned’ to achieve the desired performance. This was also the procedure in this thesis work. The resulting white noise spectral amplitudes for the North

and East directions are $S_N = S_E = 0.001$ m^2/s^7 for the North and East directions and $S_D = 0.0005$ m^2/s^7 for the Down direction. This reflects the notion that the aircraft dynamic uncertainty in the vertical direction is smaller than in the lateral directions.

With the values for S and R determined, the following natural frequency Ω for the North, East and Down directions result: $\Omega_N = \Omega_E = 2.15$ rad/sec, $\Omega_D = 0.89$ rad/sec. This roughly corresponds to a time constant of 0.5 seconds for the North and East direction, and 1.1 seconds for the Down direction. The Bode magnitude and phase plots of the transfer function Eq. (4.29) for the North and East directions are shown in Figure 4.7.

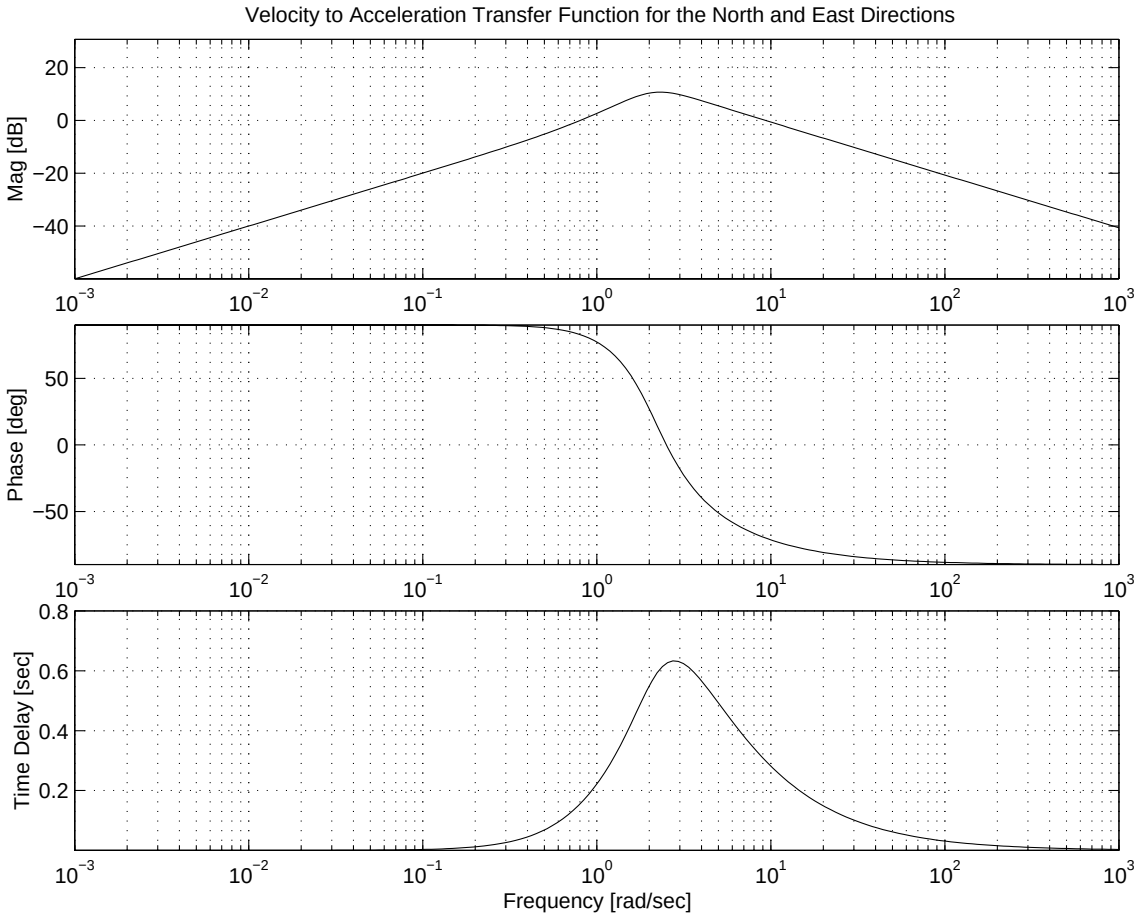

Figure 4.7: Bode Plot of Velocity to Acceleration Transfer Function for the North and East Directions

As can be seen from the magnitude plot, the filter acts as a differentiating element up to the break frequency Ω. Inputs at higher frequencies are essentially attenuated. The phase plot indicates that for low frequencies the filter has 90 deg phase lead, as expected from a differentiator. For higher frequencies, however, the filter introduces lag. For the case shown in Figure 4.7, phase lags and corresponding time delays are apparent above 0.3 rad/sec. The time delay at higher frequencies sets the limit for the usable bandwidth of the filter and thus has a profound impact on the control performance of a system closing the control loop around it.

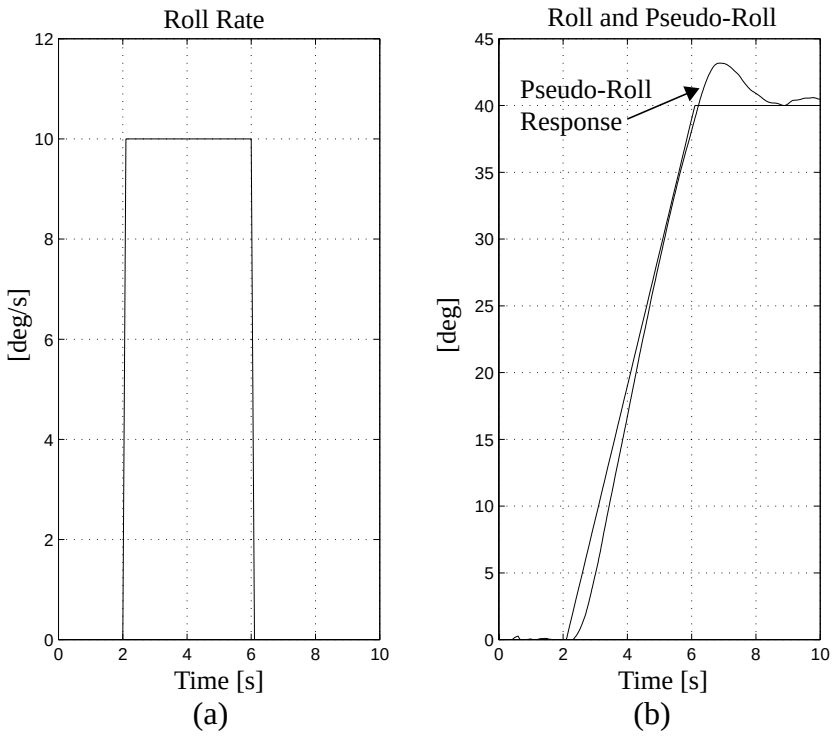

Figure 4.8: Simulated Pseudo-Roll Time Response

In order to demonstrate the time domain behavior of the filter, a pseudo-roll response during an aircraft roll maneuver is simulated using the following procedure. A simulated step in roll rate of 10 deg/sec, shown in Figure 4.8(a), is applied to a simple point mass aircraft model for a time interval of 4 seconds. This leads to an aircraft roll ramp leveling of at 40 deg shown in Figure 4.8(b). The corresponding aircraft velocity profile is fed into the Kalman filters for the North and East direction with the natural frequency $\Omega_N = \Omega_E = 2.15$ rad/sec. The filtered acceleration and the velocities are then fed into the pseudo-

attitude synthesis algorithm given in Eq. (3.31) - Eq. (3.37) and the resulting pseudo-roll response is shown in Figure 4.8(b). A time lag of 0.5 seconds is observable in the initial pseudo-roll response. Also, some overshoot behavior is noticeable.

In the next chapter, the Kalman filter is investigated in context of an aircraft control scheme which closes the loop around pseudo-attitude.

Receiver Design Considerations

From the treatment so far, it becomes apparent that the GPS velocity and acceleration bandwidth is a function of the PLL and Kalman filter designs. As seen from Eq. (4.28), the PLL bandwidth determines the noise of the PLL output and consequently the Kalman filter measurement noise. The latter, in conjunction with a receiver platform dynamic model, sets the Kalman filter bandwidth. In order to obtain optimal bandwidth and noise performance of GPS acceleration and velocity information, the design of the PLL and Kalman filter have to be addressed in an integrated manner taking the anticipated vehicle dynamics into account.

For the flight test implementation used in this research, the software of the GPS receiver was not accessible and only the external Kalman filter was adjustable to meet the desired performance.

4.5 GPS Velocity and Acceleration Errors

GPS has a number of error sources affecting its position accuracy. They are discussed in (Parkinson 1996b). The most pertinent error sources to velocity and acceleration generation are Selective Availability (SA) and receiver noise. These are discussed next.

Selective Availability

SA is the intentional degradation of the GPS system, with the objective to deny the full position and velocity and, hence, acceleration accuracy to the unauthorized user. SA induces two error components. It subjects the satellite clock frequency to a dither frequency which creates a slowly varying modulation effect on both the pseudo-range and

delta range measurements. This clock dither induced error has a correlation time of approximately 2 - 5 minutes.† In addition, SA manipulates the ephemeris data in the navigation message which results in a quasi random bias error in the pseudo-range measurement. This component of SA is, thus, of no consequence to the generation of GPS velocity and acceleration information.

SA is encrypted and can be removed by the US military or other authorized users who possess the necessary crypto key. The majority of the civilian sector, however, are unauthorized users and are therefore subjected to SA. An effective way for the civilian sector to correct for SA is through the use of Differential GPS (DGPS) where GPS reference stations with known positions calculate the error in the GPS measurements and transmit it to the nearby user. The Federal Aviation Administration (FAA) is planing to implement DGPS stations on a national level with the Wide Area Augmentation System (WAAS) in 2001 and on a local level with the Local Area Augmentation System (LAAS) in 2005 (FRP 1996). Furthermore, Federal Policy calls for the phaseout of SA in the first decade of the next century (PDD 1996). While in the long run SA does not pose a problem, it is currently corrupting GPS velocity and acceleration measurements. Its impact on the synthesis of pseudo-attitude is therefore briefly examined.

The GPS Standard Positioning Service calls for a range rate error of ≤ 2 m/sec (not-to-exceed, NTE) and for a range acceleration of $\leq 19\text{mm/sec}^2$ (NTE) for each satellite (GPS 1995). These values seem overly pessimistic, however, since they characterize the worst case conditions and not SA induced errors observed most of the time. A simulation using the Rather analytic SA model‡ (van Graas 1996) should serve as a more realistic

---

† Due to its highly correlated nature, SA can not be filtered by the Kalman filter which assumes white noise input, and effectively is not distinguishable from vehicle motion. Also, since the error dynamics are not properly known, they can not be modeled and estimated as part of the filter estimate (Axelrad 1996, van Graas 1996). A method for dealing with unknown signal dynamics is to use a Schmidt filter outlined in (Brown 1997). The method does not eliminate the error but does take it properly into account in the error covariance calculation leading to better overall performance.

‡ A model more frequently used is the second-order Gauss-Markov model which generates SA range and range rate errors. It has been adopted by the RTCA for GPS receiver-testing purposes. The main disadvantage of the model and the reason it is not used here is that both the range and range rate SA are noisier than that observed from actual satellite measurements. A noisier than actual SA range rate error introduces an unrealistic noise error in the simulated SA range acceleration.

illustration of the impact which SA has on GPS velocity and acceleration, and on the synthesis of pseudo-attitude. The Rather model reflects the impact of the clock dither frequency SA component and generates *representative* SA pseudo-range and delta range errors by integrating a random jerk step sequence three times to yield SA acceleration, range rate and range, respectively. This model was developed by observing raw SA range measurement errors and its derivatives and noticing the ramp structure in the range acceleration data.

A set of simulated SA range, range rate and range acceleration histories over the course of one hour is shown in Figure 4.9(a), (b) and (c). The standard deviations of the simulated range rate and acceleration are 0.16 m/s and 1.3 mm/s^2, and correspond closely to the values 0.12 m/s and 2 mm/s^2 of actual SA data cited in (van Graas 1996). Additional simulated SA range rate and acceleration errors, representing the SA errors of other satellites, are necessary to compute the velocity and acceleration error in North, East and Down coordinates. They are not shown here.

The velocity and acceleration errors due to SA depend on the instantaneous satellite geometry and are computed using the following fundamental error equation (Parkinson 1996b)

$$\Delta\mathbf{x} = (\mathbf{G}^{\mathrm{T}}\mathbf{G})^{-1}\mathbf{G}\,\Delta\rho_{\mathrm{SA}} \tag{4.31}$$

where $\Delta\mathbf{x}$ are the velocity and acceleration errors due to SA in North, East and Down coordinates, $\mathbf{G}$ is the geometry matrix containing the line of sight vectors from the user to the satellites, and $\Delta\rho_{\mathrm{SA}}$ are the SA range rates and accelerations, respectively, of more than four satellites.

The matrix $(\mathbf{G}^{\mathrm{T}}\mathbf{G})^{-1}$ in Eq. (4.31) is the geometric dilution of precision (GDOP) matrix and the square root of its diagonal elements are a quantitative measure of the satellite geometry. They are the North DOP, East DOP, Down DOP and a time DOP, in that order, for a vector $\Delta\mathbf{x}$ whose elements are ordered in the same way. The lower the DOP values, the smaller are the errors in the NED directions. Figure 4.9(f), (g) and (h) show

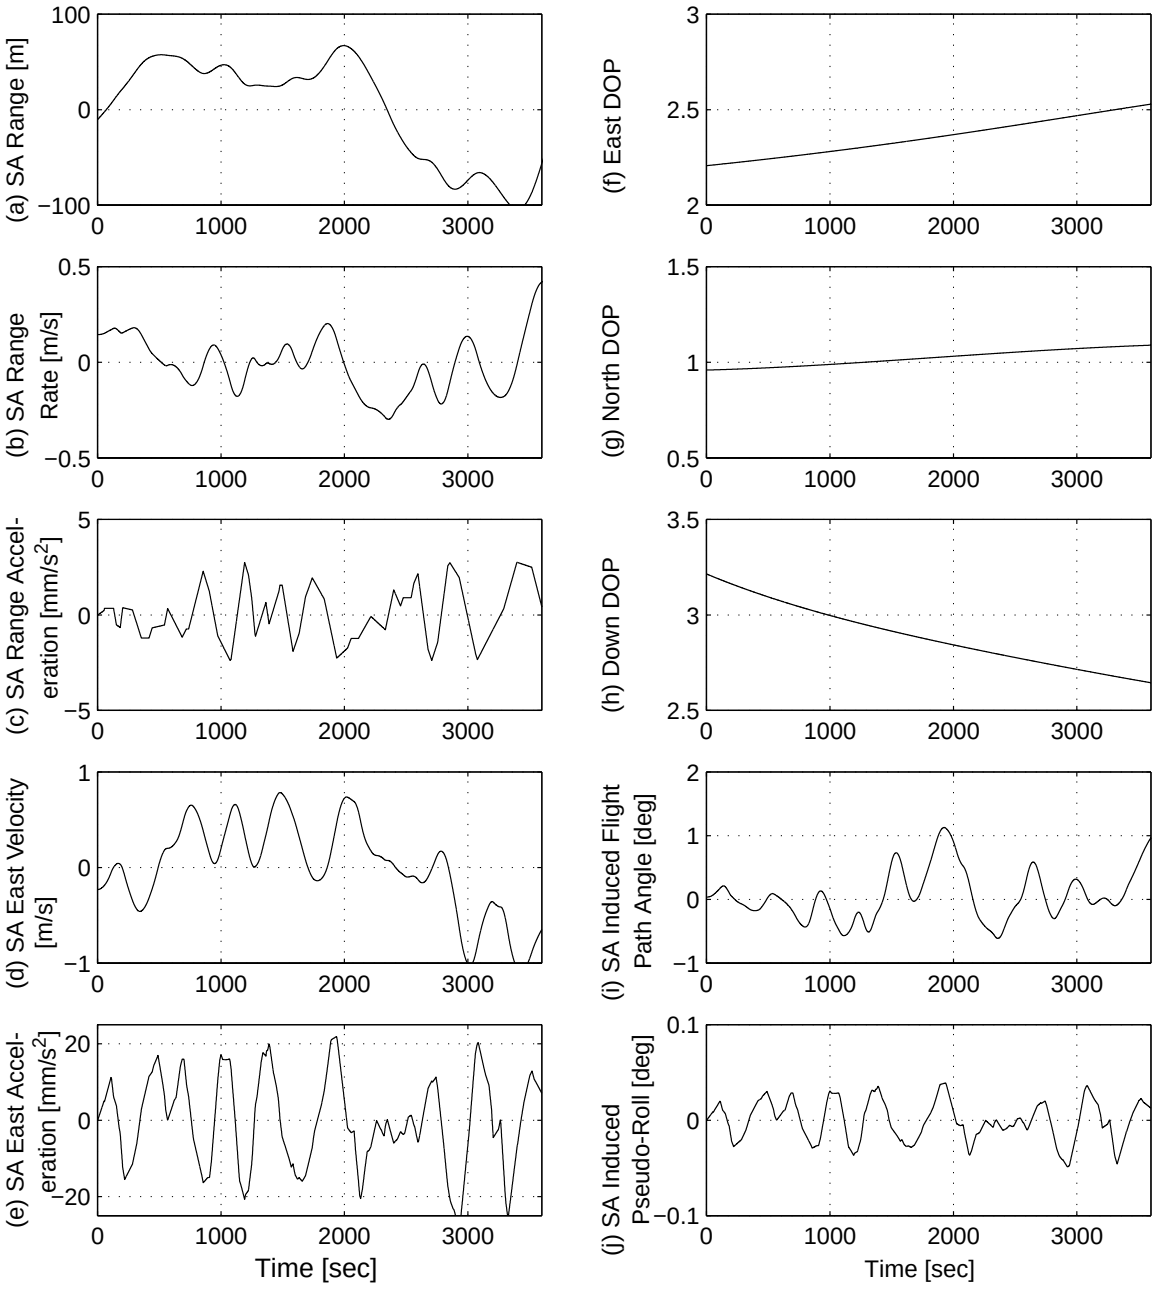

Figure 4.9: Simulated Impact of SA on Pseudo-Attitude

typical DOP values for a time interval of one hour. They were obtained from a simple GPS satellite model given in (Brown 1997). The calculated East velocity and acceleration are shown in Figure 4.9(d) and (e).

The standard deviations of the simulated velocity errors are σ_E = 0.47 m/s, σ_N = 0.21 m/s and σ_D = 0.36 m/s. The simulation closely matches the Novatel 3151R receivers nominal performance of 0.25 m/s RMS in stand-alone mode (Novatel 1996).

Figure 4.9(i) and (j) show the corresponding SA induced pseudo-attitude errors for an aircraft flying straight and level at 100 knots groundspeed in northern direction. The flight path angle error is a function of the aircraft groundspeed with the error decreasing at higher speeds. The plot in Figure 4.9(i) shows that the error does not exceed 1.2 deg for this flight condition. The standard deviation of the flight path angle error is 0.34 deg and does not constitute a significant error for most applications. However, for certain applications, such as for its unaided use in automatic landings, the SA induced flight path angle error is currently too large.

Figure 4.9(j) shows the pseudo-roll angle error due to SA. It does not exceed 0.05 deg and is, thus, negligible. This is not entirely surprising since the SA range acceleration is three orders of magnitude smaller than the gravity vector against which it is compared to compute the pseudo-roll angle. Thus, it can be concluded that SA currently does not severely limit the applicability of the pseudo-attitude concept for most applications. Moreover, long term plans call for a termination of SA (PDD 1996).

Receiver Noise

The receiver noise is mainly due to thermal noise and is uncorrelated in nature. As explained in Section 4.4, the receiver noise is determined by a bandwidth versus noise performance trade-off. For the Novatel 3151R GPS receiver, the noise performance during static tests was shown to be 0.01 m/sec RMS (North and East) and 0.1 m/sec RMS (Down).

4.6 GPS Integrity, Availability and Continuity

This section briefly highlights GPS integrity and availability issues because of their relevance to the use of GPS-based pseudo-attitude. These areas are subject to extensive on-going research and development efforts, and a full treatment of this topic is beyond the scope of this thesis.

In addition, this section outlines the effects of short satellite outages or changes in the in-view satellite constellation on the velocity and acceleration information. In this thesis, these issues are referred to as *GPS continuity* issues.†

GPS Availability and Integrity

Current GPS availability and integrity levels do not allow for the use of GPS as a sole means navigation sensor (FAA 1996).‡ Augmentation systems, such as the Wide Area Augmentation System (WAAS) and Local Area Augmentation System (LAAS), are currently under development. Once fully operational, they are expected to provide sufficient availability and integrity for GPS to be the sole means navigation system for all phases of flight from en route to precision approaches (FRP 1996).

It is evident that the novel use of single-antenna GPS as a *flight control* or *attitude* sensor is tied to different, more stringent availability and integrity requirements. This is because the flight control loop closure, by its “inner loop” nature, is more time critical than the outer guidance loop. An aircraft losing its navigation and guidance capability or being mislead by an undetected satellite anomaly can still be controlled while guidance

† It should be noted that the term *continuity* may be used differently throughout the GPS related literature.

‡ A distinction is made between a navigation system certified as *supplemental* or *sole means*. A supplemental navigation system must provide navigation information at the required accuracy level, if it can, and recognize instances when it cannot. In the latter case, the system must warn the user in a timely manner, who then switches to the primary navigation system. A sole-means navigation system must not only be able to detect anomalies, but it must also recover by removing the faulty satellite from the solution.

and separation may be obtained from radar vectors supplied by the ground. The erroneous indication or the loss of attitude information, on the other hand, may cause the aircraft to loose control within seconds of their occurrence.

It is apparent from these considerations that GPS-based attitude may be suited best as a *supplemental* attitude sensor in conjunction with additional attitude sensors. Also, the current and future GPS availability and integrity standards may allow *limited use of GPS as a sole-means* attitude sensor. Applications with lower availability and integrity standards, such as expendable unmanned aerial vehicles (UAV), may be envisioned to use GPS as a cost effective sole-means attitude sensor. Finally, the use GPS prediction software and mission planing tools enables the search for time intervals and coverage areas with sufficient high availability for sole-means GPS-based flight control. The implications of limited GPS availability and integrity on the use of GPS as an attitude sensor and possible applications are further discussed in Chapter 10.

GPS Continuity

Changes in the satellite configuration tracked by the GPS receiver may sometimes lead to jumps in the GPS navigation solution. A large discontinuity in measured GPS velocity could potentially induce a large jump in acceleration and thus pseudo-roll. Fortunately, as will be shown below, the effects of satellite configuration changes are minor most of the time.

Configuration changes may occur due to rising and setting satellites, or because a satellite is temporarily obscured from the antenna field of view by the aircraft structure (e.g. while the aircraft is banked). The jumps are the result of GPS pseudo-range and delta-range measurement errors, including primarily the effects of Selective Availability (SA). Since the measurement errors are different for each satellite, a position and velocity solution based on a changed satellite configuration may differ from a solution based on a previous configuration. The magnitude of the discontinuity is dependent on the number of

visible satellites and their distribution before and after the configuration change. A larger number of satellites tracked by the receiver and a more evenly distributed satellite configuration reduce the sensitivity of the navigation solution to configuration changes.

Based on a simplistic analysis, the effects of satellite configuration changes on GPS-based velocity and, ultimately, on pseudo-attitude appear to be benign. This is attributable to a combination of two factors. First, SA range rate errors are small compared to the velocity of conventional aircraft. Since the SA induced jumps in GPS velocity are smaller than the SA range rate errors, they are not apparent in flight path angle at typical aircraft speeds. Furthermore, the corresponding jumps in GPS acceleration are commonly much smaller than gravity against which they are compared when computing pseudo-attitude. Second, the Kalman filter estimating acceleration acts as a low-pass filter which reduces the effects of a jump in observed acceleration and prevents instantaneous changes in pseudo-roll.

While the effects of configuration changes on velocity and pseudo-attitude are benign, they may cause noticeable discontinuities in GPS position. This is due to the presence of slowly varying SA range errors which are large in amplitude.

The availability of differential corrections eliminates the effects of SA and reduces most of the impact of a changing satellite configuration. Also, mounting the GPS antenna at a location with an unobstructed field of view in all directions greatly reduces the number of satellite configuration changes.

4.7 Chapter Summary

This chapter investigated the mechanisms of GPS velocity and acceleration information generation. In addition, it highlighted the issues, trade-offs and limitations in the signal processing architecture and in the GPS system as a whole which pertain to the synthesis of pseudo-attitude. In the following, the most important points are summarized:

- Three dimensional user velocity and acceleration can be determined from the observed Doppler frequency shifts of the received GPS carrier signal. GPS receiver phase-locked loops (PLL) measure the frequency shifts in form of delta ranges. A third or higher-order PLL is necessary to track the GPS carrier signal under dynamic conditions (i.e. under line-of-sight accelerations). A PLL design trade-off exists between good tracking performance (large bandwidth) and low noise (low bandwidth).

- GPS acceleration and velocity information can be estimated from delta range measurements using a Kalman filter. Alternatively, if velocity information is provided by the GPS receiver, an external Kalman filter may be employed which estimates GPS acceleration from velocity information. In both cases, it is appropriate to include a jerk state in order to model aircraft roll rate appropriately. A trade-off involving Kalman filter bandwidth and noise in the estimates exists.

- Selective Availability (SA), which at the time of publication is the main error source affecting stand-alone GPS, does not severely limit the applicability of the pseudo-attitude concept. The SA induced error in flight path angle does not constitute a significant error for most applications. Its effect on pseudo-roll is negligible. Furthermore, Federal Policy calls for the termination of SA within the next decade.

- The novel use of single-antenna GPS as a flight control or attitude sensor requires more stringent GPS availability and integrity standards than those needed for navigation. Considering current and projected integrity and availability levels, single-antenna GPS-based pseudo-attitude may be suited best as a supplemental attitude sensor in conjunction with additional attitude sensors, or as a sole-means attitude sensor in applications with lower availability and integrity standards.

Chapter 5

Closing the Loop Around GPS Velocity Based Attitude Information

This chapter discusses the new paradigm for flight control loop closure around the inertial velocity and acceleration vector as sensed by the GPS receiver. The concept is based on controlling and sensing the velocity vector *directly*, rather than through the control of aircraft attitude, as in the traditional control approach. The treatment relies heavily on the material presented in Chapter 3 and 4.

The discussion shows how GPS velocity based attitude information can be modeled as part of the aircraft control loop and examines some of the issues of closing the flight control loop around the GPS-based velocity vector. It relies on a linearized analysis of the closed loop system consisting of aircraft dynamics, GPS instrumentation, and flight control system, and makes extensive use of the MATLAB simulation tool.

Section 5.1 gives a short introduction on the concept of velocity vector based flight control loop closure and presents related work. In Section 5.2, a linear model of the aircraft flight control loop is presented. The pseudo-attitude equations developed in Chapter 3 are linearized and the linearized models of the aircraft and GPS receiver are presented. The linearization yields decoupled models for longitudinal and lateral motion. Section 5.3 investigates the open-loop behavior of pseudo-attitude. In particular, the discussion illustrates the influence of adequate aircraft behavior around the velocity vector and the effects of uncoordinated flight and atmospheric disturbances on pseudo-roll. Section 5.4 investigates the closed-loop behavior of pseudo-attitude, and highlights some of the issues associated with a pseudo-attitude based flight control loop closure. Section 5.5 presents a case example of a pseudo-attitude based autopilot design. The example served as a basis for the autopilot design used in the flight tests described in Chapters 8 and 9. Finally, Section 5.6 summarizes the chapter.

Use of classical control methods is made throughout the chapter. In particular, the approach of closing single-input-single-output (SISO) loops in an sequential manner from the ‘inner most’ to the ‘outer most’ loop is pursued. In this thesis, the classical control approach is preferred over the modern control approach involving full state feedback because the former is more intuitive. In addition, it takes into account that pseudo-attitude, consisting of flight path angle and pseudo-roll angle, allows for a similar control loop structure as traditional attitude consisting of pitch and roll angle.

It is important to note that it is *not* the chapter’s intention to conduct a control performance comparison between traditional attitude based and pseudo-attitude based flight control loop closures. Rather, it is intended to show that it *is* possible to close the flight control loop around GPS-based velocity vector information and achieve adequate control performance to safely control the aircraft.

It should be appreciated that the characteristics of an aircraft control loop are inherently connected to the dynamics of a particular aircraft and that aircraft dynamic characteristics vary greatly between aircraft. It is thus difficult, if not impossible, to make generic quantitative statements about control loop properties which pertain to different aircraft. The considerations throughout this chapter are therefore tied to the dynamics of a particular aircraft. A Cessna 182 (C182) aircraft model is used for the analysis because the aerodynamic model is available in the open literature. In addition, the C182 resembles in dimension and weight a Piper Arrow aircraft† which was the aircraft type used for flight tests described in subsequent chapters. However, the issues treated here could easily be investigated in context of a different aircraft model.

5.1 Introduction

Recall that the velocity vector provides a more direct indication of the aircraft trajectory than traditional attitude. In this context Lambregts states that “Attitudes are not uniquely related to the flight path... . Manual control of the flight path using attitude for an inner loop control is therefore always iterative in nature” (Lambregts 1979).

† The aerodynamic model of the Piper Arrow is unavailable in the public domain.

In traditional control schemes, the pitch and roll attitude are controlled to effect a desired velocity vector and, hence, flight path change. In the traditional lateral control method, an aircraft roll angle is established in order to generate an acceleration force that, in turn, changes the velocity vector. Similarly, in traditional pitch attitude based longitudinal flight control schemes, pitch and thrust control are used to achieve a desired flight path angle and speed.

In the velocity based flight control paradigm, the flight path vector is sensed and controlled directly. In order for this approach to be successful, however, the aircraft has to be well behaved ‘around the velocity vector’. That is, unwanted aircraft modes have to be adequately suppressed and the aircraft response to control inputs has to be sufficient. This function, if necessary, is typically achieved using inner loop stability augmentation and control augmentation systems. As detailed in Chapter 2, the paradigm has the distinct advantage that is relies on low cost sensors to provide the measurements necessary to close the loops, namely GPS for the flight control loop and low grade (automotive) inertial sensors for the inner loops.

In the past, Lambregts developed a Velocity Vector Control Wheel Steering Mode in conjunction with a velocity vector aligned attitude display shown in Figure 3.8 (Lambregts 1979). The objective of the control mode and display systems is to largely reduce the need for pilot inner loop control by relying on automatic stability augmentation and, at the same time, to provide direct pilot control over the flight path. The aforementioned display is used to depict flight path angle and ground track together with other control variables, such as traditional aircraft attitude, and to close the pilot’s control loop. The velocity vector control steering mode and the necessary stability augmentation rely on traditional instrumentation consisting of an inertial navigation system (INS) and an air data system. The study indicated the substantial pilot workload reduction during manual airplane control can be achieved with the Velocity Vector Control Wheel Steering Mode.

While the concept of aircraft velocity vector control advocated in this thesis is related to the concept explored in Lambregts' work, the implementations are fundamentally different. Lambregts' velocity vector control law and stability augmentation relies on the feedback of *traditional* roll angle and requires *traditional instrumentation* including INS and air data, thus lacking the simplicity and low cost of the GPS-based pseudo-attitude approach presented in this thesis.

In the following sections, issues associated with the loop closure around velocity based pseudo-attitude are investigated. The effects of GPS as the sensing element are hereby included.

5.2 Linearization of the Aircraft Flight Control Loop

In this section, the linear representation of the aircraft flight control loop is developed. The aircraft flight control loop is shown in Figure 5.1. Aircraft velocity is measured by a single-antenna Novatel 3151R GPS receiver, and acceleration is estimated using a Kalman filter, as described in the previous chapter. Both GPS-based velocity and acceleration are available at a 10 Hz update rate and fed into the pseudo-attitude synthesis algorithm. Pseudo-attitude is used by the controller (human or autopilot) to compute aircraft actuator commands and close the loop. In the following, the elements of the loop are considered separately, and, where necessary, linearized.

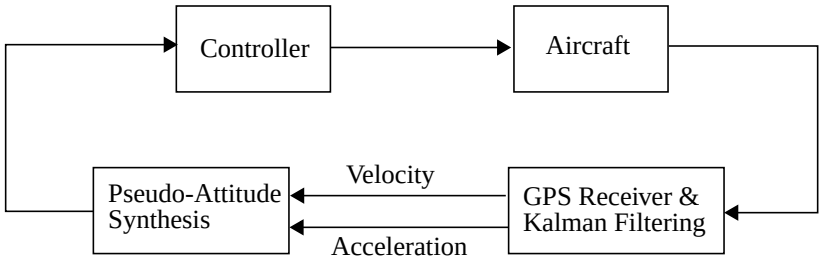

Figure 5.1: Aircraft Flight Control Loop

Linearized Aircraft Model

To obtain a linearized aircraft model, the 6 DOF non-linear aircraft equations of motion are linearized about a reference flight condition. For the treatment in this chapter, the reference flight condition corresponds to an aircraft on an approach flight path and is characterized by:

- Rectilinear, symmetric flight with wings level
- Vertical flight path angle of γ_o = -3 deg
- Constant aircraft speed of U_o = 85 knots in North direction
- No steady-state wind: $\mathbf{w} = 0$
- Initial altitude of 1500 ft.

By choosing the aircraft body x_b - axis to point into the relative wind for the nominal flight condition, one obtains the linearized aircraft model in stability axes (x_s, y_s, z_s).† The selected flight condition decouples the linearized longitudinal and lateral aircraft models completely. The linearized, longitudinal aircraft model is then given by

$$\dot{\mathbf{x}}_{Lon} = \mathbf{A}_{Lon}\mathbf{x}_{Lon} + \mathbf{B}_{Lon}\delta_e + \mathbf{G}_{Lon}\mathbf{g}_{Lon} \tag{5.1}$$

$$\mathbf{x}_{Lon} = [u, w, q, \theta, h, z_E], \qquad \mathbf{g}_{Lon} = [u_g, w_g, q_g] \tag{5.2}$$

where δ_e is the elevator input. The vector $\mathbf{x}_{Lon}$ represents the longitudinal disturbance state of the aircraft consisting of inertial velocity u in x_s direction, inertial velocity w in z_s direction, pitch rate q, pitch angle θ, longitudinal flight path deviation h, and Down position z_e in the NED frame. $\mathbf{g}_{Long}$ is the gust disturbance input vector consisting of forward gust input u_g, downward gust input w_g and pitch rate gust input q_g.

Similarly, the lateral aircraft model is given by

$$\dot{\mathbf{x}}_{Lat} = \mathbf{A}_{Lat}\mathbf{x}_{Lat} + \mathbf{B}_{Lat}\mathbf{u}_{Lat} + \mathbf{G}_{Lat}\mathbf{g}_{Lat} \tag{5.3}$$

$$\mathbf{x}_{Lat} = [v, p, r, \phi, d, \psi], \qquad \mathbf{u}_{Lat} = [\delta_a, \delta_r], \qquad \mathbf{g}_{Lat} = [v_g, p_g, r_g] \tag{5.4}$$

† This frame is a special body axes reference frame typically used in the study of small disturbances from nominal reference flight conditions.

The vector $\mathbf{x}_{Lat}$ represents the lateral disturbance state of the aircraft consisting of inertial velocity v in y_s direction, inertial roll rate p and yaw rate r, traditional roll angle ϕ, lateral flight path deviation d, and heading angle ψ. $\mathbf{u}_{Lat}$ is the control input vector consisting of the aileron input δ_a and the rudder input δ_r, and $\mathbf{g}_{Lat}$ is the gust disturbance input vector consisting of sideward gust input v_g, roll rate gust input p_g and yaw rate gust input r_g.

The matrices $\mathbf{A}_{Lon}$, $\mathbf{B}_{Lon}$, $\mathbf{G}_{Lon}$, $\mathbf{A}_{Lat}$, $\mathbf{B}_{Lat}$, and $\mathbf{G}_{Lat}$ contain the stability derivatives. They are evaluated in Appendix E for a Cessna 182 at 1500 ft. altitude and 85 knots approach speed on -3 deg glidepath (Roskam 1995). The resulting longitudinal and lateral aircraft eigenmodes are shown in Table 5.1.

Aircraft Modes	**Poles**	**Natural Frequency ω**	**Damping ζ**	**Time Constant T**
Phugoid	$-2.88 \pm 2.23i$	0.34	0.21	-
Short Period	$-0.072 \pm 0.336i$	3.64	0.79	-
Dutch Roll	$-0.683 \pm 2.6i$	2.69	0.25	-
Roll Subsidence	-9.40	-	-	0.106 sec
Spiral Mode	0.004	-	-	250 sec

Table 5.1: Longitudinal and Lateral Aircraft Modes

The short period is well damped so stability augmentation (SAS) is not required. The phugoid and dutch roll mode, however, are only lightly damped and typically require pilot control to damp oscillations. The spiral mode is slightly unstable, but with a large time constant. All the eigenmodes meet Level 1 handling qualities which allows for adequate aircraft control (Stevens 1992).

Linearized GPS Model

Aircraft velocity is measured up to a bandwidth of 15 Hz (> 90 rad) which is the receiver's carrier tracking phase-locked loop bandwidth. The bandwidth is approximately 25 times larger than the fastest relevant aircraft modes. Thus, in the context of aircraft velocity measurements, the GPS receiver can be modeled as an allpass filter.[†]

The acceleration is inferred from velocity through a Kalman filter estimator performing essentially a differentiation. The bandwidth of the inferred acceleration is $\Omega_N = \Omega_E = 2.15$ rad/sec for the North and East direction, as discussed in Section 4.4. The filter roll-off is close to the aircraft Dutch Roll mode and the filter induces a phase lag of approximately 90 deg at that frequency.

Although the actual GPS receiver provides discrete-time measurements at 10 Hz in a sample-and-hold fashion, it will be modeled as a continuous-time system. This is reasonable because the sampling period of 0.1 seconds is nearly 20 times faster than the time constants of the fastest relevant aircraft mode. This simplification allows for a control loop analysis in the Laplace domain. However, in order to cross-check this simplification, aircraft time responses, based on simulations including a sample-and-hold block in the GPS model, will be compared with the corresponding continuous-time responses.

Linearized Pseudo-Attitude

The next step is the linearization of the pseudo-attitude synthesis described by Eq. (3.30) - Eq. (3.37) in Section 3.1.4 about the given nominal flight condition. The linearization is performed in Appendix F and only the results are presented here. The linearized flight path angle γ is

$$\gamma = \theta - \frac{w}{U_o} = \theta - \alpha \tag{5.5}$$

† The latency of 50 ms in the velocity measurement (see Section 4.2.2) has a marginal effect on the closed loop performance in the frequency range considered and is thus not modelled here. Also, the impact of Selective Availability on GPS velocity and pseudo-attitude was considered previously in Section 4.5 and is not inlcuded in the discussion.

where both θ and w are elements in the longitudinal state vector. Linearized pseudo-roll $\tilde{\phi}$, expressed in stability axes, is

$$\begin{aligned}\tilde{\phi} &= \phi + \frac{f_y}{g_o \cos\gamma_o} \\ &= \phi + \frac{1}{g_o \cos\gamma_o}[U_o\dot{\beta} + U_o r - g_o\phi\cos\gamma_o] \\ &= \phi + \frac{1}{g_o \cos\gamma_o}[Y_v(v - v_g) + Y_p(p - p_g) + Y_r(r - r_g) + Y_{\delta a}\delta_a + Y_{\delta r}\delta_r]\end{aligned} \tag{5.6}$$

where f_y is the specific force in the direction of y_s, as measured, for example, by an accelerometer with its sensitive axis mounted along this axis. The gravitational acceleration g_o is 9.81 m/s^2 and the sideslip angle is $\beta = v/U_o$. The $Y_{(.)}$ are the stability derivatives of the sideforce Y with respect to the different states and inputs.

The first line in Eq. (5.6) states that pseudo-roll, to first-order, equals the sum of traditional roll and the ratio of sensed sideforce and the gravity component perpendicular to the flight path. Any experienced sideforce is thus interpreted as a pseudo-roll component. For coordinated flight, the sideforce term is zero and pseudo-roll, to first-order, is equal to traditional roll angle. The second and third equality of Eq. (5.6) show the influence of sideslip, sideslip rate, and the different aircraft states, control surface and gust inputs on pseudo-roll.

The linearization in Eq. (5.5) - Eq. (5.6) assumes that the GPS antenna is at or near the aircraft center of gravity. If the antenna is mounted at a considerable distance from the center of gravity, it senses additional lever arm motion which has to be accounted for by including the appropriate terms in the equations. This case is not considered here.

Simplified Linearized Aircraft Flight Control Loop

In order to complete the linearized aircraft flight control loop description, a useful simplification is performed. For practical implementation purposes, it is easier to switch the position of the linearized pseudo-attitude synthesis function block and the GPS receiver function block in the aircraft flight control loop. Figure 5.2 shows the simplified linearized aircraft flight control loop.†

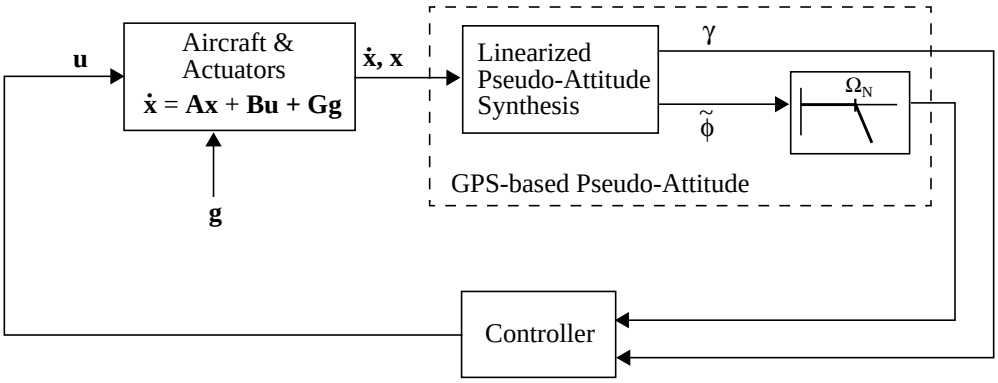

Figure 5.2: Simplified Linearized Aircraft Flight Control Loop

In the simplified control loop, pseudo-attitude is synthesized from the longitudinal and lateral aircraft states and their derivatives using Eq. (5.5) and Eq. (5.6). Flight path angle γ is computed from longitudinal velocity states. Since the GPS receiver is modeled as an allpass filter in the context of velocity measurements, no GPS induced filtering block is shown in the γ path in Figure 5.2.

Linearized pseudo-roll angle is computed from the aircraft acceleration states and actual roll angle, as indicated in Eq. (5.6). In order to properly introduce the band limiting effects of the Kalman filter acceleration estimator, linearized pseudo-roll is subsequently low-pass filtered. The low-pass filter is simply the Kalman filter estimator in Eq. (4.29) multiplied by an integration (1/s) operator to account for the fact that the input is no longer velocity but acceleration based pseudo-roll. The roll-off frequency of the resulting low pass filter is the same as for the original differentiator, namely $\Omega_N = \Omega_E = 2.15$ rad/sec.

† In principle, the lateral and longitudinal control loops are decoupled. For simplicity, they are shown together.

Pseudo-attitude is used by the controller, which may be a pilot or autopilot, to compute the necessary control inputs to the aircraft for proper inner-loop control. The aircraft elevator, aileron and rudder actuators are modeled as low pass filters with unity gains and time constants of 0.1 seconds.

5.3 Open-Loop Behavior

This section investigates the open-loop behavior of pseudo-attitude. The primary focus is on the lateral open-loop behavior examined in Section 5.3.1. Section 5.3.2 briefly discusses the longitudinal open-loop behavior.

5.3.1 Lateral Open-Loop Behavior

In previous chapters, coordinated flight and uniform wind have been assumed. In this section, open-loop simulations are performed to illustrate the effects of uncoordinated flight and atmospheric disturbances on pseudo-roll. The influence of adequate aircraft response around the velocity vector on pseudo-roll is demonstrated.

Figure 5.3(a) shows the un-augmented open-loop response of the C182 at the nominal flight condition to an aileron input of 5 deg over one second with the rudder fixed at the zero position. Figure 5.3(b) shows the open-loop response of the C182 augmented with a stability augmentation system (SAS). From top to the bottom, the plots show aircraft sideslip angle, traditional roll angle, unfiltered and filtered pseudo-roll angles, and the actuator input.

In order to separate the issues related to the pseudo-roll synthesis as shown in Eq. (5.6) from the influence of the limited bandwidth of GPS acceleration, both filtered and unfiltered pseudo-roll are shown. The unfiltered pseudo-roll angle is the pseudo-roll angle as computed in Eq. (5.6) but not bandlimited by the subsequent low-pass filter (modeling the Kalman filter). The filtered pseudo-roll angle, on the other hand, is the output of the low-pass filter and is thus the *actual control variable* available to close the loop as shown in Figure 5.2.

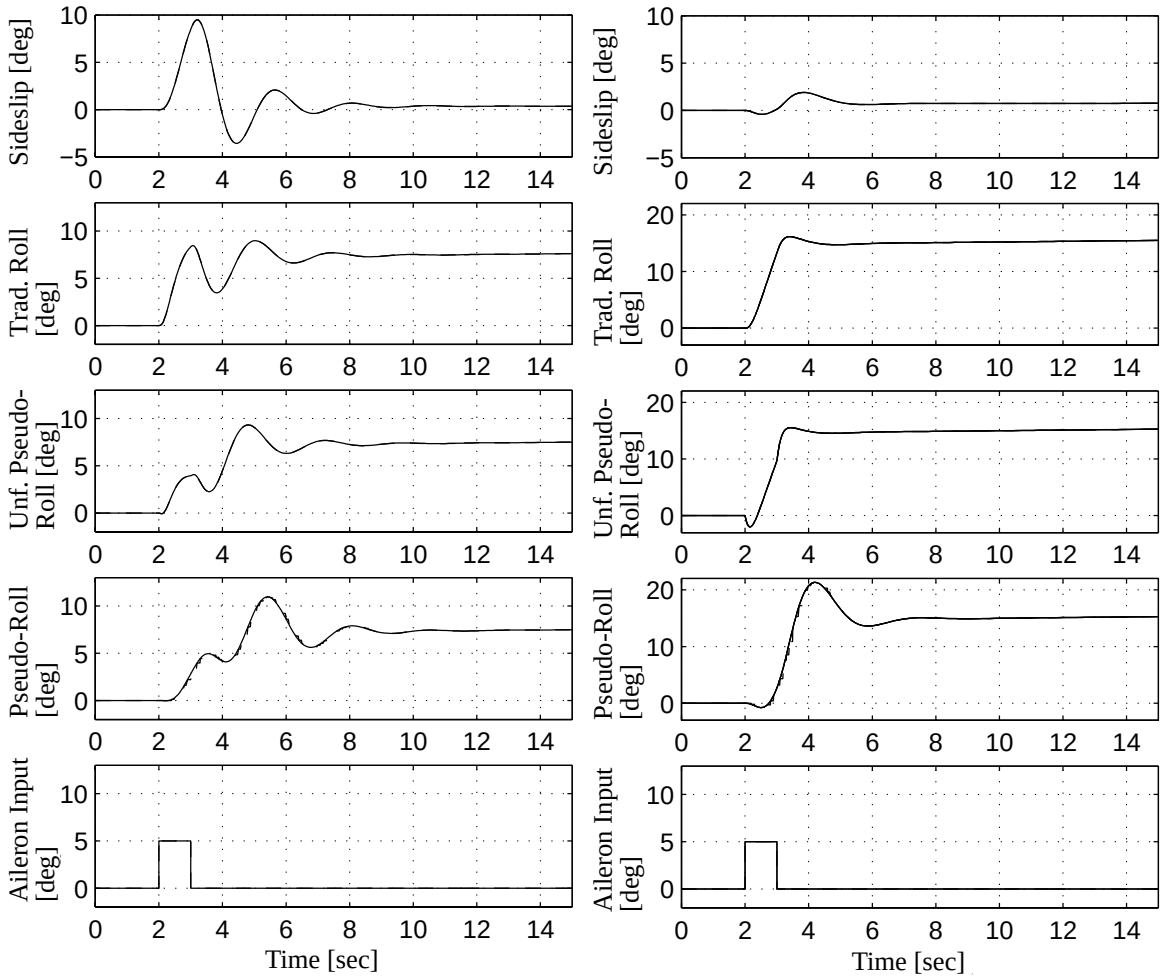

Figure 5.3: Simulated Open-Loop Pseudo-Roll Response to Aileron Input (a) Deficient Aircraft Behavior (b) Adequate Aircraft Behavior (with Augmentation)

The plot in Figure 5.3(a) shows a large initial sideslip excursion (at 2 - 4 seconds) up to approximately 10 deg due to aircraft adverse yaw, followed by moderately damped dutch roll oscillations. The large initial uncoordination causes, as expected, a considerable difference between the traditional aircraft roll and pseudo-roll angle. With the decay in sideslip angle, conventional roll and pseudo-roll response gain in similarity. The dutch roll oscillation in the traditional roll angle can also be observed in pseudo-roll. Comparing the unfiltered and filtered pseudo-roll time response reveals that the low-pass filter introduces delay of about 0.5 seconds and noticeable overshoot. This is in close correspondence to the observed Kalman filter response in Figure 4.8 (Section 4.4).

With the filtered pseudo-roll response, the corresponding discrete-time pseudo-roll response in dashed lines is shown. This response is generated using a simulation which includes a sample-and-hold block in the GPS model. As can be seen, the discrete-time pseudo-roll response is not distinguishable from the continuous-time response which verifies the initially assumed continuous model of the GPS receiver.

In Figure 5.3(b), the effects of adverse yaw and dutch roll on the aircraft response are mitigated using stability and control augmentation systems. In the current example, an Aileron-to-Rudder Interconnect (ARI) with a gain of -2 is used to reduce the effects of adverse yaw (Stevens 1992). In addition, a yaw rate r feedback with a high-pass washout filter of the form $H_{wo}(s) = -0.35s/(s+1.25)$ is used to increased dutch roll damping. This moves the dutch roll poles to $s_{dr} = -1.23 \pm 1.6i$ with a damping and natural frequency of $\zeta_{dr} = 0.61$ and $\omega_{dr} = 2.02$, respectively. The improvement in the aircraft response is immediately apparent. Adverse yaw is almost completely eliminated, and dutch roll motion is heavily damped.

With these improvements in place, traditional roll angle increases to approximately 15 deg in a steady manner with a small overshoot and no oscillations. The unfiltered pseudo-roll largely follows traditional roll angle with a slight delay of about 0.2 seconds. The initial reversal in the direction of the pseudo-roll build up is due to slight overcompensation. The filtered pseudo-roll exhibits noticeable overshoot and lags by approximately 0.5 seconds with respect to conventional roll angle. These characteristics correspond to the ones observed for the Kalman filter response in Figure 4.8. Finally, no difference between the continuous-time and discrete-time responses can be observed.

Figure 5.4(a) shows the aircraft response to a gust induced step in sideslip angle. The rudder and aileron are fixed at the zero position. The sideslip step of 1.1 deg at the nominal speed of 85 knots corresponds to a sidegust velocity of 2.8 ft/sec which slightly exceeds the standard deviation of gust velocity used in the literature (Bryson 1994). As can be seen, only a minor dutch roll motion is excited and observable in the pseudo-roll plots. Figure 5.4(b) shows the sidegust velocity response of the augmented aircraft. The increase in dutch roll damping is immediately apparent.

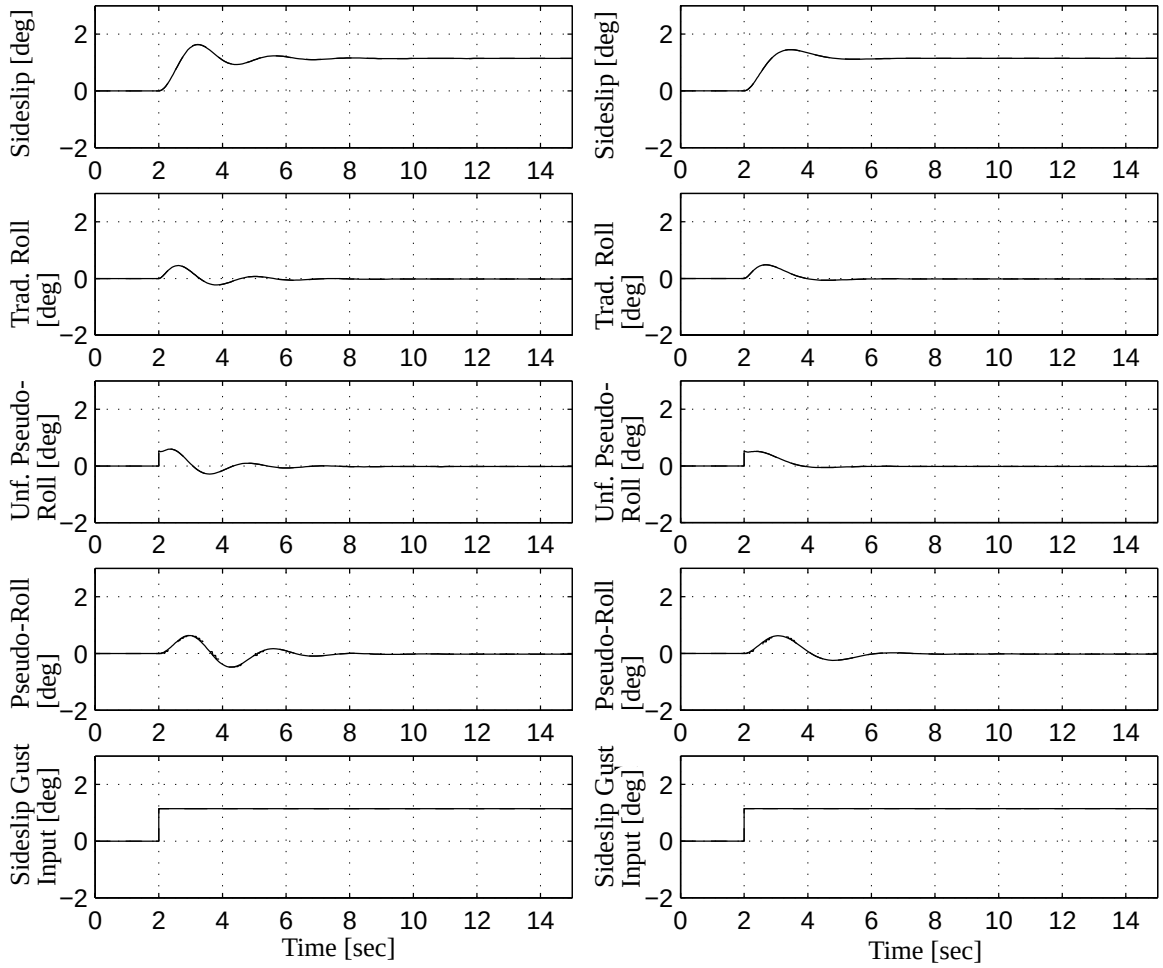

Figure 5.4: Simulated Open-Loop Pseudo-Roll Response to Atmospheric Disturbances (a) Deficient Aircraft Behavior (b) Adequate Aircraft Behavior (with Augmentation)

5.3.2 Longitudinal Open-Loop Behavior

In some instances, the damping of the longitudinal short period mode is insufficient. A flight path angle γ feedback may cause the short period to deteriorate rapidly and drive it unstable when the controller gain is increased.† Consequently, the loop has to be closed with low gains resulting in a low bandwidth system.

† The total system damping remains constant (McRuer 1973). Flight path angle feedback increases phugoid damping at the expense of reduced short period damping.

To mitigate the short period damping deficiency, a pitch rate signal can be fed back in a pitch damper SAS. This hardly alters the phugoid motion (McRuer 1973). With an improved short period damping in place, the γ feedback loop can be closed with increased gains resulting in a higher bandwidth system.

In case of the C182, the short period is sufficiently damped and no pitch damper is necessary.

5.4 Closed-Loop Behavior

In this section, the use of pseudo-attitude to close the flight control loop is examined for the C182. Section 5.4.1 examines issues of lateral loop closure using pseudo-roll, and Section 5.4.2 discusses longitudinal loop closure using flight path angle.

5.4.1 Lateral Loop Closure

A pseudo-roll command loop is obtained by feeding back the pseudo-roll angle of the aircraft and differencing it with the commanded pseudo-roll to obtain an error signal. The error signal is multiplied by a gain and, if necessary, appropriately filtered to produce aircraft aileron commands which control the aircraft in order to minimize the error signal.

Figure 5.5 shows the root locus and the bode plots of the pseudo-roll command loop using the augmented aircraft. In the current example, only a simple gain of 0.15 and no additional filtering is used. The ‘o’, ‘x’, and ‘*’ represent the position of the open-loop zeros, open-loop poles and closed-loop poles using the current gain, respectively. Figure 5.5(d) is a zoomed-out version of the root locus shown in (c). The dutch roll poles are roughly canceled by nearby zeros and this mode is therefore not present in the roll motion, as seen in the previous section. The most distinct feature in the root locus is the presence of a *non-minimum phase zero*, causing the complex pole pair introduced by the Kalman filter to move to the right half plane. The non-minimum phase behavior sets an *inherent limit to the achievable closed-loop bandwidth* of the system. Figure 5.5(a) and (b) show the corresponding bode magnitude and phase plots. With the gain used, a crossover frequency of 0.5 rad/sec and a phase margin of 75 deg is achieved.

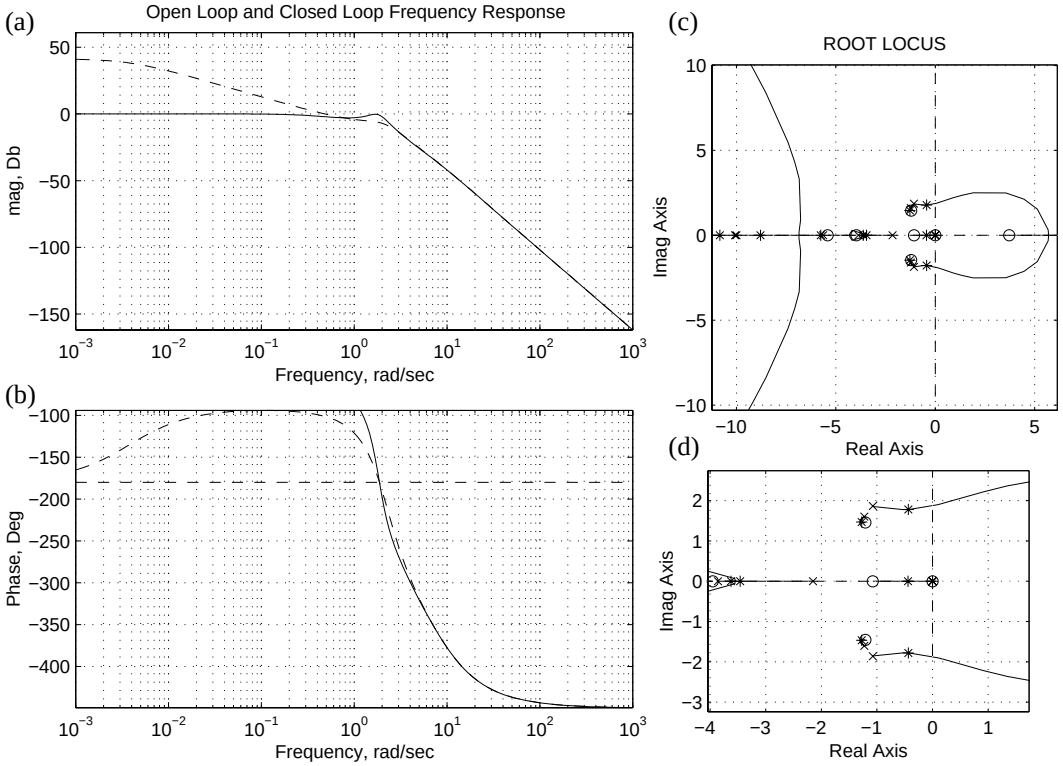

Figure 5.5: Bode and Root Locus Plots of Pseudo-Roll Command Loop

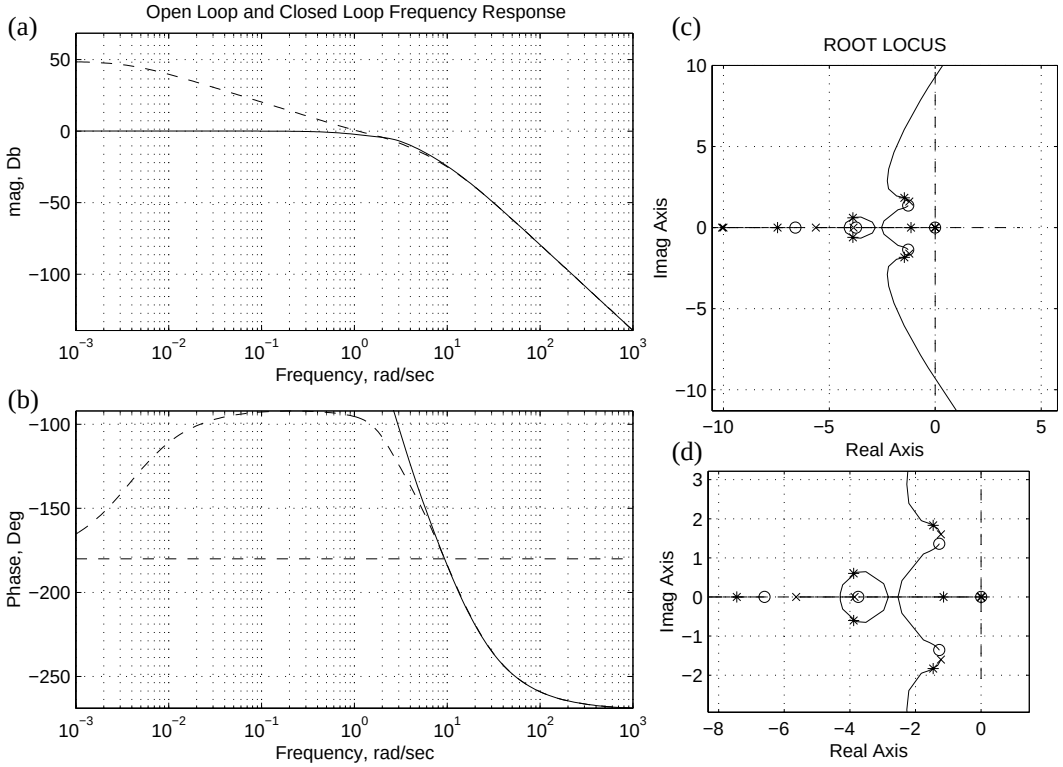

Figure 5.6: Bode and Root Locus Plots of Traditional Roll Command Loop

For comparison, Figure 5.6 shows the root locus and bode plots of a traditional roll angle command loop for the same gain. As can be seen, the control loop is minimum phase and a higher bandwidth is achievable than in the pseudo-roll case. For the current example, a crossover frequency of 1 rad/sec and a phase margin of 85 deg is obtainable. Traditional roll feedback allows for a higher bandwidth system since traditional roll *effects* a velocity vector change, while pseudo-roll is a *measure* of this change, and constitutes thus an outer-loop variable.

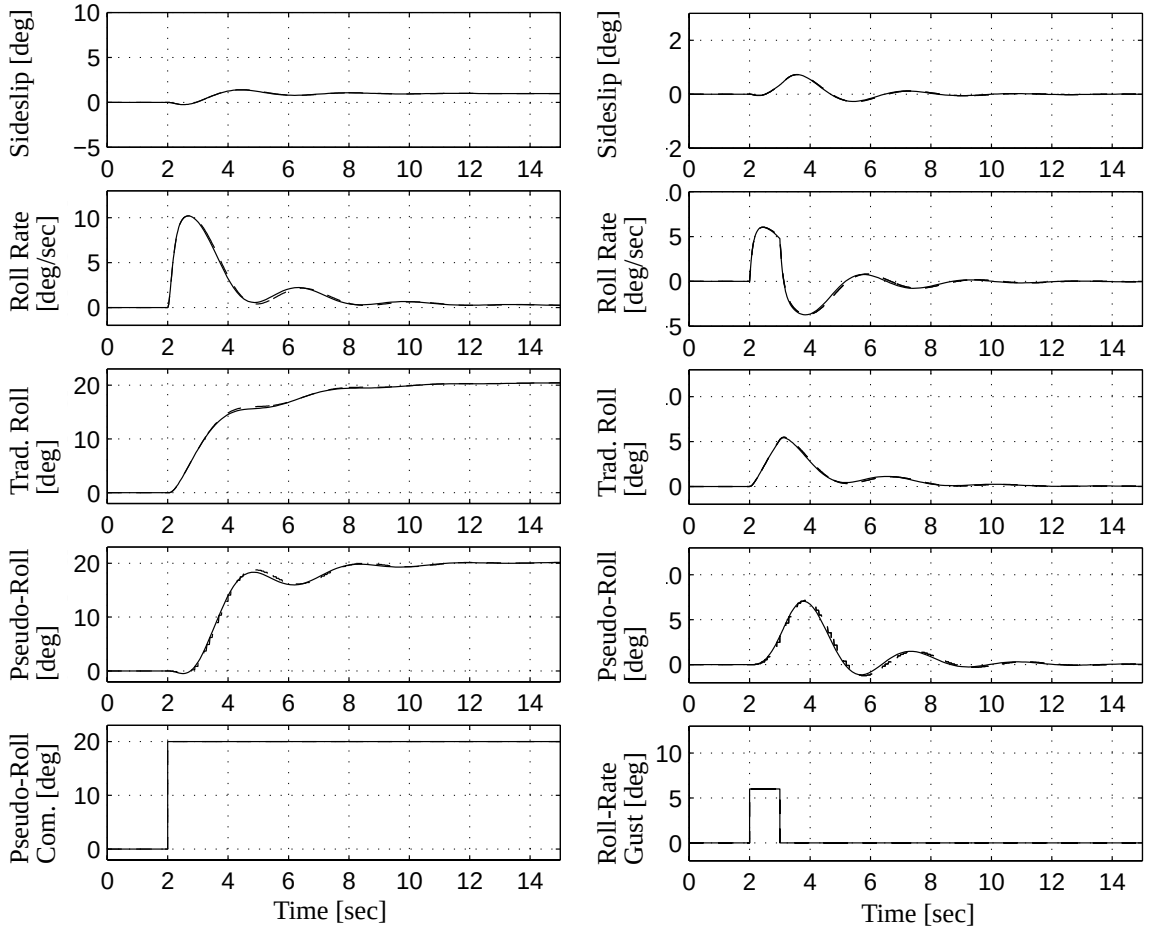

Figure 5.7: Simulated Time Responses of Pseudo-Roll Command Loop to Step Command and Gust Inputs

Figure 5.7(a) shows closed-loop time responses to a pseudo-roll command input of 20 deg. Figure 5.7(b) shows closed-loop time responses to a 5 deg/sec roll rate gust input (as they arise when a gradient exists in the vertical gust velocity distribution over the wing

spanwidth). Shown are the sideslip angle, roll rate, traditional roll angle, filtered pseudo-roll, and pseudo-roll command or roll rate gust input, respectively. Due to the control augmentation, the flight uncoordination is limited to less than 2 deg of sideslip. Pseudo-roll largely follows traditional roll angle with the overshoot less pronounced than in the open-loop case. Similar to the open-loop case, it lags by approximately 0.5 seconds. Finally, in all the plots, the continuous-time and discrete-time (in dashed lines) responses are indistinguishable.

5.4.2 Longitudinal Loop Closure

Figure 5.8 shows the bode and root locus plots of a flight path angle command loop. The loop is obtained by feeding back flight path angle and differencing it with the commanded flight path angle. The difference may be filtered, if necessary, and multiplied by a gain to produce the elevator command necessary to minimize the deviation. For comparison, Figure 5.9 shows the bode and root locus plots of a traditional pitch command loop. In both figures, a gain of -0.4422 was used to obtain the closed loop poles represented by '*' in the root locus plots. No stability or control augmentation was applied. The 'o' and 'x' represent the position of the open-loop zeros and poles, respectively.

The root locus plots indicate that both flight path angle and pitch angle feedback provide *path damping* (i.e. phugoid damping) and allow for the control of the aircraft trajectory. At the same time, both feedbacks destabilize the short period mode. Similar to the lateral case discussed above, the velocity based flight path angle feedback leads, for the aircraft and flight condition considered here, to a *non-minimum phase zero* which typically reduces the achievable bandwidth. Since for this case the zero is far from the imaginary axis, its effects on the achievable bandwidth are relatively benign. Pitch feedback, on the other hand, leads to a minimum phase transfer function.

Examination of the bode plots reveal a similar behavior. A comparison of the bode plots show that the pitch rate feedback allows for a slightly higher bandwidth if no additional compensation is used. For flight path angle feedback, a crossover frequency of

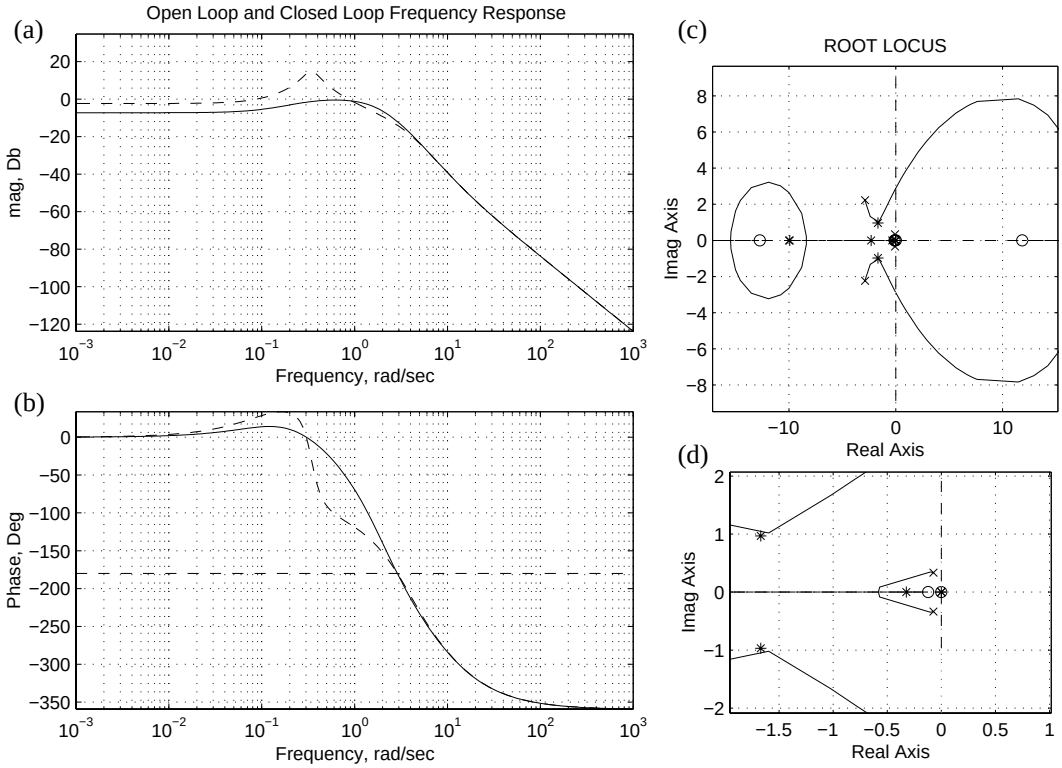

Figure 5.8: Bode and Root Locus Plots of Flight Path Angle Command Loop

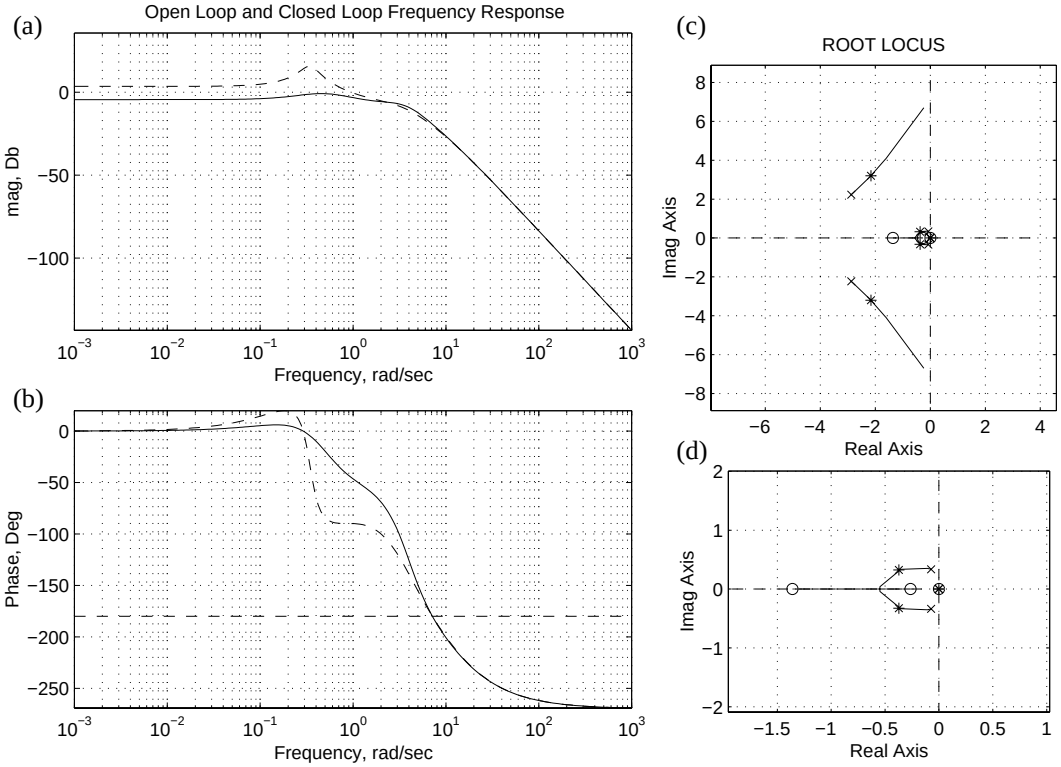

Figure 5.9: Bode and Root Locus Plots of Pitch Command Loop

approximately 1 rad/sec and a phase margin of 65 deg can be obtained, while for pitch angle feedback a crossover frequency of 2 rad/sec and a phase margin of 80 deg is achievable. Pitch feedback allows for a higher bandwidth system since the pitch angle *effects* a velocity vector change, while flight path angle is a *measure* of this change, and constitutes thus an outer-loop variable.

The attainable bandwidth of flight path angle feedback seems sufficient in the case illustrated here. In some instances, the achievable bandwidth may not suffice and lead to pilot induced oscillation (PIO). Typically, additional compensation in form of a pitch damper increases the bandwidth in these cases, as discussed in Section 5.3.2.

A time domain response of the flight path angle command loop is presented in the next section in the context of an autopilot case example

5.5 Case Example: Pseudo-Attitude Based Autopilot

This section illustrates the use of pseudo-attitude in an autopilot logic for approach conditions. This autopilot development served also as the basis for the autopilot design used in the flight tests described in Chapter 8 and 9.

The autopilot logic uses the lateral and longitudinal deviations from the desired flight path and its rate of change to form a pseudo-attitude command which is fed to the respective pseudo-attitude command loops discussed earlier. Figure 5.10 shows the lateral and longitudinal autopilot laws. They were designed using classical design methods and simulations.

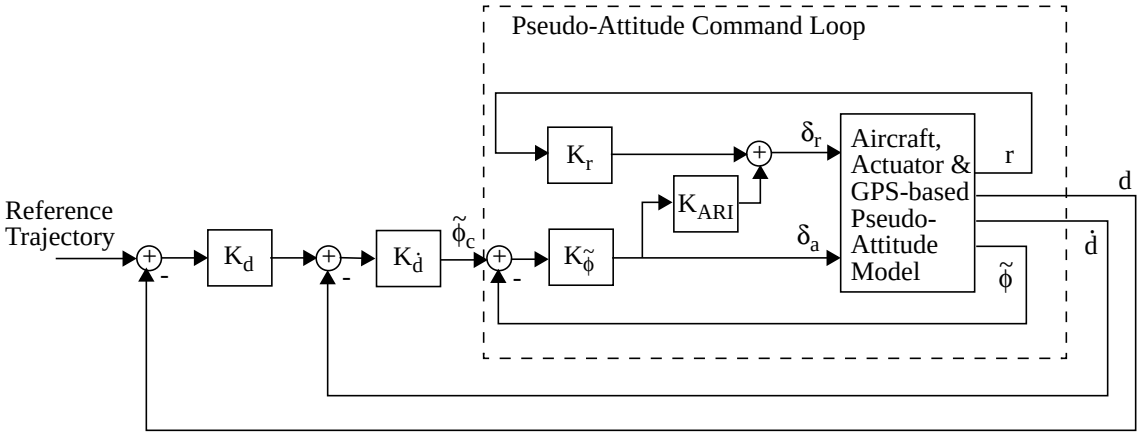

(a)

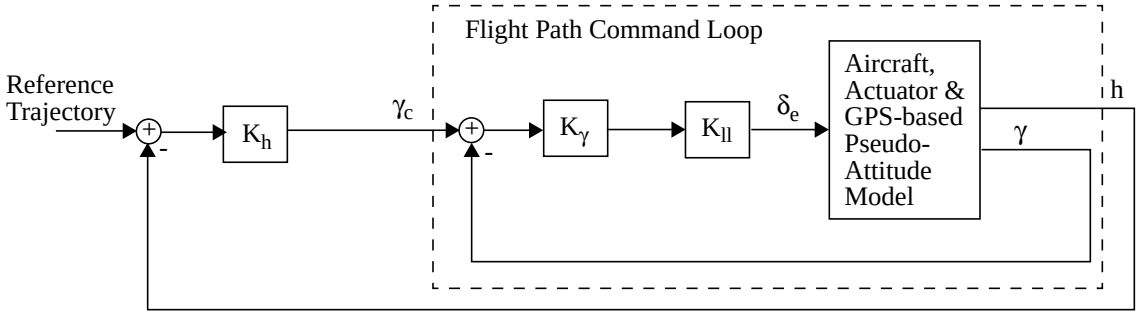

(b)

Figure 5.10: Lateral and Longitudinal Autopilot Logic

Lateral Autopilot Logic

The pseudo-roll command loop presented in Section 5.4.1 is used to build the lateral autopilot guidance logic. It includes a yaw-damper and an Aileron-to-Rudder Interconnect (ARI) to eliminate the effects of uncoordinated flight. The lateral deviation rate from the desired flight path, denoted $\dot{d}$, is calculated from the lateral aircraft states, $\mathbf{x}_{Lat}$, given in Eq. (5.4) as

$$\dot{d} = v + U_o \cos\gamma_o \psi \tag{5.7}$$

The lateral autopilot logic is then given by:

$$AP_{lat} = K_{\tilde{\phi}} [K_{\dot{d}} (-K_d d - \dot{d}) - \tilde{\phi}] \tag{5.8}$$

The particular gains used for the lateral guidance law are $K_{\hat{\phi}} = 0.15$, $K_d = 0.18$ and $K_{\dot{d}} = 0.02$.

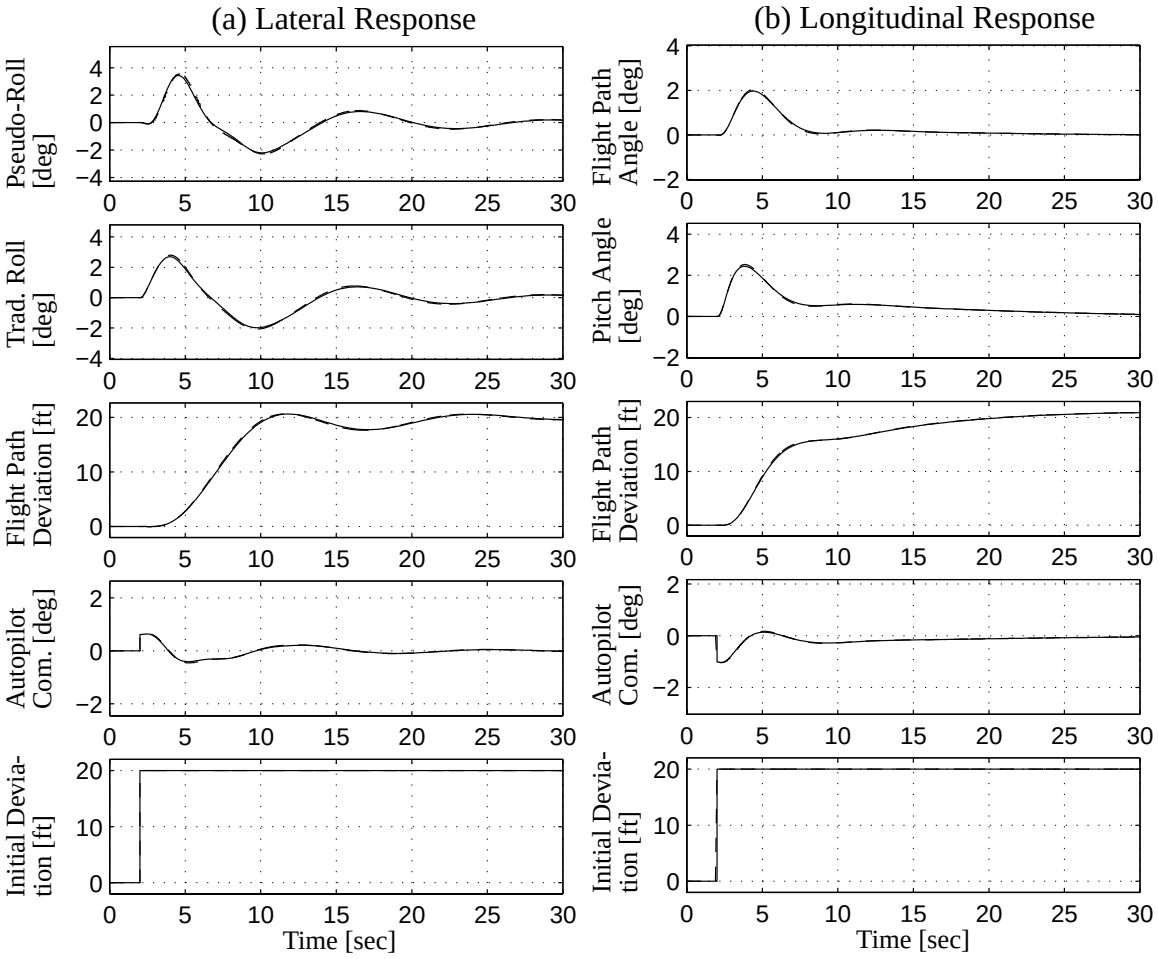

Figure 5.11: Autopilot Response to (a) Lateral and (b) Longitudinal Displacements

Figure 5.11(a) shows the lateral aircraft responses to an initial 20 ft step deviation. Shown are the pseudo-roll and traditional roll response, the actual lateral deviation, the actuator command generated by the autopilot, and the initial step input, in that order. Also plotted in the figure are the time responses of the discrete-time simulation. Due to close correspondence with the continuous-time responses they are not discernible in the figure. Figure 5.11(a) indicates that an adequate response time is achievable using the pseudo-attitude command loop. Finally, pseudo-roll largely follows traditional roll with a slight delay, as expected.

Longitudinal Autopilot Logic

Figure 5.10(b) shows the longitudinal autopilot logic. It is based on the flight path angle command loop discussed in Section 5.4.2. An additional lag-lead compensator, denoted K_{ll}, has been included in the flight path angle command path to improve the low-frequency gain deficiency which is visible in the bode plot of Figure 5.8. The resulting actuator command is then given by:

$$AP_{lon} = K_{ll} K_{\gamma}(-K_h h - \gamma) \tag{5.9}$$

The particular gains used for the longitudinal guidance law are $K_h = 0.002$ and $K_{\gamma} = -0.44$, and the lag-lead compensator was chosen as $K_{ll} = (s+0.1)/(s+0.001)$.

Figure 5.11(b) shows the time response of the longitudinal guidance law to an initial vertical displacement of 20 ft from the desired trajectory. It shows the flight path angle and pitch angle response, the vertical distance to desired trajectory, the autopilot command to the elevator actuator, and the initial deviation, in that order. The discrete-time simulation results in the same time responses and they are therefore not visible in the plots.† As can be seen, an adequate time response is obtainable using this guidance law.

5.6 Chapter Summary

This chapter illustrated that it is possible to close the flight control loop around GPS-based velocity information and achieve adequate control performance to safely control the aircraft.

If necessary, adequate aircraft response and increased bandwidth are achievable using stability and control augmentation systems, such as pitch and yaw damper. These can be implemented by feeding back high-frequency inertial quantities obtained from low grade inertial sensors.

---

† A discrete-time lag-lead of $H(z) = (1-0.99z^{-1})/(1-0.9999z^{-1})$ was used for this simulation.

Chapter 6

Flight Test Setup

A flight test system was implemented to demonstrate the single-antenna GPS-based pseudo-attitude concept and some of its applications. The tests were primarily conducted in the single-engine four-seat Piper Arrow aircraft shown in Figure 6.1.

Figure 6.1: Flight Test Aircraft

The objectives of the flight tests documented in this thesis were as follows:

- Demonstration of the single-antenna GPS-based pseudo-attitude synthesis and the comparison with traditional attitude: This was intended to serve as a proof-of-concept of the GPS-based pseudo-attitude synthesis.

- Experimental evaluation of the pilot usability of the pseudo-attitude system for aircraft control: This evaluation aimed at demonstrating the ability of the pilot to fly the pseudo-attitude display and close the loop around GPS-based pseudo-attitude.

Contingent on the successful achievement of these objectives the following additional objectives were considered:

- Flight demonstration of a pseudo-attitude based autopilot logic: This was intended to demonstrate the feasibility of an autopilot system that relied on pseudo-attitude and therefore only on information obtained from a single-antenna GPS receiver.

- Flight demonstration of a pseudo-attitude based tunnel-in-the-sky trajectory guidance system: This aimed at demonstrating the feasibility and evaluating the performance of a perspective trajectory guidance system which relied entirely on information obtained from a single-antenna GPS receiver.

The implementation of the flight test system hardware and software are discussed in Section 6.1. This is followed by a description of initial ground and flight tests of the flight test system in Section 6.2. Finally, Section 6.3 gives an overview on the flight tests performed.

6.1 Flight Test System

The primary purpose of the flight test system was the generation and display of GPS velocity vector based pseudo-attitude in real-time. In addition, the flight test system had to accommodate the real-time execution and display of pseudo-attitude based autopilot and guidance functions. Additional sensors were included to obtain a measurement of the complete aircraft state as a reference. This was useful for in-flight and post-flight evaluation. A sampling frequency and display update rate of 10 Hz were chosen to satisfy the controllability and perceptual requirements outlined in Section 3.2.3.

To accommodate changing flight test objectives, the flight test system was designed in a modular manner. This allowed for easy addition or removal of sensors to meet the needs of a particular flight test. In addition, the flight test system was designed to be entirely self-contained and independent from aircraft systems in order to minimize the potential for interference with the latter. The next two sections present the flight test system hardware and software in more detail.

6.1.1 Flight Test System Hardware

The flight test system was implemented using a modular design in order to adapted to the different flight test objectives. A *full configuration* of the flight test system, including a GPS receiver and additional sensors serving as a reference, was necessary to meet the first flight test objective which aimed at the evaluation of the pseudo-attitude synthesis and its comparison to traditional attitude. A disadvantage of this configuration, however, was the cumbersome installation it required. A scaled-down, *portable configuration* was implemented to meet the other flight test objectives. This configuration centered around a GPS receiver only and allowed for an easy and fast installation. Both configurations are presented next.

Full Configuration Flight Test System

A block diagram of the full configuration flight test system is shown in Figure 6.2. The elements of the system were the instrumentation, a laptop computer and a display, and the necessary power supply.

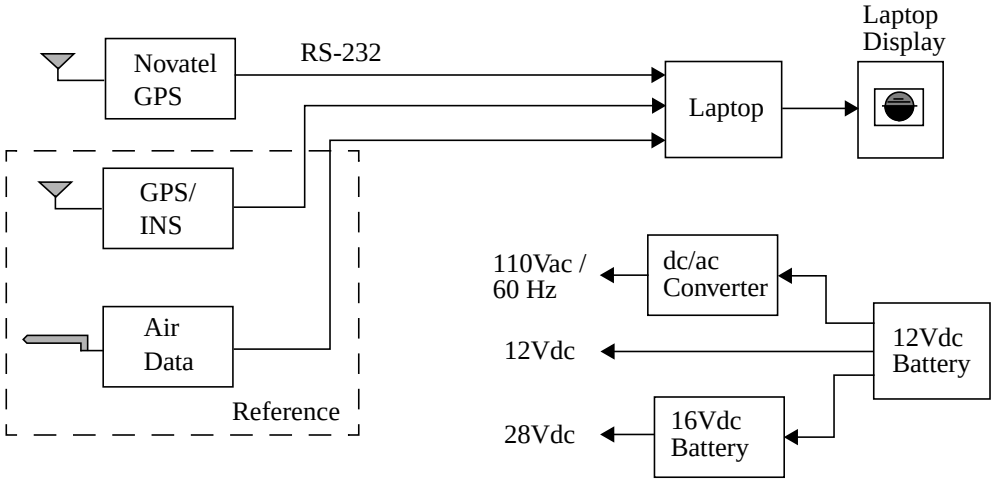

Figure 6.2: Block Diagram of Full Configuration Flight Test System

The instrumentation consisted of a Novatel 3151R GPS receiver as the primary GPS velocity source, a Rockwell C-MIGITS GPS/INS unit as a reference for traditional attitude, and a Cambridge Aero Instruments S-NAV air data unit providing air data for

post-flight analysis. Information from the GPS/INS and the air data systems were not used in the computation of pseudo-attitude but provided reference data for post-flight evaluation of the GPS-based pseudo-attitude system.

The Novatel GPS receiver was selected because its velocity update rate of 10 Hz met the display update rate requirements. The receiver had 12 parallel channels and was therefore able to track virtually all visible GPS satellites at any point in time. The receiver channels used a third-order phase-locked loop (PLL) with a carrier tracking bandwidth of 15 Hz (Novatel 1995, 1996). This was essential to track the GPS signals during high dynamic aircraft maneuvers. In addition, it enabled the measurement of velocity information up to that frequency and thus sensed all the anticipated rigid body dynamics of the aircraft. The receiver could also operate in real-time differential mode and accept differential corrections in standard RTCM format. The Novatel 3151R had an Eurocard form factor and was contained in a Novatel PowerPak chassis providing the power supply.

The GPS/INS unit used as a “truth” attitude reference was a 5-channel C/A code Miniature Integrated GPS/INS Tactical System (C-MIGITS) developed by Rockwell (now Boeing). The C-MIGITS integrated a 5-channel NavCore V C/A code GPS receiver with a GIC-100 inertial measurement unit (IMU) in a tightly coupled mechanization. The IMU used tactical grade piezoelectric multisensor technology which relied on the Coriolis effect to sense acceleration and rates. Attitude accuracy was specified by the manufacturer to be 4 mrad (1σ)† and the measured attitude was thus more than adequate to serve as a reference (Martin 1994).

Air data was measured by a compact sailplane racing S-NAV air data computer. The instrument housed pressure, temperature and flow rate transducers. A thermistor and pitot-static probe allowed measurements of outside air temperature, altitude, and indicated and true airspeed. Measurement uncertainties were estimated by the manufacturer to be 100 ft. in pressure altitude and 3-5 knots in airspeed.

---

† Heading errors were specified to be greater than 4 mrad, depending on the time since the last horizontal maneuver.

Besides the velocity data, the Novatel 3151R GPS receiver transmitted position, time, the number of satellites, the dilution of precision (DOP), and receiver status information. The GPS/INS unit transmitted traditional attitude, velocity and position data as well as receiver status information. For later flight tests, aircraft accelerations and turn rate outputs were added. The air data unit provided, barometric altitude and airspeed. Appendix G gives a detailed list of the data acquired.

The Novatel GPS velocity data and the GPS/INS attitude, position and velocity data were transmitted at an update rate of 10 Hz to a Pentium 100 MHz based laptop over standard RS-232 serial links. The air data unit transmitted at a fixed 0.5 Hz update rate. The laptop contained a Pentium 100 MHz 32 MB RAM, a 2.1 GB hard drive, five RS-232 serial ports (one internal and four based on a PC card). The laptop executed the flight test system software responsible for the synthesis and display of pseudo-attitude and pseudo-attitude based guidance information at a sample rate of 10 Hz.

A 12Vdc rechargeable sealed lead-acid battery with 80 Ampere-Hour (Ah) capacity powered the flight test system. The Novatel GPS receiver and S-NAV air data unit had an input voltage range of 10-36Vdc and 11-16Vdc, respectively, and were powered directly by the battery. The GPS/INS unit required an input voltage of 28Vdc. This was provided by connecting 8 small 2Vdc / 2.5Ah rechargeable sealed batteries in series with the main battery. To power the laptop, the 12Vdc battery voltage was transformed to 110V / 60 Hz using a dc/ac converter and fed into the laptop's adapter input. In addition, the laptop could be powered by its build-in battery.

To protect against inadvertent short circuits, the current draw from the main battery was limited using a 10A fuse. With the exception of the GPS/INS unit, the flight test system using the main battery could be operated for several hours. The operation of the GPS/INS unit was limited to approximately 1 hour of operation after which the unit automatically shut down due to overheating.

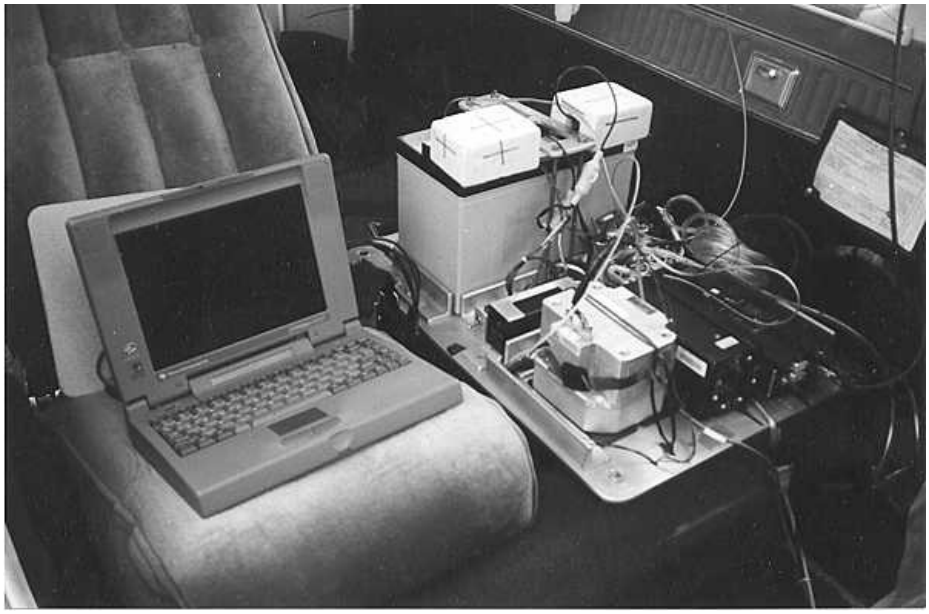

Figure 6.3: Flight Test Pallet

The flight test instrumentation and batteries were fixed to a test pallet which in turn was mounted to the aircraft structure at the location of the left back seat. Figure 6.3 shows the pallet mounted in the flight test aircraft.

The Novatel GPS and the MIGITS GPS antenna were fixed to the wind shield inside the crew cabin approximately 30 cm forward of the center of gravity. This resulted in a satellite visibility of typically 7 to 8 satellites during flight tests. The laptop screen with the pseudo-attitude display was held in front of the pilot flying the airplane close to the center of the instrument panel. Figure 6.4 shows the location of the different flight test system components in the aircraft.

A video camera was included in the flight test system to obtain a visual record of the pseudo-attitude performance.

Portable Configuration Flight Test System

For many of the flight tests, a portable version of the flight test system was employed. It consisted of the Novatel GPS receiver and the laptop only. The receiver was powered by 6 2Vdc/2.5Ah batteries in series connection, while the laptop used its internal battery. This configuration had the advantage of not requiring the cumbersome installation of the test

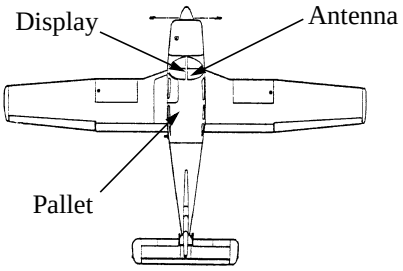

Figure 6.4: Flight Test Configuration

pallet. However, the laptop -without additional battery- was limited to approximately 1 hour of operation. Figure 6.5 shows the block diagram of the portable configuration flight test system.

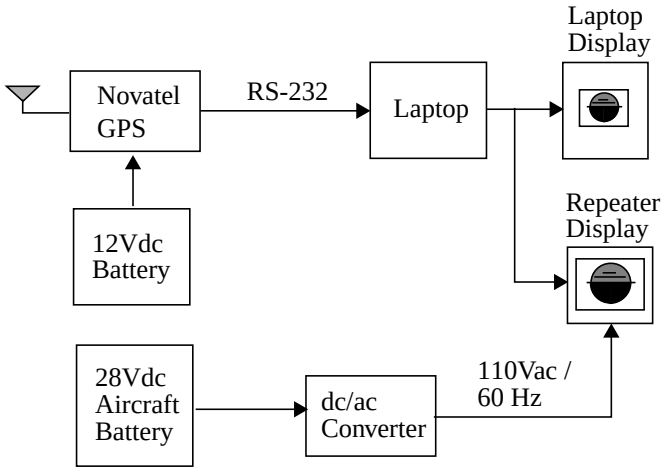

Figure 6.5: Block Diagram of Portable Configuration Flight Test System

For some of the later flight tests, a repeater display was installed in front of the pilot flying the aircraft. The display was a 10.4” Active Matrix Color Liquid Crystal Display (AMLCD) from Saphire Industrial and was operated in VGA mode (640 x 480 pixels). The display was two to three times brighter compared to traditional laptop displays and allowed, thus, for better display readability in sun light conditions. When mounted, the display obscured all the panel instruments with the exception of the airspeed indicator. The display required 110V/60Hz power input. This was supplied by transforming the aircraft’s 28Vdc cigarette lighter voltage using a dc/ac converter. Figure 6.6 shows the portable configuration flight test system and the repeater display.

Figure 6.6: Portable Configuration Flight Test System with Repeater Display

Differential GPS

The initial flight test setup called for the use of differential GPS (DGPS) for the demonstration of pseudo-attitude based *trajectory guidance*. DGPS increases the accuracy of GPS to the sub-10 meters level by eliminating Selective Availability (SA) and secondary error inducing effects such as iono- and tropospheric delays. The use of DGPS was considered in order to obtain accurate *position* data which enabled the comparison between GPS-based trajectory guidance systems and standard guidance systems such as ILS. Without differential corrections, the accuracy of GPS is better than 100 meters horizontally and 150 m vertically 95% of the time (FAA 1996).

A Northstar GPS differential beacon receiver and B-field antenna were used to obtain differential corrections from ground stations operated by the US Coast Guard. Interference caused by the aircraft's alternator, however, prevented the continuous reception of the corrections. Differential GPS was therefore *not* used for the flight demonstrations. This led to an offset between the approach flight path determined by GPS and the approach flight path indicated by the ILS. In addition, changing satellite configurations caused sometimes jumps in the GPS position solution. Implications of these effects are discussed in subsequent chapters.

It is important to note that, for the purpose of the flight demonstrations, the lack of differential corrections to the GPS data was not considered fundamental. The flight tests were intended to demonstrated the flyability and to evaluate the performance of a single-antenna GPS-based perspective guidance system, not to characterize well understood and correctable errors in the GPS data. It was assumed that future implementations of single-antenna GPS-based guidance systems would include some means of acquiring and incorporating the differential correction signal.

6.1.2 Flight Test System Software

The flight test system software was responsible for the synthesis and the display of pseudo-attitude and pseudo-attitude based guidance information at an update rate of 10 Hz. In particular, the flight test system software performed 1) the initialization of the flight test system, 2) the acquisition and the storage of the sensor data, 3) the Kalman filtering to estimate GPS-based acceleration, 4) the synthesis and display of pseudo-attitude, and 5) the calculation and display of pseudo-attitude based autopilot and trajectory guidance information†. The software provided a user interface to monitor the status of the sensors and to display custom system settings, such as the Kalman filter constants and the selected display options.

The flight test software ran in a standard Windows 95 environment. The environment is an event-based but not real-time operating system and was not designed for flight critical functions. In particular, it did not give user applications full low-level access to timing and interrupt control. However, for the purpose of flight testing the pseudo-attitude system, the capabilities of the operating system were sufficient. In addition, due to its ubiquitous nature, development tools were readily available. The actual flight test system software was programmed in C using the library routines available in Microsoft Visual C++.

---

† The displays used to represent pseudo-attitude based autopilot and aircraft trajectory guidance information will be discussed in the description of the flight tests in Chapter 8 and 9.

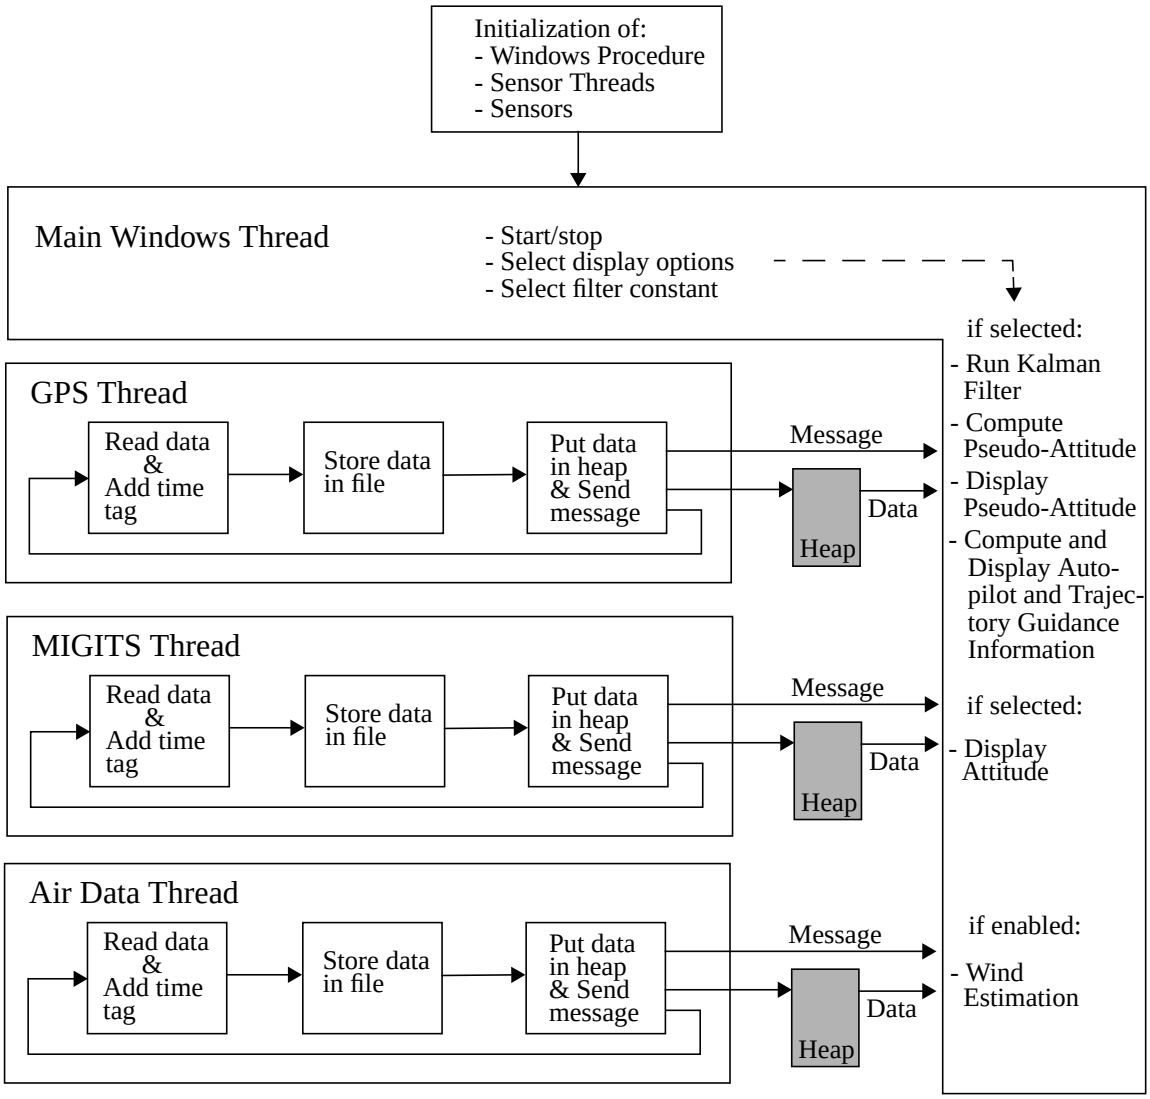

Figure 6.7: Software Architecture of Flight Test System

Figure 6.7 shows a diagram of the software architecture. The architecture was based on a multi-threaded approach and consisted of the Windows Procedure as the main thread, and of three sensor threads. The four threads were running concurrently with each other in an event-based and, thus, asynchronous manner. If the flight test system was operating in the portable configuration, the unused sensor threads (MIGITS and Air Data) were disabled.

Upon start up of the software, a series of initialization steps had to be performed. First, the RS-232 communication links were established. Next, the Windows Procedure, sensor threads and files for the sensor data were created. Finally, the MIGITS, if installed, was initialized by supplying the best estimates of position, velocity, date and time, and heading.

The main thread was responsible for the overall control of the flight instrumentation software and particularly for the processing and the display of pseudo-attitude and pseudo-attitude based guidance information. It controlled the start and the termination of the data acquisition and the selection of different user settings. User settings included the selection of the display options, the Kalman filter time constants and the display orientation.

The three sensor threads controlled the RS-232 ports interfacing with Novatel GPS receiver, MIGITS and the S-NAV air data unit, respectively. They were responsible for the acquisition and storage of the data from the respective sensors. The three sensor threads were similar in structure and only differed in the formatting applied to the incoming data. The threads were monitoring the input port for data received over the RS-232 links. The incoming data was read and time tagged to allow for the synchronization of different sensor data. The data was then reformatted to obtain common units across the different sensor data, and stored in files on the hard drive for post-processing purposes. Next, the data was put on a memory heap and a message from the respective sensor thread was sent to the main thread.

Upon receiving the respective message, the main thread retrieved the data from the heap and initiated a series of functions. Figure 6.8 summarizes the different functions and displays options available and depicts the information flow necessary to drive them.

In the default case of data retrieved from the Novatel GPS sensor thread, the Kalman filter and the pseudo-attitude synthesizing functions (described in Eq. (3.30) - Eq. (3.37), Eq. (4.23) - Eq. (4.27) and in Appendix C.4) were called, and the pseudo-attitude display was updated. Figure 6.9 shows the pseudo-attitude display as it appeared on the screen.

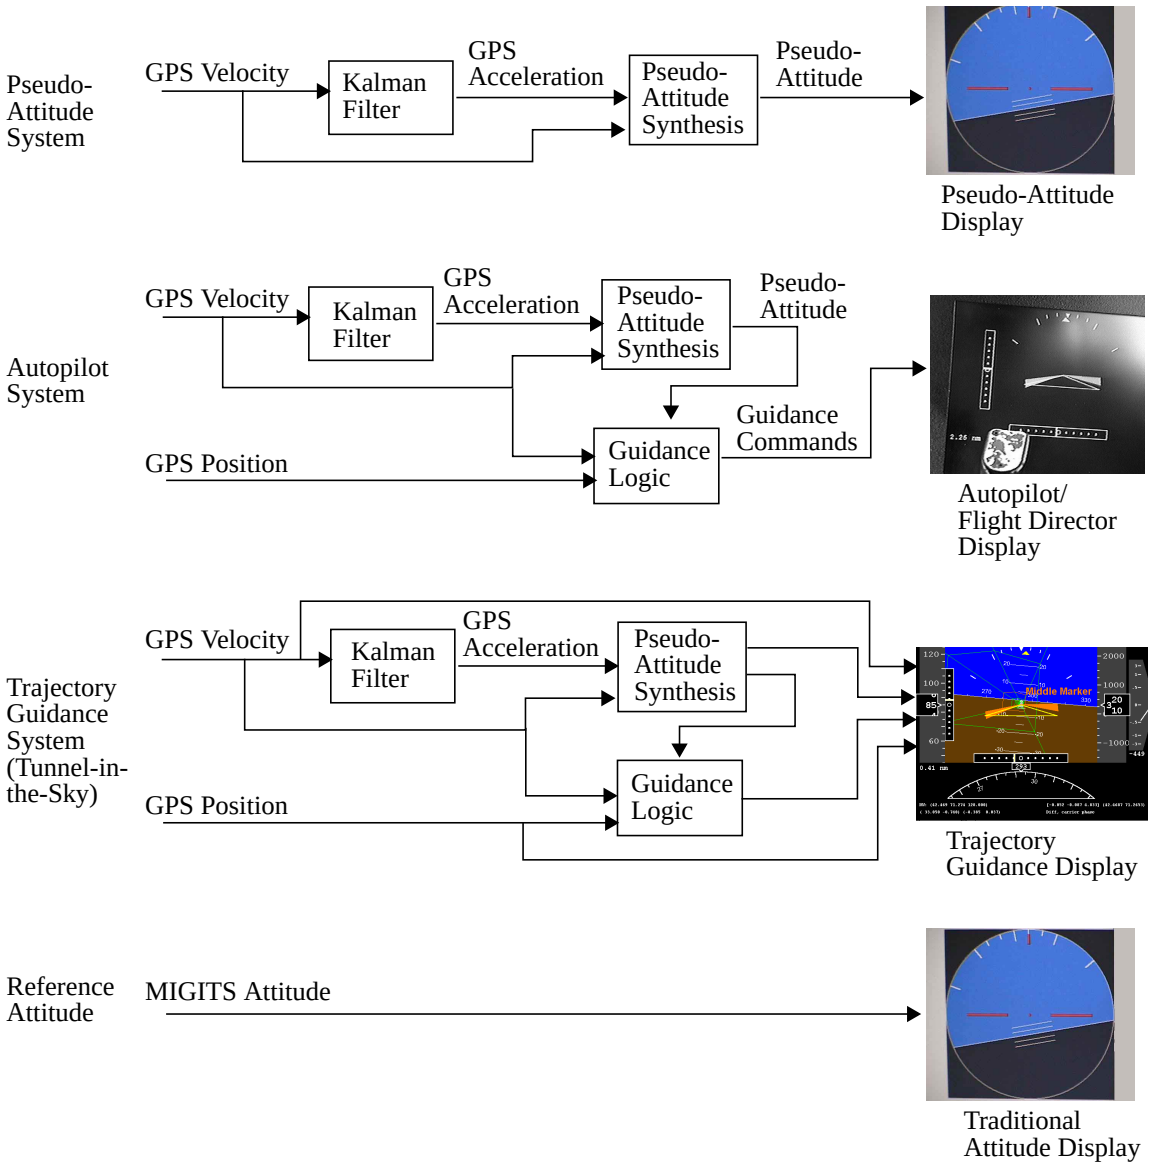

Figure 6.8: Summary of Flight Test System Functions and Displays

For flight tests investigating pseudo-attitude based autopilot and trajectory guidance logic, additional functions calculating and displaying the control and guidance information were executed.

In cases where the MIGITS was installed and selected as the display source, the main thread retrieved and displayed traditional attitude data from the heap upon receiving the respective message. Finally, if the air data unit was installed, the main thread retrieved air data from the heap and, if enabled, computed an in-flight wind estimate.

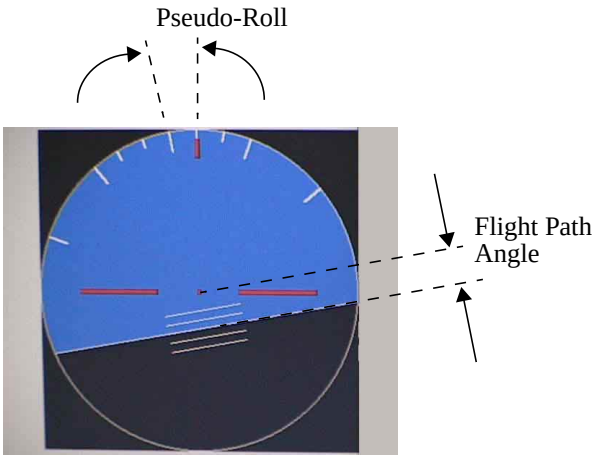

Figure 6.9: Pseudo-Attitude Display

6.2 Initial Ground and Flight Tests

Initial static tests were performed to check for the proper operation of the sensors and software. In addition, Novatel GPS velocity data was acquired under static conditions to characterize noise performance of the receiver. Next, automobile tests were conducted to test the system in a dynamic environment. For that purpose, the fully equipped test pallet was put in a car and the GPS antennas and pitot-static tube were fixed to the roof. Driving curved highway sections served as a simulation of wide aircraft turns and the proper dynamic behavior of the sensor data was observed. As expected, pseudo-roll angle indications were observed during these maneuvers. In a final ground test, the test pallet was mounted into the flight test aircraft to check the installation and the satellite visibility.

In a separate development, the acceleration estimating Kalman filter gains were initially tuned using the simulations outlined in Section 4.4 and shown in Figure 4.8. The filter gains obtained this way were subsequently fine tuned in a preliminary flight test to best satisfy pilot acceptance. The fine tuning was tied to a trade-off between display latency and perceptual requirements. The former required a minimum of filtering to reduce the latency in the display indication, whereas the latter called for enough filtering to obtain a smooth display indication.

Similarly, the gains for the autopilot and trajectory guidance system were initially tuned using a Cessna 182 simulator. The gains were then adjusted during an initial flight test.

6.3 Summary of Flight Tests

The flight test location was Hanscom Airport located in Bedford, MA. Two flight test series were conducted to achieve the four flight test objectives stated at the beginning of this chapter and shown in the first column of Table 6.1.

The first flight test series comprised 6 test flights and addressed the first two objectives. The flight tests consisted of simple coordinated and uncoordinated maneuvers. and allowed the comparison of pseudo-attitude and traditional attitude and the evaluation of pilot usability under dynamic conditions. The majority of tests were flown in a Piper Arrow using both the full and the portable flight test system configuration. Additional aircraft used were a Piper Warrior, a Mooney M20E and Cessna 310.

The second flight test series was made up of 3 flights all of which were flown in a Piper Arrow using the portable configuration of the flight test system and the repeater display. These tests were conducted under approach conditions and served to meet the last three objectives, namely the evaluation pilot usability of pseudo-attitude under approach conditions, and the demonstration of pseudo-attitude based autopilot control and trajectory guidance. Much of the work associated with the second flight test series was performed by Amonlirdviman (1998) and Walker (1998) in close collaboration with the author of this thesis.†

Table 6.1 gives an overview of the flight test objectives, the maneuvers performed to achieve them, the number of subjects pilot which performed the tests, and the type of aircraft and flight test system configuration used.

---

† The work was part of an undergraduate project jointly supervised by the author of this thesis and Dr. R. John Hansman.

Objectives	Flight Test Series	Maneuvers	Number of Subjects	Flight Test System Configuration	Aircraft Type
Proof-of-Concept, Comparison to traditional attitude	#1	Coordinated and uncoordinated flight	2	Full	Piper Arrow
Pilot usability of pseudo-attitude	#1	Coordinated and uncoordinated flight	6	Full / Portable	Piper Arrow, Piper Warrior, Mooney, Cessna 310
	#2	Approach	1	Portable, Repeater Display	Piper Arrow
Pseudo-attitude based autopilot system	#2	Approach	2	Portable, Repeater Display	Piper Arrow
Pseudo-attitude based aircraft trajectory guidance system	#2	Approach	3	Portable, Repeater Display	Piper Arrow

Table 6.1: Summary of Test Flights

The flight tests, their specific setup and protocol, and their results are presented in subsequent chapters in greater detail. Chapter 7 covers the flight tests addressing the first two objectives, namely the evaluation of the pseudo-attitude synthesis and its comparison with traditional attitude as well as the evaluation of pilot usability of pseudo-attitude. In Chapter 8, the flight tests demonstrating pseudo-attitude based autopilot control are discussed. Finally, Chapter 9 focuses on flight tests demonstrating a pseudo-attitude based aircraft trajectory guidance system.

Chapter 7

Experimental Evaluation of Pseudo-Attitude

This chapter discusses the experimental evaluation of the pseudo-attitude system under several different flying conditions. The objectives of the experimental evaluation were twofold. The first objective was an in-flight demonstration of pseudo-attitude and a comparison between pseudo-attitude and traditional attitude. This was intended to serve as a proof-of-concept of the pseudo-attitude system.

The second objective was the assessment of pilot usability of the pseudo-attitude system. This was intended to demonstrate the ability of a range of pilots to close the flight control loop around pseudo-attitude. This ultimately demonstrated the ability of the pseudo-attitude system to act in a functionally equivalent manner as the traditional attitude indicator in the aircraft.

The flight evaluation was composed of several flight test sessions. In an initial flight test the pseudo-attitude synthesis was evaluated by comparing it to the conventional pitch and roll attitude measured by the GPS/INS unit. In this as well as in subsequent flight tests, IFR-rated pilot subjects evaluated the usability of pseudo-attitude. These flight tests comprised different coordinated and uncoordinated flight maneuvers. Section 7.1 and Section 7.2 discuss the flight test protocol and the results of the experimental flight evaluation.

In a separate flight test, pseudo-attitude was demonstrated under ILS approach conditions to illustrate the usability of pseudo-attitude for approach tasks. Section 7.3 presents the demonstration of the pseudo-attitude based ILS approach. Finally, Section 7.4 gives a discussion of the results and conclusions.

7.1 Flight Test Setup and Flight Test Protocol

The first two flight tests used the full flight test system configuration, while subsequent tests relied on the portable configuration.

The pilots which served as experimental subjects were required to have an instrument rating in order to lessen the influence of learning effects on the outcome of the evaluation. Their flight experience was recorded at the outset of the flight test.

The flight test protocol included a pre-flight briefing explaining the principle of pseudo-attitude synthesis and the flight test procedures. The entire flight test was performed under simulated IFR conditions where the pilot had her/his view obscured with a hood. Once in the air, the pilot subjects flew the pseudo-attitude display for several minutes to familiarize themselves with the display.

After familiarization, the pilot subjects performed a sequence of coordinated flight maneuvers using the traditional panel mounted mechanical attitude indicator. The sequence of maneuvers consisted of straight and level flight for 30 seconds, two shallow and two steep turns for a heading change of 90 deg, a 500 ft climb and a 500 ft descent during a 360 deg turn. Subject ratings of the traditional attitude indicator using the modified Cooper-Harper scale (see Appendix H) were obtained and served as a baseline for the comparison. The traditional attitude indicator was then covered and the pilot subjects repeated the procedure, this time using the synthesized pseudo-attitude display shown on the laptop screen held in front of them, and Cooper-Harper ratings of the pseudo-attitude system were obtained. For the purpose of this evaluation, shallow and steep turns were defined as bank angles of approximately 20 and 45 deg, respectively.

The pilot subjects also performed an uncoordinated slip maneuver with full rudder deflection while maintaining a straight flight path, using the traditional attitude indicator. Cooper-Harper ratings were obtained serving as a baseline for the comparison. After covering the traditional attitude indicator, the pilots repeated the maneuver using the synthesized pseudo-attitude display and Cooper-Harper ratings were obtained. This concluded the flight test.

7.2 Flight Test Results

Results from the flight tests are shown below, including a comparison of single-antenna GPS-based pseudo-attitude and traditional attitude measured by the MIGITS unit. In addition, the pilot subject evaluations of the pseudo-attitude display are presented. The ground track of a typical flight test sequence is shown in Figure 7.1.

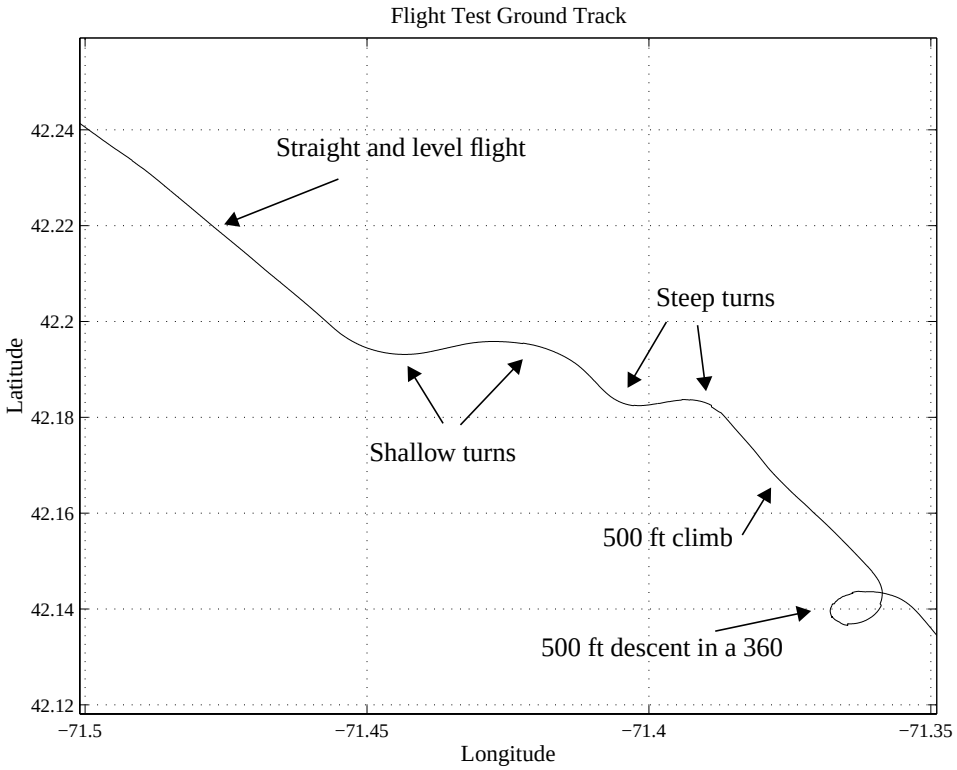

Figure 7.1: Ground Track of Flight Test Sequence

7.2.1 Comparison of Pseudo-Attitude and GPS/INS Reference Attitude

Comparisons of pseudo-attitude and traditional roll and pitch angles measured by the GPS/INS unit are shown in Figure 7.2. The data was taken in conditions of moderate turbulence with a wind magnitude of 28 knots.

Figure 7.2(a) shows the comparison of pseudo-roll and traditional roll for straight and level flight in moderate turbulence, where level flight was difficult to achieve. Figure 7.2(b) shows the comparison for shallow turns and Figure 7.2(c) for steep turns. As can be seen, the synthesized roll angle corresponds closely to the traditional roll angle.

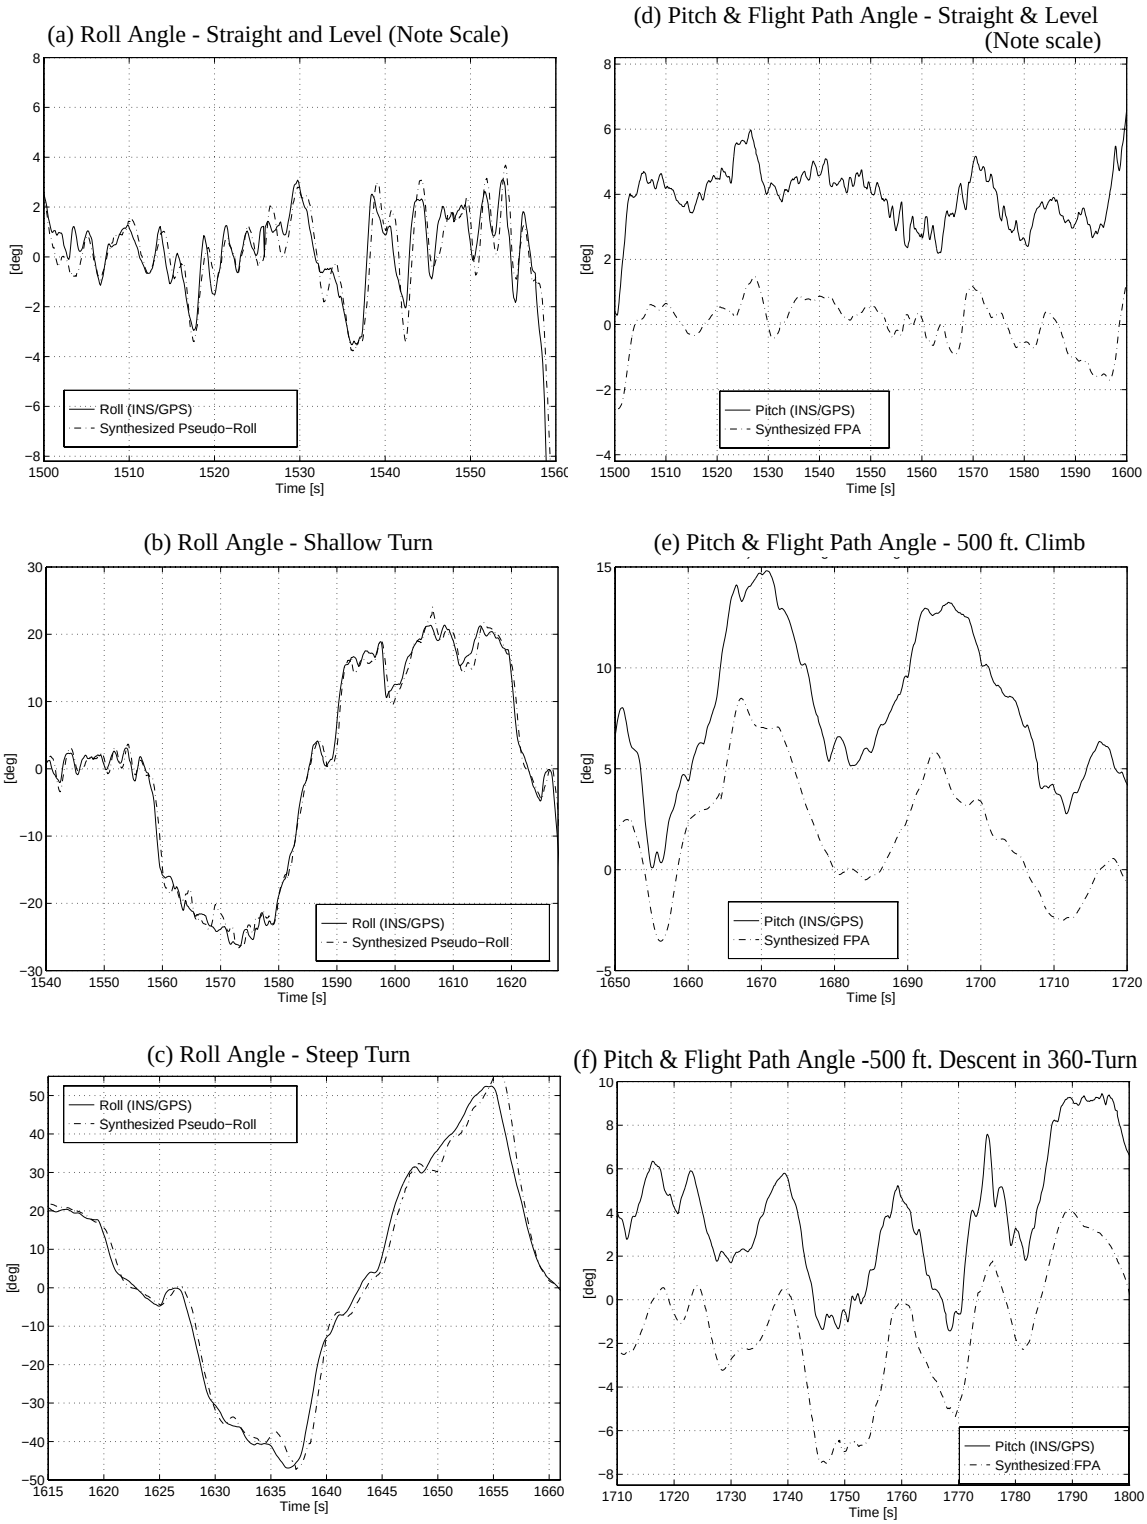

Figure 7.2: Comparison of Pseudo-Attitude and GPS/INS Reference Attitude

A slight overshoot behavior in pseudo-roll can be seen at instances where traditional roll changes abruptly. In addition, a lag of less than 0.5 seconds is observable. Both the overshoot and the lag correspond to the Kalman filter characteristics discussed previously in Chapter 4 and 5. A lower delay time and better overshoot behavior may be achievable with a different filter tuning at the expense of more noise in the estimate.

Figure 7.2(d) shows the comparison of synthesized flight path angle and reference pitch angle during straight and level flight in the presence of moderate turbulence. Figure 7.2(e) shows the comparison for a 500 ft climb and Figure 7.2(f) for a 500 ft descent during a 360 deg turn. The plotted flight path angle largely follows the pitch angle with an approximately constant offset. The difference between them is due to aircraft angle of attack. A high frequency content is observable in the pitch angle which is not present in the flight path angle. This is because pitch attitude is a control variable of higher bandwidth and is adjusted to achieve a desired flight path angle.

Figure 7.3: Comparison of Traditional Attitude and Pseudo-Attitude

Figure 7.3 shows a comparison of the traditional attitude and pseudo-attitude displays during an approximately 30 deg level turn in one of the flight tests. As can be seen, the pseudo-roll indication on the laptop screen (on the right) corresponds closely to the traditional roll indication on the panel mounted artificial horizon (on the left).

7.2.2 Subjective Evaluation of Pilot Usability of Pseudo-Attitude

The evaluation of pilot usability of the pseudo-attitude system is based on flight tests performed by six IFR-rated pilots. Their flying experience, the wind conditions during the flight tests, and the aircraft type used are listed in Table 7.1.

Pilots	Instrument Hours	Wind Conditions During Flight Test	Aircraft Type
Subject A	250	150 deg, 28 knots	Piper Arrow
Subject B	400	260 deg, 44knots	Mooney M20
Subject C	1500	260 deg, 26 knots	Piper Arrow
Subject D	150	260 deg, 26 knots	Piper Arrow
Subject E	30	310 deg, 25 knots	Piper Warrior
Subject F	110	300 deg, 20 knots	Cessna 310

Table 7.1: Summary of Subject Pilot Flight Experience, Weather Conditions and Aircraft Type Used

Cooper-Harper subjective evaluations of the pseudo-attitude system and the traditional attitude indicator for the different flight maneuvers are shown in Table 7.2. The Cooper - Harper scale ranges from 1 to 10, where 1 is the highest rating.

Pilots	Coordinated Flight		Slip	
	Attitude Indicator	Pseudo-Attitude Display	Attitude Indicator	Pseudo-Attitude Display
Subject A	2	2	-	-
Subject B	3	3	5	3
Subject C	3	roll 3, fpa 6	3	3
Subject D	3	4	2	3
Subject E	3	3	4	3
Subject F	2	2	2	1

Table 7.2: Cooper-Harper Subjective Evaluation of Pilot Usability[a]

a. The Cooper-Harper scale ranges from 1 to 10, where 1 is the highest rating.

Based on the six pilot evaluations, no statistically significant difference exists between the two displays. The examination of the individual ratings reveals that in the case of coordinated flight, the pseudo-attitude system performed equivalently to the traditional attitude indicator for the majority of pilots. It should be noted that for the evaluation of pseudo-attitude the pilot subjects had an unusual instrument scan pattern due to the position of the laptop screen held in front of them. A more detailed explanation of the modified Cooper-Harper scale is given in Appendix H.

One subject (pilot C) objected to some of the display features in the pitch ladder (insufficient space between bars on the pitch ladder) and as a consequence voluntarily separated the Cooper-Harper ratings for flight path angle and pseudo-roll.

For the slip maneuver, the pseudo-attitude display shows a more consistent rating with four out of five pilots giving it an equal or better rating. This result may be attributable to the fact that to accomplish the maneuver, the pilot simply has to track zero pseudo-roll angle in order to fly a straight flight track.

7.2.3 Additional Results

Additional flight maneuvers were performed during flight tests with the full flight test system in order to investigate the behavior of pseudo-attitude in uncoordinated flight. In particular, the effects of severe yawing and steady sideslip on pseudo-roll were demonstrated.

Figure 7.4 shows the pseudo-roll response to a severe yawing maneuver. As a comparison traditional roll angle is shown. In addition, the estimated sideslip angle†, the velocity vector heading and the aircraft heading angle are indicated.

At about 3845 seconds into the test, the pilot pushed the left rudder to the limit while flying level, thereby initiating severe yawing to the left. This led to a decrease in heading and a build up of positive sideslip angle. The latter caused a side acceleration in the negative body-y_b axis which translated into a negative pseudo-roll angle and a decreasing GPS velocity vector heading.

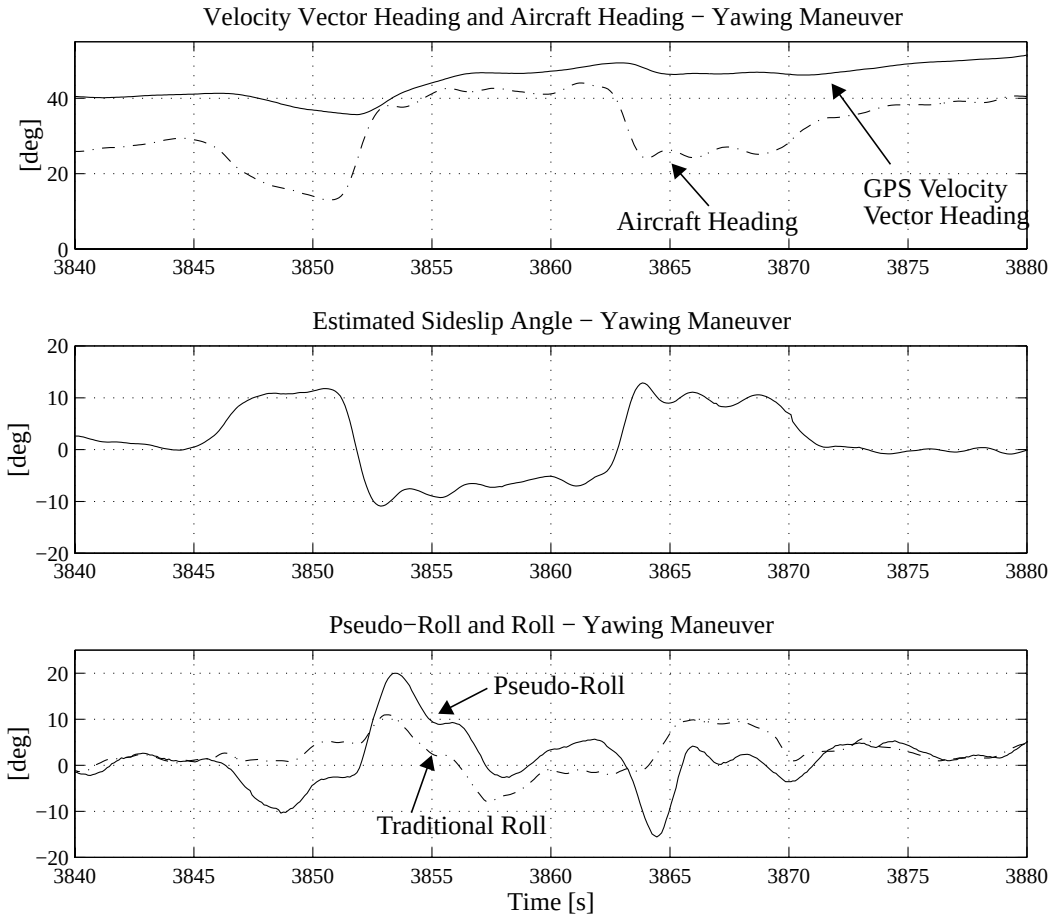

Figure 7.4: Pseudo-Roll Response to Yawing Maneuver

At 3852 seconds into the flight, the reverse process was initiated when the pilot pushed the right rudder to the limit. A sudden build-up in heading and negative sideslip angle can be observed. As a result, an acceleration in the positive body-y_b axis is apparent which led to significantly larger pseudo-roll angle indication than the actual roll angle. At the same

† The fligth test system did not include provisions to measure sideslip angle directly during the tests. Rather, it was estimated from the flight test data as follows: 1) An average wind vector was estimated by differencing the velocity vector with respect to the air from the GPS velocity vector. The former was obtained from airspeed and heading data while the aircraft was flying coordinated. The wind vector was filtered using a 3 minutes moving average filter to obtain an average wind estimate and remove high frequency components. 2) While the aircraft was flying uncoordinated, the *instantaneous* velocity vector with respect to the surrounding air was obtained by subtracting the average wind vector from the GPS velocity vector. 3) The sideslip angle was then obtained by taking the angle between the instantaneous heading angle and the direction of the velocity vector with respect to the air.

time, the pseudo-roll angle indicated the increase in GPS velocity vector heading. The pilot repeated the yawing maneuver one more time at 3862 seconds into the flight with similar results

It can be observed that for steady sideslip angles the difference between pseudo-roll and traditional roll is nearly a constant bias, while a change in sideslip, i.e. the presence of sideslip rate, induces a corresponding change in pseudo-roll. This is also in correspondence with the linearized description of pseudo-roll given in Eq. (5.6), that is

$$\tilde{\phi} = \phi + \frac{f_y}{g_o \cos\gamma_o} \tag{7.1}$$

where f_y is the specific side force, as measured, for example, by an accelerometer, γ_o is the reference flight path, and g_o, the gravitational acceleration, is 9.81 m/s^2. As can be seen the difference between traditional roll angle and pseudo-roll is due to f_y which is - to first-order - linearly related to a function of sideslip and sideslip rate.

Figure 7.5 shows the pseudo-roll response during skidding and slipping turns. At 3890 seconds into the flight, the pilot pushed the left rudder to the limit and entered a left turn. Immediately, a positive sideslip angle of 10 deg was established† which led to a nearly constant negative offset between pseudo-roll and traditional roll. At 3900 seconds into the flight, the pilot leveled off and turned immediately into a right turn while maintaining the sideslip. This caused the offset between pseudo-roll and traditional roll to change sign.

It is important to note that the presence of steady sideslip did not reduce the pilot's ability to control the aircraft using pseudo-attitude.‡

† The estimated sideslip angle, as shown in Figure 7.5 is slowly decreasing over the course of this maneuver. The author believes that this does not reflect the actual sideslip. Rather, it is thought to be an effect of the sideslip estimation.

‡ In principle, the constant offset in pseudo-roll can mentally be removed by the pilot.

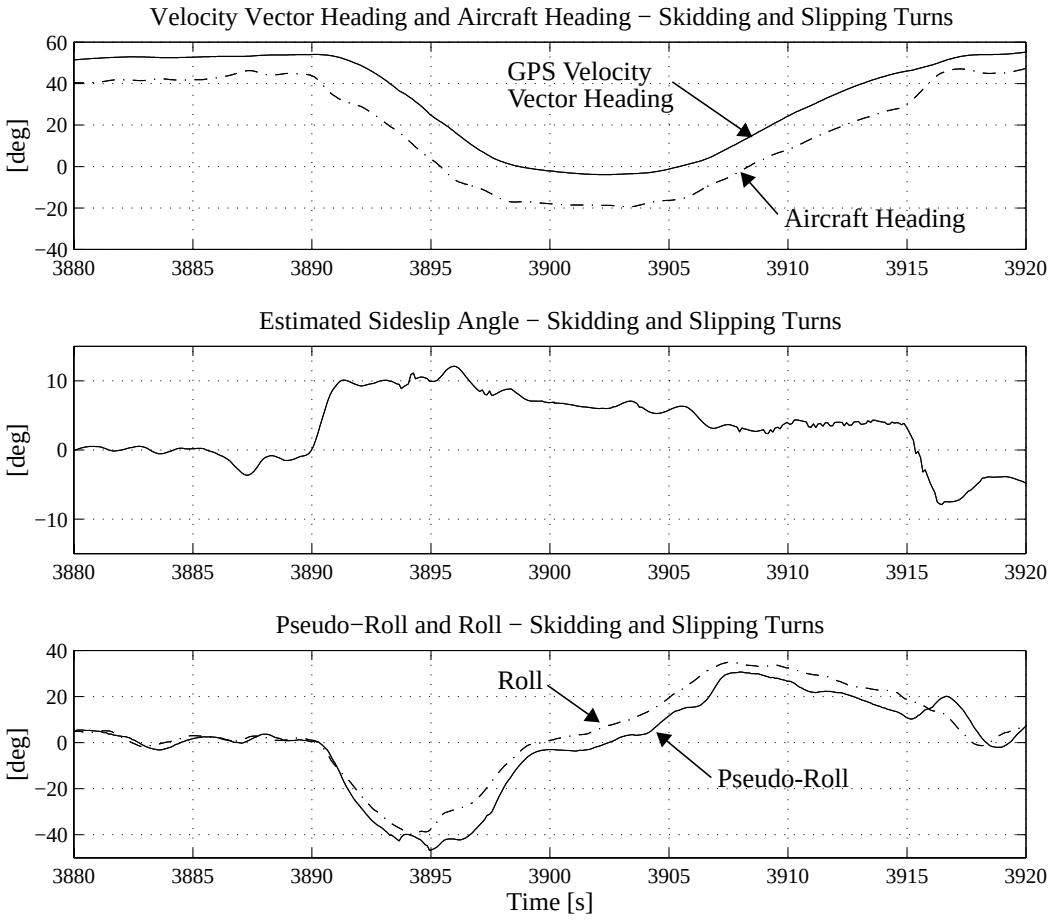

Figure 7.5: Pseudo-Roll Response During Skidding and Slipping Turns

7.3 Demonstration of Pseudo-Attitude Based ILS Approach

In order to demonstrate the usability of pseudo-attitude for approach tasks, an Instrument Landing System (ILS) approach was flown using the pseudo-attitude display instead of the traditional attitude indicator. The pseudo-attitude based ILS approach was compared to a standard ILS approach flown using the traditional attitude indicator.

7.3.1 Flight Test Setup and Flight Test Protocol

The ILS approaches were flown in a Piper Arrow using the portable flight test configuration and the 10.4” repeater display. On the repeater display the pseudo-attitude was shown. When mounted, the repeater display obscured all flight instruments with the exception of airspeed. In addition, the ILS indicator was visible to the side of the repeater display.

The preflight briefing at the outset of the flight test included an explanation of the objectives of this demonstration. The entire flight test was performed with the pilot wearing a hood thereby simulating instrument flight rules (IFR) conditions.

Once in the air, the pilot had some training on the pseudo-attitude display prior to flying the ILS approaches. After receiving clearance from air traffic control, a standard ILS approach with the traditional attitude indicator was flown which served as a baseline. After completion of the regular ILS, the pilot performed the pseudo-attitude based ILS approach. A missed approach was performed once the aircraft reached an altitude of approximately 500 ft.

The approach was videotaped to obtain a visual record of how well the aircraft was aligned with the runway. After the approach, a subjective rating of the guidance system was obtained from the pilot using the modified Cooper-Harper scale (see Appendix H).

7.3.2 Results of Flight Demonstration

One pilot with low experience completed the pseudo-attitude based ILS approach. In addition, she completed a standard ILS using traditional attitude. A summary of the pilot's flying experience and of the weather condition on the testing day is given in Table 7.3.

Pilots	Piper Arrow hours	Total hours	Instrument hours / Approaches	Wind conditions during flight
Subject A	10	200	50 / 100	350 deg, 12 knots, gust 18

Table 7.3: Summary of Subject Pilot Flight Experience and Weather Conditions

Both approaches were successfully completed. Control was never in question. Figure 7.6 shows the flight path of the two approaches to Runway 29. No substantial differences between them can be observed. Both flight paths feature slight oscillations in the lateral direction. In the vertical direction, the conventional attitude based approach path is offset from the reference path shown as a dotted line. This is primarily due to the slowly varying error induced by Selective Availability (SA). Because of its correlation time of 2-5 minutes, the observed offsets may vary during the time two approaches are

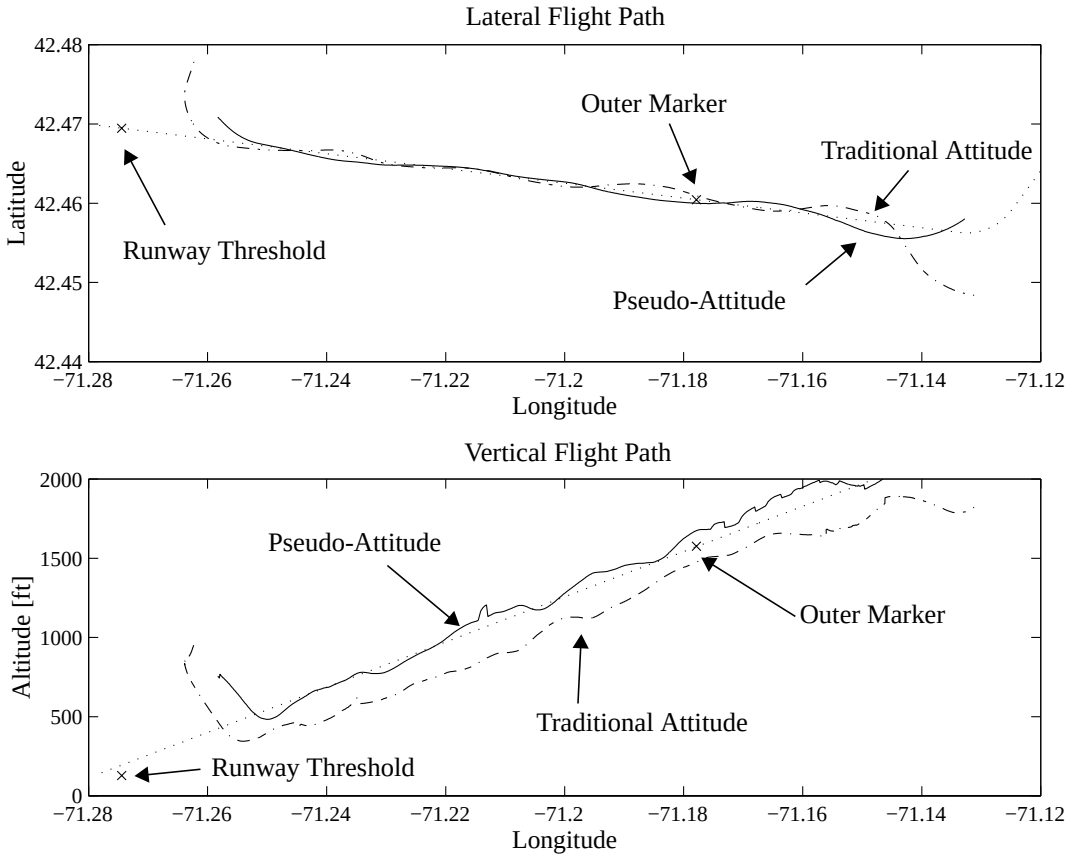

Figure 7.6: Flight Path of ILS Approaches

flown.† In addition, jumps in the vertical data of the pseudo-attitude based ILS approach are observable. The jumps were due to changes in the satellite configuration tracked by the GPS receiver which led to a new GPS position solution. The configuration changes were partly because of rising and setting satellites. Most of the time, however, they were caused by changes in aircraft attitude which brought satellites with low elevation angles into and out of the GPS antenna field of view.

Figure 7.7 shows the lateral and vertical tracking errors in more detail. For comparison, a 100 ft lateral and 70 ft vertical tunnel window is shown. No substantial differences between the pilot tracking performance of the two ILS approaches are

† However, it is assumed here that during the time an approach is flown, the error does not change singnificantly.

observable. Large lateral deviations which extend up to four times the tunnel width are apparent. In the vertical direction, the variation for both approaches are comparable to the tunnel height.

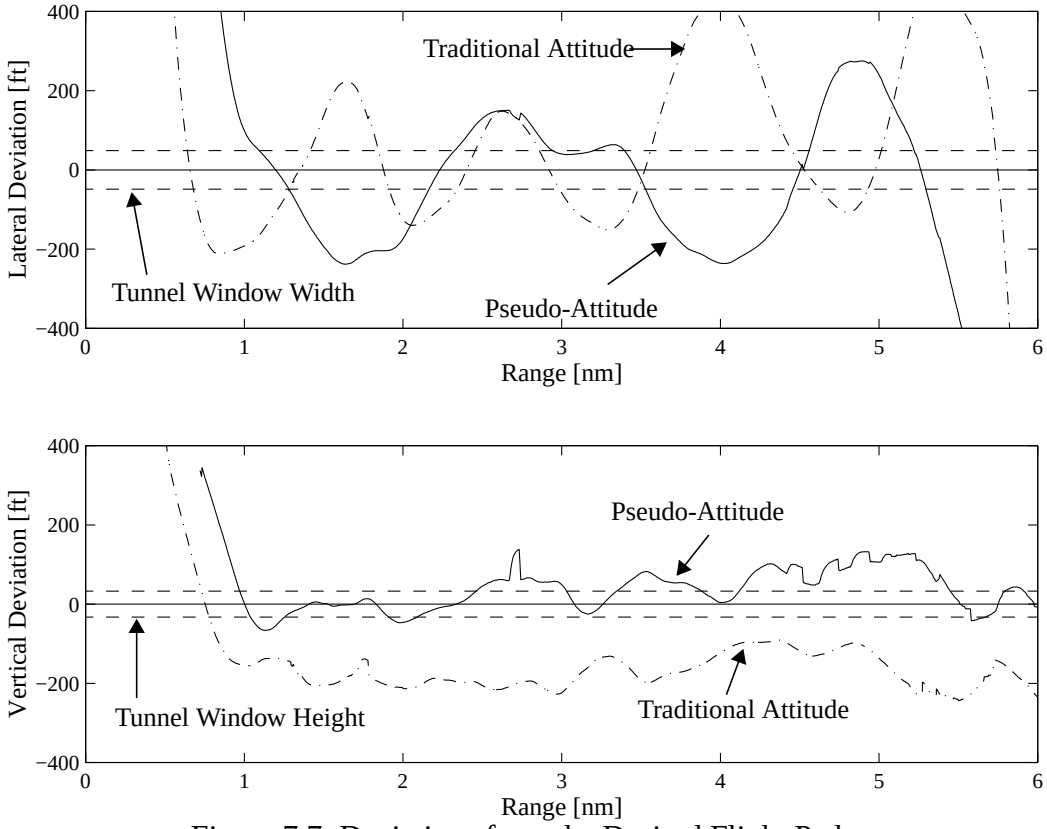

Figure 7.7: Deviations from the Desired Flight Path

Table 7.4 indicates the standard deviation (σ) and peak-to-peak (Pp) deviation of the flight performance of the two approaches. For this evaluation, 50 seconds of data (512 samples) were selected from each approach in a common altitude range which extended from 800 ft to 1500 ft. The data was filtered using a 4th-order Butterworth filter with 0.4 Hz break frequency to mitigate the effects of the jumps which GPS satellite configuration changes introduced into the data. As expected, the values for the two approaches do not differ substantially from one another.†

Also given in Table 7.4 are the Cooper-Harper ratings for the traditional attitude and pseudo-attitude displays under the experienced approach conditions. The subject pilot gave the traditional attitude and the pseudo-attitude based ILS guidance systems a rating

Approach	Lateral Deviation [ft]		Vertical Deviation [ft]		Cooper - Harper
	σ	Pp	σ	Pp	
Conventional Attitude Based ILS (Baseline)	201	569	39	136	4
Pseudo-Attitude Based ILS	129	390	40	180	5

Table 7.4: Deviations and Cooper-Harper Ratings for ILS Approaches

of 4 and 5, respectively. The rating of 4 indicated that the traditional attitude based ILS guidance system had minor but annoying deficiencies, while the rating of 5 implied that the pseudo-attitude based system had moderately objectionable deficiencies.

The pilot indicated that the higher Cooper-Harper rating for the pseudo-attitude based ILS guidance system was because of transient pseudo-roll angle excursions which occurred in response to strong cross winds gusts. At the time of testing gusts of 18 knots were present.

7.4 Conclusions

Several flight tests, consisting of coordinated and uncoordinated flight maneuvers as well as of an ILS approach, were successfully flown. These flight tests indicated that pseudo-attitude is equivalent in performance to conventional attitude with no subjective or substantial objective differences between them. The results demonstrated the ability of the pilot to close the flight control loop around pseudo-attitude.

† For the statistical evaluation of the data it was assumed that the SA induced offset was constant and that its influence was eliminated by considering the variations in the data. In reality, however, SA is slowly varying with a correlation time of 2-5 minutes and may therefore appear as a flight technical error in the approaches flown using the ILS. Nonetheless, a typical SA induced one-sigma error in velocity of 0.2 m/s gives raise to a change of 10 m during the 50 seconds time interval considered for the data evaluation. This error seems not to reduce the validity of the results, considering that the ILS approaches had deviations from the desired flight path of more than 100 m.

Chapter 8

Demonstration of Pseudo-Attitude Based Flight Director / Autopilot Approach Guidance Logic

In the previous chapter, the ability of pilots to control the aircraft using pseudo-attitude was successfully demonstrated under several different flight conditions. Both subjective and objective results indicated that pseudo-attitude allows the pilot to achieve equivalent performance as traditional attitude.

While this indicated that a pilot can close the flight control loop around pseudo-attitude, it was not a priori evident that an autopilot may do so as well. The question was thus raised as to whether GPS-based pseudo-attitude could potentially be used to drive an aircraft autopilot system.

If realized, this capability would provide the larger General Aviation (GA) community with autopilot or flight director guidance since it is based solely on a readily-available single-antenna GPS receiver. Furthermore, recent developments of miniaturized unmanned aerial vehicles (micro UAV) call for a minimum of flight instrumentation in terms of mass, size and power consumption. A single-antenna GPS receiver providing information for navigation, guidance *and flight control* may, thus, be ideally suited to meet these goals.

Consequently, the objective of the flight tests discussed in this chapter was to demonstrate the *feasibility* of an autopilot system which relied on pseudo-attitude and other information obtained from a single-antenna GPS receiver.

In the following sections, the particular setup for this flight demonstration and the flight test protocol are outlined. Next, results from the flight test are presented, and a discussion and conclusions are given.

8.1 Flight Test Setup

The flight tests were conducted in the Piper Arrow shown in Figure 6.1 using the portable flight test configuration and a repeater display discussed in Chapter 6.

The flight test aircraft available was not equipped with a dual-axes autopilot nor were actuators installed. Since the purchase and installment of the necessary equipment was not feasible under the scope of this flight demonstration, a different approach was pursued. This approach treated the *pilot* as an *actuator* promptly reacting to visual autopilot commands given on the display in front of him in form of standard flight director command bars. Figure 8.1 shows the display on which the autopilot commands were conveyed to the pilot. The display depicted the aircraft symbol in form of a triangle fixed at the center of the display and flight director command bars in the shape of a shallow inverted 'V'. The latter were a visual translation of the autopilot commands. Their vertical and rotational deviation from the aircraft symbol, supplied to the pilot implicitly the flight path and pseudo-roll angle necessary to establish the aircraft on the desired trajectory. The pilot was directed to capture and maintain the desired flight path by "flying into the V-symbol" such that the command bars were aligned with the edges of the aircraft symbol.†

It was asserted that if a pilot, without obtaining any other visual or sensorial inputs, performed adequately in this pursuing task, then a pseudo-attitude based autopilot logic driving traditional actuators would be feasible.

The autopilot guidance laws were designed to allow the capture and tracking of a desired trajectory. Figure 8.2 shows the block diagrams of the autopilot logic. They were simpler in structure than the ones presented in Chapter 5 to allow for easy in-flight adjustment of the gains. The longitudinal control law blended longitudinal tracking error

† Due to an error in the programming, the pseudo-roll angle indicator and the localizer and glideslope indicators were not removed from the screen. These indicators were at the periphery of the screen, however, and the pilots indicated that they were able to ignore them and focus only on the flight director during the tests.

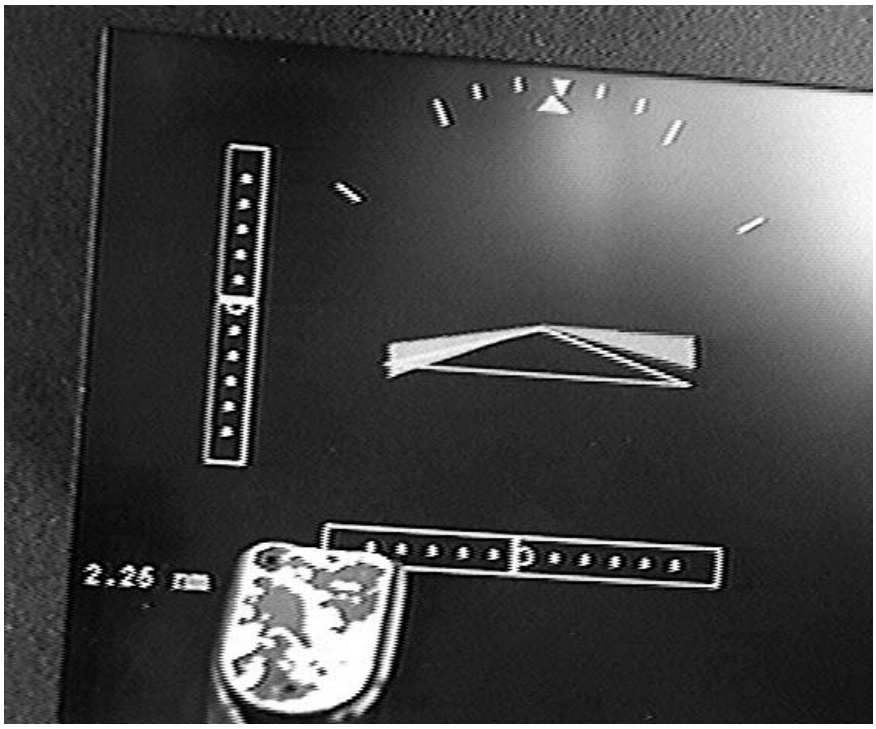

Figure 8.1: Autopilot Command Display

and flight path information to generate a longitudinal flight path command. The lateral control law combined lateral tracking error and error rate as well as pseudo-roll angle to yield a pseudo-roll command. The actual control laws were given by:

$$\begin{aligned} FD_{long} &= -K_{\gamma} \cdot (-K_h h - \gamma) \\ FD_{lat} &= (-K_d d - \dot{d}) \cdot K_{\dot{d}} - \tilde{\phi} \end{aligned} \tag{8.1}$$

where FD is the flight director command, h is the tracking error in longitudinal direction, and γ is the flight path angle. d and $\dot{d}$ are the tracking error and error rate in lateral direction, respectively, and $\tilde{\phi}$ is the pseudo-roll angle. K_h, K_{γ}, K_d, and $K_{\dot{d}}$ are the respective gains. The gains were set in initial tests using a Cessna 182 simulator and further adjusted in preliminary flight tests. The values used were as follows: $K_h = 0.002$, $K_{\gamma} = -0.4422$, $K_d = 0.16$, and $K_{\dot{d}} = 0.0162$. The guidance logic was added as a routine to the main thread, as shown in Figure 6.7.

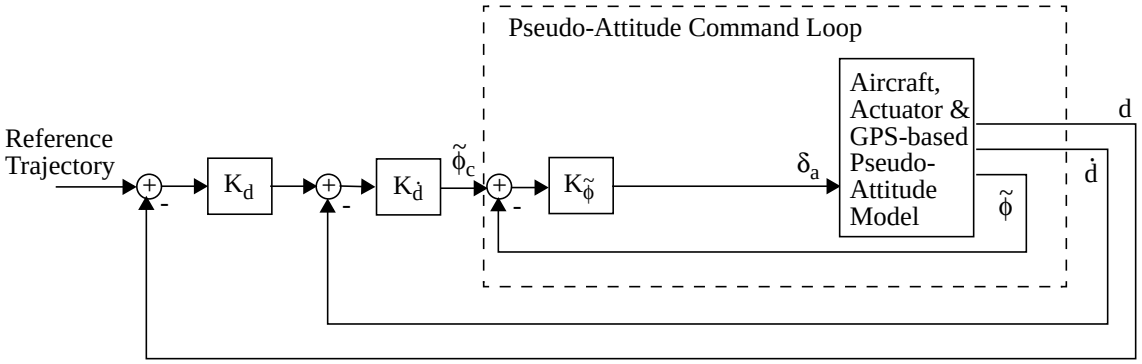

(a)

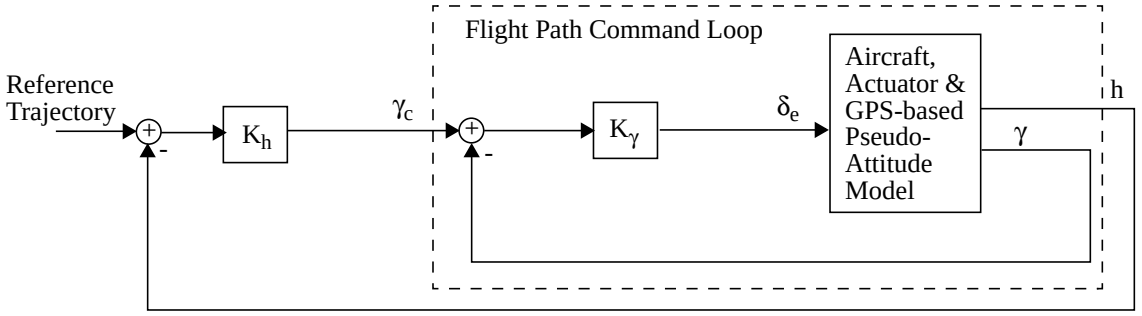

(b)

Figure 8.2: Block Diagram of Lateral and Longitudinal Autopilot Logic

8.2 Flight Test Protocol

The autopilot guidance system was demonstrated by flying approaches to Runway 29 at Hanscom Airport located in Bedford, MA. The pilots which served as experimental subjects were recruited from the MIT community and all had recent instrument flying experience. Their flying experience in the Piper Arrow aircraft, hours of instrument flying, number of instrument approaches as well as other background information were recorded.

The preflight briefing at the outset of the flight test included an explanation of the objectives of the demonstration and a presentation of the autopilot command display. The pilot was instructed to limit his focus on the autopilot command display and to follow the displayed commands accurately and promptly. The pilot wore a hood during the test flight in order to constrain his field of view to the attitude command display.

Once in the air, the pilot had some opportunity to train with the attitude command display. After completion of the familiarization the approach with the attitude command display was flown.

For each approach, clearance from air traffic control had to be obtained. All the approaches were videotaped in order to obtain a visual record of how well the aircraft was aligned with the runway. Typically, a missed approach was flown once the aircraft reached an altitude of approximately 500 ft.

8.3 Results and Discussion

Two approaches were flown with the attitude command display. Both were flown during the same flight test, one by the principal subject pilot and the second by the pilot acting previously as a safety pilot seated in the right seat. After the principal subject pilot completed the approach he acted as a safety pilot and the display was mounted in front of the pilot in the right seat. Table 8.1 shows the flying experience of the subject pilots and the wind conditions during the flight test.

Pilots	**Piper Arrow hours**	**Total hours**	**Instrument hours / Approaches**	**Wind conditions during flight**
Subject A	600	5200	250 / 600	350 deg, 12 knots, gust 18
Subject B	10	200	50 / 100	350 deg, 12 knots, gust 18

Table 8.1: Summary of Subject Pilot Flight Experience and Weather Conditions

Both approaches were successfully flown to the missed approach point. Aircraft control was never in question. Figure 8.3 shows the flight path of the two approaches and the reference flight path (dotted line). As can be seen, both subject pilots closely followed the desired flight path once passing the outer marker.

Figure 8.4 shows an out-of-the window view of the runway during the approach (of subject B) and the corresponding attitude command display.[†] The approach flight path is clearly seen to be aligned with Runway 29.

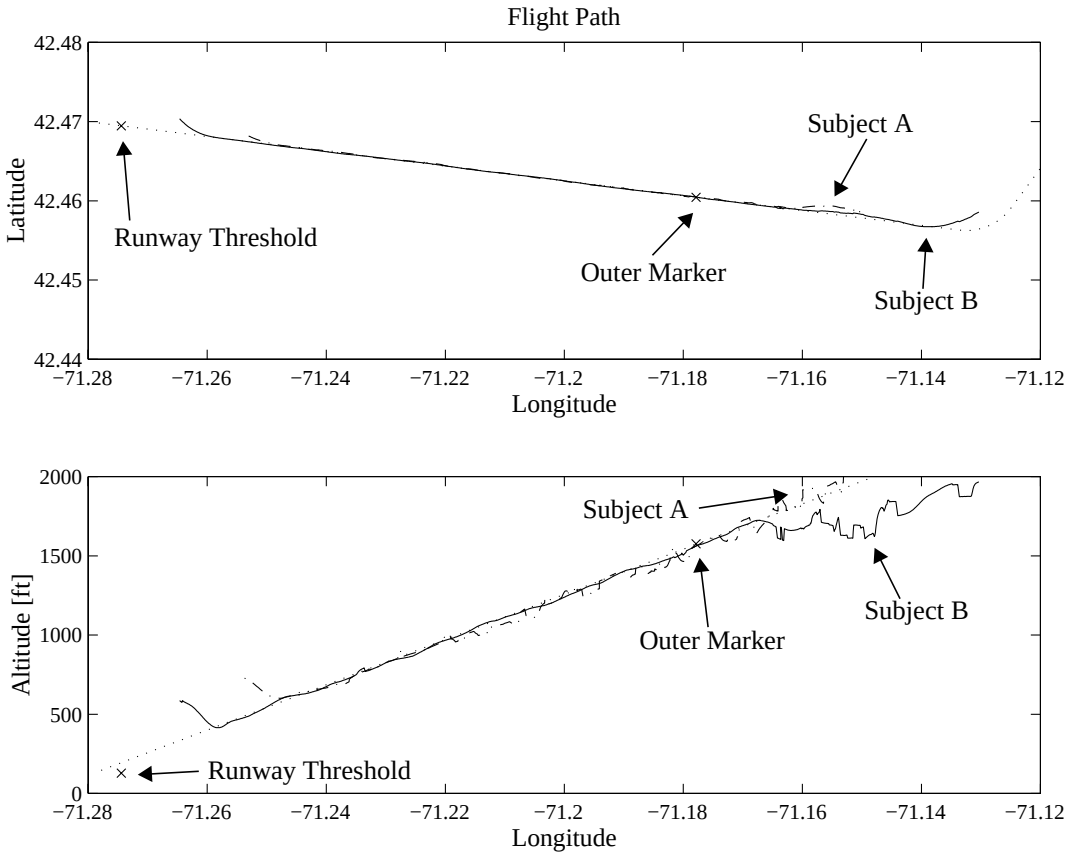

Figure 8.3: Approach Flight Path

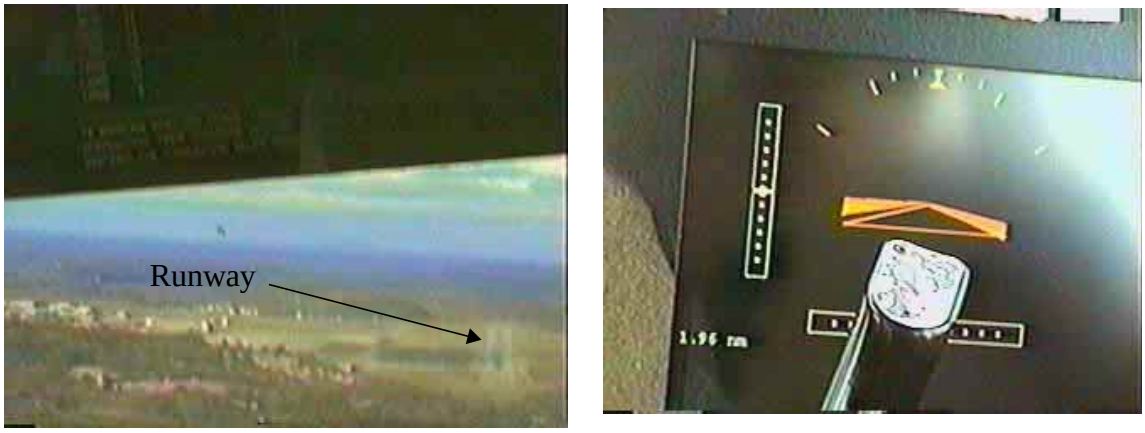

Figure 8.4: a) Out-of-the-Window View (b) Corresponding Attitude Command Display

† The two pictures are screen shots from the flight test video and are captured within three seconds of each other.

The lateral and vertical tracking performance is plotted in more detail in Figure 8.5. For comparison, a 100 ft lateral and 70 ft. vertical tunnel window is shown. It is apparent that the deviations stayed well within this ‘imaginary’ tunnel for the most part of the approaches.

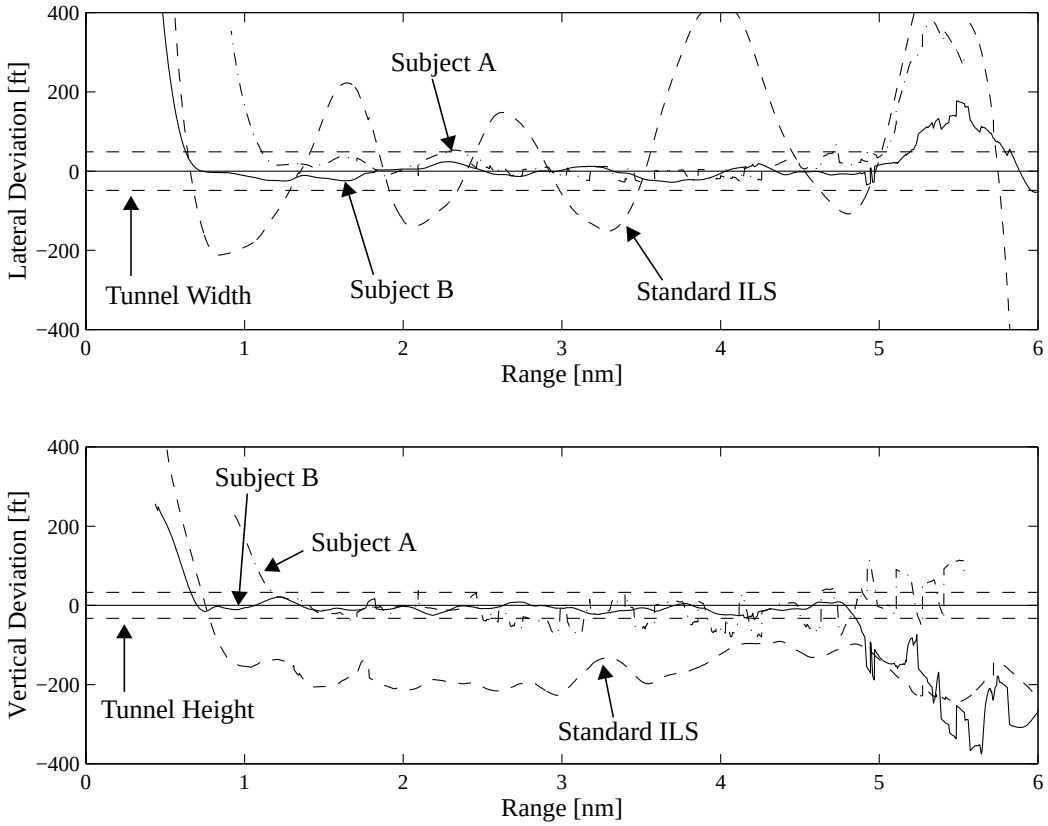

Figure 8.5: Deviations From the Desired Flight Path

In addition, the deviations of an ILS approach flown with traditional attitude (see Section 7.3) are shown in Figure 8.5. The comparison reveals two interesting differences. First, the oscillation amplitude of the ILS based approach flight path deviations in the lateral direction were in some instances more than four times larger than the corresponding tracking error obtained using the ‘human-actuator’ autopilot. The variations in vertical direction seem comparable, however. Second, it is apparent that the oscillation frequency of the tracking errors obtained with human-actuator autopilot are

higher than the oscillation frequency of deviations of the traditional attitude based ILS approach. This suggests that the human-actuator based command system results in higher bandwidth control system.

In the vertical direction, the conventional attitude based ILS approach path is offset from the reference path shown as a dotted line. This is primarily due to the slowly varying error induced by Selective Availability (SA). In addition, small jumps in the GPS data are observable in Figure 8.3 and Figure 8.5. These are particularly apparent in the altitude data and are the result of configuration changes in the satellites tracked by the GPS receiver, as discussed in the previous chapter.

Table 8.2 gives a summary of the standard deviations (σ) and peak-to-peak (Pp) values of the tracking error for the two human-actuator autopilot approaches and the ILS approach†. For this evaluation 50 seconds of data (512 samples) were selected from each approach in a common altitude range which extended from 800 ft to 1500 ft. The data was filtered using a 4th-order Butterworth filter with 0.4 Hz break frequency to mitigate the effects of the jumps which GPS satellite configuration changes introduced into the data. Both the standard deviation and the peak-to-peak value were considered.

Approach	**Lateral Deviation [ft]**		**Vertical Deviation [ft]**	
	σ	**Pp**	**σ**	**Pp**
Autopilot: Subject A	20	83	25	108
Autopilot: Subject B	14	52	8	40
Conventional Attitude Based ILS	201	569	39	136

Table 8.2: Standard Deviation and Peak-to-peak Value of Tracking Error

† For the statistical evaluation of the *ILS data* it was assumed that the SA induced offset was constant and that its influence was eliminated by considering the variations in the data. In reality, however, SA is slowly varying with a correlation time of 2-5 minutes and may therefore appear as a flight technical error in the approaches flown using the ILS. This problem did not exist for the GPS-based approaches, since SA induced errors would immediately be reflected as changes in the location of the assumed reference flight path used in the autopilot logic. Nonetheless, a typical SA induced one-sigma error in velocity of 0.2 m/s gives raise to a change of 10 m during the 50 seconds time interval considered for the data evaluation. This error seems not to reduce the validity of the results, considering that the ILS approaches had deviations from the desired flight path of more than 100 m.

8.4 Conclusions

These results as well as observations made during the flight tests indicated the feasibility of a pseudo-attitude based autopilot system which relied on pseudo-attitude and, consequently, on information entirely obtained from a single-antenna GPS receiver. Using the pseudo-attitude based autopilot system, the subject pilots achieved better tracking performance and higher control bandwidth than with traditional ILS approach guidance.

Chapter 9

Demonstration of Pseudo-Attitude Based Tunnel-in-the-Sky Trajectory Guidance Systems Using Single-Antenna GPS

A number of recent studies proposed “tunnel-in-the-sky” displays as alternatives to conventional guidance displays (Barrows 1995, Theunissen 1997). Tunnel-in-the-sky trajectory guidance systems provides the pilot with a perspective flight path display which depicts the outside world in form of a horizon and the desired flight path in form of tunnel gates. The horizon information is typically obtained from an attitude sensing instrument. Additional guidance cues, such as deviation indicators, flight director cues or trajectory predictors, may be included in these displays.

Tunnel-in-the-Sky trajectory guidance systems rely on instantaneous aircraft attitude, velocity and position information to provide trajectory guidance to the pilot. Thus far, only sensors such as Inertial Navigation Systems (INS), Attitude and Heading Reference Systems (AHRS) in conjunction with additional navigation sensors, or, more recently multi-antenna GPS receiver had the capability to provide this information. However, these sensors and, in some cases, the need to fuse the information from different sensors render these systems costly and prevent their use in General Aviation (GA) aircraft.

The availability of position, velocity and now pseudo-attitude information from a s*ingle-antenna GPS receiver,* eliminates these shortcomings and reduces the cost of these systems sufficiently to make them available to the larger GA community. An inexpensive GPS-based trajectory guidance system could provide guidance for complicated flight paths and allow precision approaches to be flown at airports that do not provide ILS capability. It hereby reduces the risks of accidents at these airports such as controlled flight into terrain (CFIT).

What fundamentally distinguishes the approach taken in this chapter from the previous implementations is the use of single-antenna GPS-based pseudo-attitude to drive the horizon on the perspective flight path display. Consequently, the *objective* of the flight tests discussed in this chapter was to demonstrate the feasibility and evaluate the performance of a perspective trajectory guidance system which relied entirely on information provided by a single-antenna GPS receiver.

In the following section, the particular setup for this flight demonstration is outlined. In Section 9.2 the flight test protocol is detailed. Section 9.3 presents the objective and subjective results of the flight tests. In Section 9.4 conclusions are given.

9.1 Flight Test Setup

The flight tests were conducted in a Piper Arrow using the portable flight test system configuration and the repeater display.

Three systems for providing approach guidance were evaluated and compared in this flight test demonstration: the standard Instrument Landing System (ILS) as a baseline, a single-antenna GPS-based system with a tunnel-in-the-sky display, and a single-antenna GPS-based system with a combined tunnel-in-the-sky and flight director display.

Standard ILS

The standard aircraft instrument panel with an ILS instrument served as the baseline for the comparison with the single-antenna GPS guidance systems. The standard aircraft instrument panel provided, among other information, conventional attitude, heading, altitude, vertical speed and airspeed indications. The standard ILS display indicated the position of the aircraft with respect to a vertical beam (glideslope) and a horizontal beam (localizer) transmitted from the runway.

Single-Antenna GPS System with Tunnel-in-the-Sky Display

The tunnel-in-the-sky display gave the pilot a perspective view of the outside world in form of a horizon and showed the desired flight trajectory in form of a tunnel. It took advantage of the three-dimensional positioning capabilities of GPS to locate the aircraft with respect to the desired flight path. In the implementation used in this flight demonstration, the horizon was driven by pseudo-attitude synthesized from GPS velocity measurements. The trajectory guidance system generated a set of tunnel 'gates' that were displayed to the pilot and represented the desired approach trajectory. In order to stay on course, the pilot had to fly through the gates.

Figure 9.1 shows the tunnel-in-the-sky display.† The tunnel was superimposed on the horizon which also featured a pseudo-roll marker, a flight path angle ladder and a ground track heading indication. The aircraft symbol was shaped like a triangle and fixed at the center of the horizon display. It indicated the pseudo-attitude of the aircraft with respect to the horizon.

The tunnel size corresponded to an area of 100 ft horizontally and 70 ft vertically. The tunnel featured a specially colored gate (orange) indicating the middle marker. Glide slope and localizer deviation indicators were shown to the left and at the bottom of the horizon, respectively. Additional display information included an altimeter tape and vertical speed indicator at the right side and a ground speed tape at the left side of the display. A separate ground track heading compass indicator and the distance to runway were provided at the bottom of the display. It should be noted that the entire display was driven by information obtained from a single-antenna GPS receiver. The flight director command bars shown in the center of Figure 9.1 were not included in the tunnel-in-the-sky display.

---

† The display partially adopted the symbology used by Barrows (1995). In particular, the tunnel gates indicated the necessary aircraft roll angle which was calculated based on assumed flight velocity and turn radius.

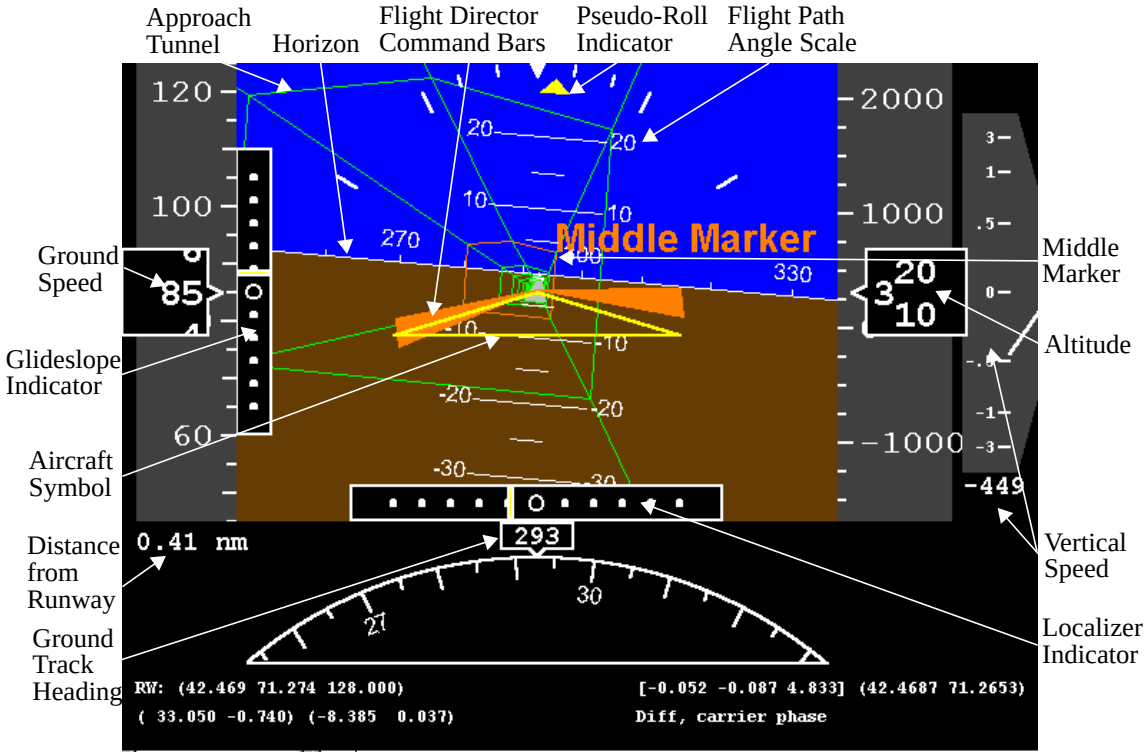

Figure 9.1: Combined Tunnel-in-the-Sky and Flight Director Display[a]

a. For the case when the tunnel-in-the-sky display was used alone, the flight director command bars were not included. The display was otherwise identical.

Single-Antenna GPS System with Combined Tunnel-in-the-Sky and Flight Director Display

This system augmented the tunnel-in-the-sky display with an additional flight director. The flight director command bars had the shape of a shallow inverted ‘V’, as shown in the center of Figure 9.1. They commanded the flight path and pseudo-roll angle necessary to establish the aircraft in the tunnel by their vertical and rotational position, respectively. The pilot captured and maintained the desired flight path by “flying into the V-symbol” such that the command bars were aligned with the edges of the aircraft symbol.

The flight director command bars were driven by the same guidance logic which was used for the demonstration of the pseudo-attitude based autopilot system, described in the previous chapter (see Eq. (8.1)).†

9.2 Flight Test Protocol

The pilots which served as experimental subjects to evaluate the GPS-based trajectory guidance systems were recruited from the MIT community. They were required to have recent instrument flying experience in order to lessen the influence of learning effects on the outcome of the evaluation. Their flying experience in the Piper Arrow aircraft, hours of instrument flying, number of instrument approaches as well as their experience with guidance systems were recorded.

At the outset of each flight test, the pilot received a preflight briefing on the GPS-based trajectory guidance system and the different displays. He was instructed that his task during the flight test was to minimize his deviations from the desired flight path. Once in the air, the pilot was given the opportunity to fly and familiarize himself with the GPS guidance system, the tunnel-in-the-sky display and the flight director. The pilot was then allowed to fly one practice approach using the GPS guidance system with the combined tunnel-in-the-sky and flight director display. This display option exposed the pilot to all the elements of the GPS-based trajectory guidance system. Due to time constraints during each flight test, practice approaches using each of the displays individually was not feasible. In addition, the pilot's recent instrument experience reduced the need for an ILS practice approach.

The entire flight test was flown under simulated instrument conditions in order to limit the factors influencing the pilot's flight performance to the guidance system provided during the experiment. To accomplish this, the pilot wore a hood and his field of view was therefore restricted to the instrument panel. He could not see outside of the aircraft nor could he receive any external references during the experiment.

---

† For the first two flight tests, slightly different gains were used: $K_h = 0.002$, $K_\gamma = -0.4422$, $K_d = 0.2793$, and $K_{\dot{d}} = 0.0224$. These gains were subsequently modified to the ones given in Eq. (8.1) to make the flight director more responsive.

After training and familiarization was completed, the pilot performed the actual flight test approaches according to the test matrix shown in Table 9.1. Each pilot flew the three guidance systems in the order indicated to counterbalanced the learning effects associated with testing human subjects.† When the pilot flew the standard ILS approach, the repeater display was removed.

Pilots	Standard ILS (Baseline)	Tunnel-in-the-Sky (GPS)	Tunnel-in-the-Sky with Flight Director (GPS)
Subject 1	3	1	2
Subject 2	1	3	2
Subject 3	2	1	3

Table 9.1: Flight Test Matrix

For each approach, clearance from air traffic control had to be obtained. All the approaches were videotaped in order to obtain a visual record of how well the aircraft was aligned with the runway. Typically, a missed approach was flown once the aircraft reached an altitude of approximately 500 ft. After each approach, a subject rating for the particular guidance system flown was obtained from the subject pilot using the modified Cooper-Harper scale.

After the aircraft had landed, the pilot subject was given a post flight survey and a debrief. The survey asked to comment on the different displays and included an Analytic Hierarchy Process (AHP) to assess the overall subjective preference of the displays (Yang 1995).

9.3 Results and Discussion

Three pilots served as experimental subjects to evaluate the GPS-based trajectory guidance system. Their flight experience and the wind conditions during their flights are listed in Table 9.2.

† Because of only three test flights, the test matrix was partially counterbalanced

Pilots	Piper Arrow hours	Total hours	Instrument hours / Approaches	Wind conditions during flight
Subject 1	25	550	40 / 40	270 deg, 8 knots
Subject 2	600	5200	250 / 600	100 deg, 5 knots, gust 10
Subject 3	10	200	50 / 100	350 deg, 12 knots, gust 18

Table 9.2: Summary of Subject Pilot Flight Experience and Weather Conditions

Results from the flight tests are shown below. They include qualitative observations and statistical analysis of the flight performance as well as subjective evaluations of the different trajectory guidance systems.

During flight tests, the approach clearance from air traffic control was often received just outside the localizer outer marker. Because clearance was obtained at differing points along the desired flight trajectory, the data considered in the following discussion were limited to the portion of the approach flown after the aircraft had stabilized on the localizer and glideslope. Typically, a missed approach was flown once the aircraft reached an altitude of approximately 500 ft, although in some cases the missed approach was made at higher altitude.

9.3.1 Qualitative Observations of Flight Performance

All the approaches were successfully flown to the missed approach point or the point were Air Traffic Control (ATC) requested a missed approach because of traffic. Aircraft control was never in question. Figure 9.2 shows an out-of-the window view of the runway during the approach (of subject 3) and the corresponding tunnel-in-the-sky display.† The approach flight path is clearly seen to be aligned with Runway 29.

A representative example of the flight performance observed is shown in Figure 9.3. It depicts the flight paths of the approaches flown by subject pilot 3 to Runway 29. This figure offers a qualitative look at the flight performance during the approaches. The data

† The two pictures are screen shots from the flight test video and are captured within three seconds of each other.

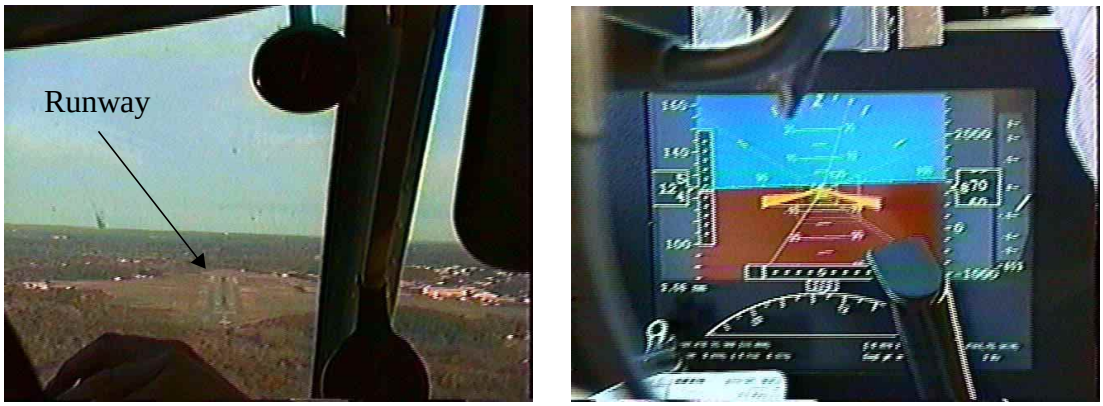

Figure 9.2: (a) Out-of-the-Window View (b) Corresponding Tunnel-in-the-Sky Display

shown as well as observations made during all of the flight tests suggest that the flight performance of approaches flown using both GPS guidance systems was at least comparable to and perhaps better than that the flight performance of approaches flown using the standard ILS. This becomes more evident when looking at the deviations from the desired flight path instead. Figure 9.4 shows the lateral and vertical deviations of the same approaches. In the figure, the tunnel dimensions are indicated in dashed lines.

For the approaches flown using both GPS guidance systems, the aircraft for the most part stayed within the tunnel. The approach flown using the standard ILS, however, shows oscillations in lateral tracking performance with much larger amplitude which in certain cases extend to four times the tunnel dimensions. The variations in vertical deviation appear comparable to the ones experienced with both GPS guidance systems.

It can be further observed that the approach paths based on the GPS guidance system contain oscillations at higher frequency as compared to the long period oscillation of the ILS based approach. This suggests that both GPS trajectory guidance systems allow for a higher bandwidth control system. The control system includes the aircraft, the instrumentation and the display interface, and the human pilot. The bandwidth which a human pilot can achieve is a function of both the degree to which the control skill is practiced and the nature of information provided to the pilot about the state of the aircraft and the world outside (Nagel 1988). This suggests that the tunnel-in-the-sky display (with or without the flight director) provided feedback which allowed for a higher control

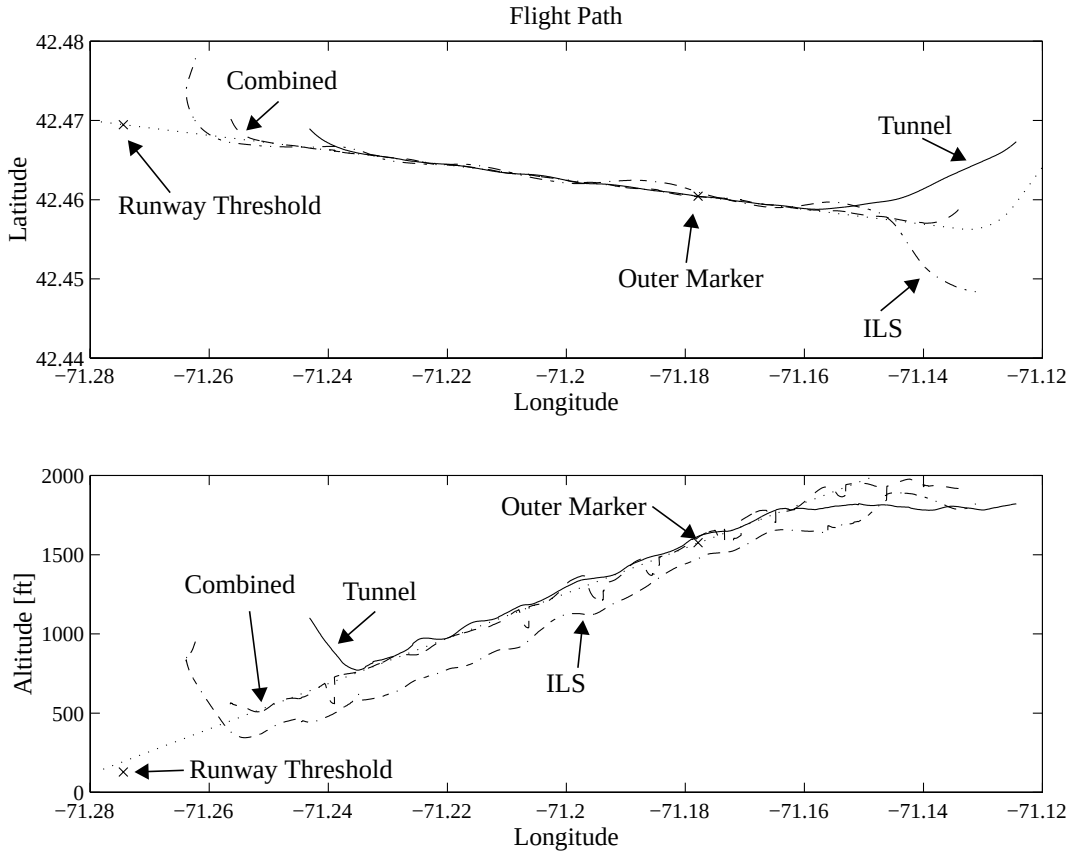

Figure 9.3: Representative Approach Flight Path

bandwidth and, at the same time, resulted in increased physical workload for the pilot since more frequent adjustments to the flight path were made. The results from one of the subject surveys implied that the increase in control bandwidth may also result from a decrease in the pilot's mental workload since the aircraft position relative to the desired flight path no longer needed to be inferred from the ILS needles.

Two problems, caused by the lack of differential GPS corrections, are observable in Figure 9.3 and Figure 9.4. Large jumps in the GPS data, in particular apparent in the altitude data, led to an instantaneous change of the tunnel position on the display. This caused the aircraft to suddenly fly outside the tunnel and prompted the pilot to reacquire the tunnel. These jumps were the result of a configuration change of the satellites tracked by the GPS receiver which led to a new GPS position solution. The configuration change was partly due to rising and setting satellites. Most of the time, however, it was caused by changes in attitude which brought satellites with low elevation angles into and out of the

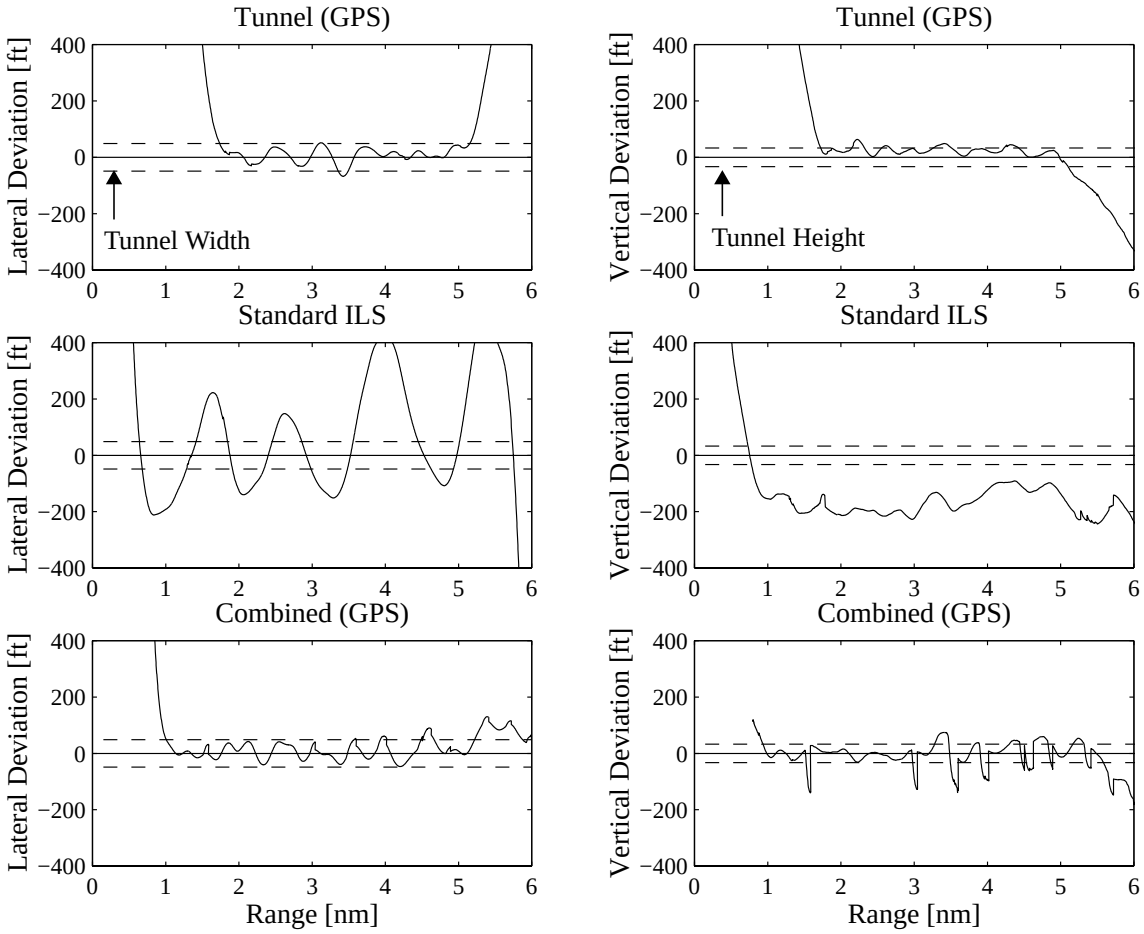

Figure 9.4: Representative Deviations from the Desired Approach Flight Path

GPS antenna field of view. This was aggravated by mounting the antenna inside the crew cabin where the it did not have an unobstructed field of view in all directions. The jumpiness was particularly apparent during one of the flight tests in which the pilot vigorously chased the tunnel thereby causing an increasing number of altitude jumps. The second problem was the nearly constant offset between the ILS and the GPS trajectories, especially apparent in the vertical data. This was assumed to be mainly due to the presence of SA.†

† The availability of differential corrections eliminates the effect of SA and reduces most of the impact of a changing satellite configuration on the position solution. Also, mounting the GPS antenna outside the aircraft and limiting its field of view to exclude GPS satellites with low elevation angles reduces greatly the number of satellite configuration changes.

9.3.2 Analysis of Flight Performance

In this section, the flight performance is evaluated quantitatively using approximately 50 seconds of data (512 sample) in a common altitude range of all the approaches. The data was filtered using a 4th-order Butterworth filter with 0.4 Hz break frequency to mitigate the effects of the jumps which GPS satellite configuration changes introduced into the data. In addition, non-physical jumps in the data recorded during the ILS approach were removed. This was not possible for the data recorded during the approaches using the GPS trajectory guidance systems since the jumps were apparent to the pilot and he compensated for them.

The metrics considered for the evaluation of the flight performance are based on the variations in the tracking error rather than on the mean. The performance metrics used are the standard deviation and the peak-to-peak value of the lateral and vertical tracking error. The peak-to-peak value gives an indication of the worst case excursions but does not include information about how long these maximal deviations last.

Table 9.3 shows the standard deviations (σ) and peak-to-peak (Pp) values of the lateral and vertical tracking errors for the three guidance systems. Also shown, are the altitude ranges chosen for the data analysis. All the values are expressed in feet.

Pilots	**Altitude Range [ft]**	**Lateral Deviation [ft]**						**Vertical Deviation [ft]**					
		ILS (Baseline)		**Tunnel (GPS)**		**Combined (GPS)**		**ILS (Baseline)**		**Tunnel (GPS)**		**Combined (GPS)**	
		σ	**Pp**	σ	**Pp**	σ	**Pp**	σ	**Pp**	σ	**Pp**	σ	**Pp**
Subject 1	800-1300[a]	132	381	30	97	8	33	16	72	21	103	13	80
Subject 2	800-1500	58	196	9	32	6	25	26	245	8	38	15	56
Subject 3	800-1500	201	569	30	119	26	103	39	136	12	46	47	211

Table 9.3: Flight Performance Summary

a. Due to incomplete data acquisition, the data analyzed for the ILS approach was taken from 300-700 ft.

One way to compare the performance of approaches flown with two trajectory guidance systems is to consider the difference of their respective performance metrics, i.e.

$$\begin{aligned} \text{Difference}_{\sigma} &= \sigma_{\text{Guidance System A}} - \sigma_{\text{Guidance System B}} \\ \text{Difference}_{Pp} &= Pp_{\text{Guidance System A}} - Pp_{\text{Guidance System B}} \end{aligned} \tag{9.1}$$

This assumes that the difference in performance is only due to the different guidance systems with the error induced by the pilot's flying skills remaining constant.

Alternatively, the ratio of a performance metric may be used as a comparison for the approaches flown with two of the guidance systems, i.e.

$$\begin{aligned} \text{Ratio}_{\sigma} &= \frac{\sigma_{\text{Guidance System A}}}{\sigma_{\text{Guidance System B}}} \\ \text{Ratio}_{Pp} &= \frac{Pp_{\text{Guidance System A}}}{Pp_{\text{Guidance System B}}} \end{aligned} \tag{9.2}$$

This approach models the effects of the pilot skill level as a constant gain† and gives an indication of the percentage performance improvement obtained when using one system over the other.

Table 9.4 shows the results of a pairwise comparison among the three guidance systems. The entries are the difference and the ratio of the performance metrics for two systems, averaged over all three pilot subjects. The first column is based on the standard deviation performance metric, the second on the peak-to-peak metric. For instance, the peak-to-peak lateral deviations for the ILS approaches were on average more than eight times larger than the corresponding values for the combined tunnel-in-the-sky and flight director guidance system.

To examine the statistical significance of these comparisons, a 5% one-tailed t-test was performed. The comparisons for which the differences in flight performance were significant are highlighted in bold in Table 9.4. The approaches flown using both GPS-

† That is, a more skilled pilot will always perform better than a less skilled pilot regardless of the guidance system used.

Lateral Deviation						
	ILS vs.Tunnel		**ILS vs. Combined**		**Tunnel vs. Combined**	
	σ	**Pp**	**σ**	**Pp**	**σ**	**Pp**
Difference [ft]	**107**	**299.7**	**117**	**329**	9	29
Ratio	**5.9**	**5.0**	**11.2**	**8.3**	2.1	1.8
Vertical Deviation						
	ILS vs.Tunnel		**ILS vs. Combined**		**Tunnel vs. Combined**	
	σ	**Pp**	**σ**	**Pp**	**σ**	**Pp**
Difference [ft]	13	89	2	35	-11	-53
Ratio	2.4	3.3	1.3	2.0	0.8	0.7

Table 9.4: Pairwise Comparison of Flight Performance

based guidance systems show significantly better lateral tracking performance compared to the ILS based approaches. On the other hand, no significant difference in lateral tracking performance can be observed between the two GPS-based systems. Also, no significant difference in vertical tracking performance for any of the guidance systems is apparent. These results are in good correspondence with the qualitative observations made above.

An important point has to be considered before drawing conclusions from these results, namely the lack of an independent measurement of the aircraft position. The aircraft position data for all the approaches was obtained from the same GPS receiver that provided the information for the GPS-based trajectory guidance systems. Since the GPS measurements were obtained in a stand-alone non-differential mode they were affected primarily by SA. This had two immediate consequences.

First, the flight path flown using the standard ILS was measured by GPS and exhibited thus an error in form of an offset. For the statistical evaluation of the data it was assumed that this offset was constant and that its influence was eliminated by considering the variations in the data. In reality, however, SA is slowly varying with a correlation time of

2-5 minutes and may therefore appear as a flight technical error in the approach flown using the ILS.† Nonetheless, a typical SA induced one-sigma error in velocity of 0.2 m/s gives raise to a change of 10 m during the 50 seconds time interval considered for the data evaluation. This error seems not to reduce the validity of the results, considering that some of the ILS approaches had deviations from the desired flight path of more than 100 m.

Second, changing satellite configurations gave raise to jumps in the recorded GPS data. For the approaches flown using the ILS, these jumps could be filtered during post-processing. This was not possible for the approaches flown with the GPS guidance systems since the pilots could observe the jumps and compensate for it. These jumps, therefore, appear to degrade the obtained flight performance for the GPS-based guidance systems more than for the ILS approaches.

9.3.3 Subjective Evaluation of Trajectory Guidance Systems

The subjective ratings given by the pilots for each of the guidance systems using the modified Cooper-Harper scale are summarized in Table 9.5. The Cooper-Harper scale assigns numerical values between 1 and 10 to each configuration tested, with 1 indicating highly desirable flying qualities and 10 indicating deficiencies great enough to cause loss of control. The modified Cooper-Harper Scale is given in Appendix H.

Based on the three pilot evaluations, no significant difference between any of the guidance systems can be observed. A rating of 4 was given in the majority of cases indicating that the display characteristics had minor but annoying deficiencies which required moderate pilot compensation. Two subjects gave the tunnel-in-the-sky guidance option a rating of 5 because they objected to the GPS induced jumps of the tunnel which they experienced while flying this approach. This rating corresponds to moderate objectionable display deficiencies requiring the pilot to apply considerable compensation in order to achieve adequate performance.

---

† The same was not true for the approaches flown with GPS, since slowly varying errors appear as slowly varying changes in the reference flight paths for which the pilot could compensate.

Pilots	Standard ILS (Baseline)	Tunnel-in-the-Sky (GPS)	Tunnel-in-the-Sky with Flight Director (GPS)
Subject 1	4	5	4
Subject 2	4	4	4
Subject 3	4	5	4

Table 9.5: Cooper-Harper Ratings for the Three Guidance Systems[a]

a. The Cooper-Harper scale ranges from 1 to 10, where 1 is the highest rating.

The Analytical Hierarchy Process (AHP) for the three guidance options was performed and the results are shown in Table 9.6 (Yang 1995). The AHP allows for the ranking of multiple systems under evaluation in the order of preference. It hereby considers the relative size of the intervals between the ranking. The AHP is performed through a series of paired comparisons that are recombined to produce an overall weighted ranking (see Appendix H).

Pilots	Standard ILS (Baseline)	Tunnel-in-the-Sky (GPS)	Tunnel-in-the-Sky with Flight Director (GPS)
Subject 1	0.219	0.067	0.715
Subject 2	0.097	0.202	0.701
Subject 3	0.261	0.083	0.657
Average	**0.192**	**0.117**	**0.691**
Ratio	**1**	**0.61**	**3.60**

Table 9.6: Results of the Analytical Hierarchy Process (AHP)

Following the conversion of these values to qualitative descriptions suggested by Yang (1995), a weak preference for the combined tunnel-in-the-sky and flight director display exists compared to the standard ILS guidance option. No clear preference between the GPS-based tunnel-in-the-sky and the ILS is apparent.

It is interesting to add that all the pilots commented on the compelling nature of the tunnel-in-the-sky display, causing them to focus on the tunnel, and on the lack of distance-to-runway feedback it conveys. The former may be mitigated by assigning altitude posts to

the tunnel gates, while the for the latter a moving map may be useful. Furthermore, all pilots indicated that the flight director helped them to gauge the control input necessary to reestablish the aircraft in the middle of the tunnel.

9.4 Conclusions

These flight test results *demonstrated the feasibility* of a tunnel-in-the-sky trajectory guidance system which was based on pseudo-attitude for inner-loop control and, consequently, relied on information entirely provided by a single-antenna GPS receiver.

Qualitative observations as well as the quantitative assessment of the data collected during the approaches suggested that the GPS-based trajectory guidance systems (with or without flight director) allowed for a significant reduction in the lateral tracking errors when compared to the ILS system. No significant differences were found between the two GPS-based guidance systems, nor among all the three guidance systems for vertical flight performance.

Although Cooper-Harper ratings from all the pilots indicated a similar level of deficiency among the three trajectory guidance systems, there existed a preference by the pilots for the combined tunnel-in-the-sky and flight director display.

Chapter 10

Summary and Conclusions

This chapter summarizes the work documented in this thesis, draws conclusions and discusses implications and applications.

10.1 Summary

This thesis developed and demonstrated a novel methodology in which velocity vector information measured by GPS is used to synthesize attitude information and close the flight control loop. The methodology combines a novel velocity vector based flight control paradigm with the benefits of high-quality GPS measurements.

The novel flight control paradigm senses and controls the velocity vector directly rather than through attitude as in traditional flight control schemes. The notion of *pseudo-attitude* as a useful representation of velocity vector based control variables has been introduced and its synthesis presented. It was shown to act as a surrogate to traditional attitude. Pseudo-attitude consists of flight path angle in longitudinal direction and pseudo-roll angle in the lateral direction. The latter is defined as the bank angle corresponding to the observed lateral rate of change of the velocity vector and is calculated from the aircraft acceleration vector. It was demonstrated that for coordinated flight pseudo-roll angle closely corresponds to traditional roll angle and that, therefore, similar control strategies can be employed as for traditional roll angle. In addition, the effects of uncoordinated flight and atmospheric disturbances on pseudo-roll have been investigated.

In addition, this thesis introduced a novel *pseudo-attitude display.* It depicts pseudo-attitude in a manner similar to traditional attitude. However, it distinguishes itself from a traditional attitude display in that its aircraft symbol is referenced to the velocity vector rather than to the aircraft fuselage centerline. The display shows flight path angle instead of pitch angle and pseudo-roll instead of traditional roll angle, and conveys thus a direct feedback of the flight path state.

The measurement of high-quality GPS velocity and acceleration information necessary to synthesize pseudo-attitude was discussed. A Kalman filter with a process model that includes velocity, acceleration and jerk states was proposed to accommodate the correlated nature of aircraft roll rate. The issues, trade-offs and limitations in the signal processing architecture, in the receiver design and in the GPS system as a whole which are pertinent to the synthesis of GPS velocity based flight control were highlighted. In particular, bandwidth issues, the effects of GPS error sources and the limited integrity and availability were discussed.

Next, a linearized aircraft flight control loop including the linearized models of a Cessna 182 aircraft, the GPS receiver and the pseudo-attitude synthesis was developed. This served as a tool to investigate open- and closed-loop behavior of GPS velocity based flight control. In addition, a pseudo-attitude based autopilot guidance logic was presented as a case example.

A flight test system was implemented to demonstrate the pseudo-attitude based flight control concept in flight. The system generated and displayed pseudo-attitude, and pseudo-attitude based autopilot and guidance functions in real-time and was controlled by a Pentium based laptop computer executing the flight test software.

Flight demonstrations of the pseudo-attitude synthesis were successfully conducted and showed the validity of this concept. Multiple flight evaluations of pilot usability of the pseudo-attitude system demonstrated the ability of the pilot to fly the pseudo-attitude display and to close the loop around GPS velocity based information. Furthermore, an ILS approach using pseudo-attitude was demonstrated. In fact, no subjective or substantial objective differences in flight performance were found between pseudo-attitude and traditional attitude in all the tests.

Next, flight tests of a pseudo-attitude based autopilot logic were successfully performed in approach conditions. They demonstrated the feasibility of an autopilot system which relied on pseudo-attitude and, consequently, on information entirely obtained from a single-antenna GPS receiver.

Finally, flight demonstrations of two pseudo-attitude based tunnel-in-the-sky displays, were successfully performed in approach conditions. They illustrated the feasibility of an advanced trajectory guidance system which relied on information entirely obtained from a single-antenna GPS receiver. The two GPS velocity based guidance displays were found to allow for a improved lateral tracking capability compared to the standard ILS guidance. Moreover, subjective data indicated a pilot preference for one of the GPS-based displays over the standard ILS.

10.2 Conclusions

Pseudo-attitude was demonstrated to act in a functionally equivalent manner as traditional attitude with no subjective or substantial objective differences to traditional attitude. GPS velocity based attitude information has a number of beneficial attributes. It relies entirely on solid-state integrated circuit (IC) technology and allows for implementations with lower weight, size, and power consumption and at typically lower cost than traditional attitude sensing instrumentation. At the time of publication, GPS receiver on a single chip begin to emerge. Furthermore, because no alignment or specific aircraft information is required, the implementation as a stand-alone hand-held configuration is particularly appealing. In fact, the instrumentation used throughout a large part of the flight tests documented in this thesis was portable.

However, in order to successfully apply the GPS velocity based flight control concept, its limitations and issues have to be taken into account. Currently, the most limiting factors are the availability and integrity of GPS. Efforts such as the Wide Area Augmentation System (WAAS) are underway to mitigate these issues. While these efforts are directed to enable the use of GPS as a *sole means navigation system*, the use of GPS as a *sole means attitude and flight control sensor for piloted aircraft* is, based on the projected integrity and availability levels, likely premature. Its *supplemental use* in these aircraft as well as its s*ole means use in unpiloted vehicles*, however, have a number of important implications and create unique opportunities for new applications.

Integrity of Cockpit Systems

The supplemental use of GPS velocity based attitude information in piloted aircraft has a potentially profound impact on the integrity of cockpit systems and flight instrumentation. GPS-based pseudo-attitude is then used as an additional level of redundancy in conjunction with traditional attitude systems.

- *Example Application: Backup Attitude Indicator for General Aviation (GA) Aircraft.* The typical GA attitude redundancy architecture is thought to be insufficient since it imposes a prohibitively large mental burden on the pilot in case the primary attitude system fails. The pilot is then required to infer aircraft attitude from the remaining instrumentation (an approach commonly referred to as “flying with needle, ball and airspeed”). On the other hand, a pilot in a GA aircraft equipped with a GPS receiver may obtain GPS velocity based pseudo-attitude indication in addition to primary navigation information. This attitude source constitutes, thus, the second level of redundancy and allows for a human-centered fault detection and reconfiguration task which greatly reduces the pilot’s mental workload and increases aircraft safety.

GPS-based pseudo-attitude constitutes a source of attitude information that is functionally *independent* from attitude measured by traditional inertial sensor based systems. Its use in conjunction with traditional system, thus, not only adds an additional level of redundancy, but provides *dissimilar* redundancy. It is ideally suited for automatic fault detection isolation purposes and has the potential to detect common mode failures in configurations which consist of multiple equivalent primary attitude systems.

- *Example Application: Pseudo-Attitude based Cross-Reference or Tiebreaker.* Business- or commuter-sized aircraft typically are equipped with single- or dual-redundant primary attitude determination systems. In these architectures, pseudo-attitude may serve as a cross-reference or tiebreaker, thereby increasing the capability for fault detection and isolation and, thus, cockpit integrity.

GPS Centered Flight Instrumentation Architecture

The use of high-quality GPS velocity measurements in conjunction with the velocity based flight control paradigm creates the opportunity for a flight instrumentation architecture that is primarily centered around a single-antenna GPS receiver. The novel flight control paradigm is based on sensing and controlling the velocity vector and its rate of change directly, rather than through aircraft attitude as in traditional control schemes. Velocity and acceleration vector information is, however, completely observable from single-antenna GPS carrier Doppler frequency shift measurements. As a result, in this architecture, both the flight control and guidance loops are closed around feedback obtained from a single-antenna GPS receiver. As such, this architecture is associated with lower weight, size, and power consumption and with lower cost than traditional instrumentation architectures.

In order for the velocity vector based control paradigm to be successful, the aircraft has to be well behaved ‘around the velocity vector’. If necessary, stability and control augmentation loops with inertial sensor feedback can be used to achieve damping of unwanted high frequency aircraft modes and improved control responses. The fact that primarily high-frequency components of accelerometer and gyro outputs are fed back, makes their biases and drift rates less significant and allows the use of low-cost automotive grade inertial sensors for this task.

The concept of using inertial sensors to achieve an improved GPS velocity based flight control loop closure is in contrast to traditional integrated INS/GPS based instrumentation concepts where GPS is used to augment the inertial measurements.

- *Example Application: Instrumentation for small UAVs*. A GPS centered instrumentation architectures can be used in small, expendable unmanned aerial vehicles (micro UAV) which currently emerge. The reduction in instrumentation weight and volume is critical for the success of this development. Also, the operational lifetime of these aircraft is envisioned to be short and the limited GPS availability is therefore less critical.

Availability of Affordable Trajectory Guidance

The availability of pseudo-attitude enables a number of flight control and guidance display applications to be implemented solely based on single antenna GPS information. These applications traditionally have relied on the availability of attitude sensing instrumentation and radio-navigation aids or inertial navigation systems (INS) for the necessary attitude, position and velocity reference. Consequently, the potential use of some of the applications have been limited to business-size or larger aircraft only.

The use of pseudo-attitude as the horizon reference enables these applications to be driven by single-antenna GPS information only. This has significant system integration and cost advantages and consequently allows for their implementation in General Aviation (GA) aircraft as well.

Guidance display such as flight director and tunnel-in-the-sky displays can be implemented using single antenna GPS-based position, velocity and pseudo-attitude information. Similarly, pseudo-attitude based autopilot modes such as ground track and altitude capture and hold or approach capture and hold can be realized using GPS-based position, velocity and pseudo-attitude information. A GPS-based autopilot is particularly suited for the control of small scale, expendable UAVs where weight and power are critical.

- *Example Applications: Pseudo-Attitude Based Guidance and Autopilot Systems.* This thesis demonstrated examples of single-antenna GPS-based guidance displays and autopilot modes in flight. A tunnel-in-the-sky and a combined flight director and tunnel-in-the-sky guidance display, as well as a pseudo-attitude based approach autopilot scheme have been implemented and demonstrated in flight.

References

Ackermann, J. (1983), *Abtastregelung*, 2nd ed., Springer Verlag, Berlin.

Amonlirdviman, K. (1998), *Experimental Evaluation of Trajectory Guidance Systems Using Single Antenna GPS*, Final Research Report 16.622, Dec. 8.

Axelrad, P., and Brown, R.G. (1996), *GPS Navigation Algorithms*, GPS: Theory and Application, ed. Parkinson and Spilker, AIAA Progress in Astronautics and Aeronautics Vol. 163, pp. 409-433.

Barrows, A.K., Enge, P., Parkinson, B.W., and Powell J.D. (1995), *Flight Tests of a 3-D Perspective-View Glass-Cockpit Display for General Aviation Using GPS*, ION GPS-95, Palm Springs, CA, Sept. 12-15.

Boeing (1997), *C-MIGITS II*^{TM} Miniature Integrated GPS/INS Tactical System Specifications, Rev 3/97.

Brown, R.G., and Hwang, P.Y.C. (1997), *Introduction to Random Signals and Applied Kalman Filtering*, 3rd ed., John Wiley & Sons, New York.

Bryson, A. E. (1994), *Control of Spacecraft and Aircraft*, Princeton University Press.

Cambridge Aero Instruments (1988), *S-NAV Manual*, Version 2.

Case, D. (1996), *Angle-of-Attack Sensor Analysis and Calibration Method for Heavy Jet Transport*, 15th AIAA/IEEE Digital Avionics Systems Conference, Atlanta, Georgia, Oct. 27-31.

Cohen, C.E. (1996), *Attitude Determination*, GPS: Theory and Application, ed. Parkinson and Spilker, AIAA Progress in Astronautics and Aeronautics Vol. 163, pp. 519-538.

Etkin, B. (1972), *Dynamics of Atmospheric Flight*, John Wiley & Sons, New York.

Etkin, B. (1996), *Dynamics of Flight: Stability and Control*, 3rd ed., John Wiley & Sons.

Federal Aviation Administration (FAA) (1996), *GPS Transition Plan*, July 26.

Federal Radionavigation Plan (FRP) (1996), Department of Defense and Department of Transportation.

Friedland, B. (1986), *Control System Design: An Introduction to State-Space Methods*, McGraw-Hill.

Frost, W. (1983), *Flight in Low-Level Wind Shear*, NASA CR-3678, March.

Gebre-Egziabher, D., Hayward, R.C., and Powell, J.D. (1998), *A Low-Cost GPS/Inertial Attitude Heading Reference System (AHRS) for General Aviation Applications*, IEEE PLANS 98, Palm Springs, CA, April 20-23, pp. 518-525.

Gelb A. (1974), *Applied Optimal Estimation*, The Analytical Sciences Corp.

GPS (1995), *Standard Positioning Service Signal Specification*, 2nd ed, June 2.

Hauser, J., and Hindman, R. (1997), *Aggressive Flight Maneuvers*, Proc. of the 36th IEEE Conference on Decision and Control (CDC), San Diego, CA, December 10-12.

Hayward, R.C., Gebre-Egziabher, D., Schwall, M., Powell, J.D., and Wilson, J. (1997), *Inertially Aided GPS Based Attitude Heading Reference System (AHRS) for General Aviation Aircraft*, ION-GPS 97, Kansas City, MO, September 16-19, pp. 289-298.

Henderson, R.O. (1996), *A Study of GPS Based Attitude Indicators and Instrument Update Rates*, MSRP Report, MIT Aeronautical Systems Laboratory, August 16.

Henderson, R.O. (1997), *A Study of GPS Based Attitude Indicators and Instrument Update Rates*, AIAA Mid-Atlantic Region I Student Conference, Old Dominion University, April.

Hollister, W.M., Chivukula, V., and Toma, B. (1990), *Winds Aloft Measurement and Airspeed Calibration Using Loran*, AIAA Guidance and Control Conference, Portland, OR, August.

Kaplan, E.D. ed. (1996), *Understanding GPS: Principles and Applications*, Artech House Publishers.

Kato, O., and Sugiura, I. (1986), A*n Interpretation of Airplane General Motion and Control as Inverse Problem*, AIAA J. Guidance, Vol. 9, No. 2, March-April, pp.198-204.

King, D.W. (1993), *Handling Qualities Effects of Display Latency*, Proc. of Piloting Vertical Flight Aircraft, pp.4.43-4.54, San Francisco, CA.

Kornfeld, R.P., Hansman, R.J., and Deyst, J.J. (1998a), *Single-Antenna GPS-Based Aircraft Attitude Determination*, NAVIGATION: Journal of The Institute of Navigation, Vol. 45, No. 1, Spring.

Kornfeld, R.P., Hansman, R.J., and Deyst, J.J. (1998b), *Preliminary Flight Tests of Pseudo-Attitude Using Single-Antenna GPS Sensing*, 17th Digital Avionics Systems Conference (DASC), 31 Oct.-6 Nov., Bellevue, WA.

Lambregts, A.A. (1979), *Development of a Control Wheel Steering Mode and Suitable Displays that Reduce Pilot Workload and Improve Efficiency and Safety of Operation in the Terminal Area and in Windshear*, (AIAA Paper 79-1887CP), Guidance and Control Conference, Boulder, CO, August 6-8, pp. 609-620.

Levy, J.L. (1996), *Suboptimality of Cascaded and Federated Kalman Filters*, Proceedings of the 52nd Annual Meeting of the Institute of Navigation, Cambridge, MA, June 19-21, 1996, pp. 399-407.

Mansour, M. (1990), *Digitale Regelsysteme*, Vorlesungsskript 6. Semester, ETH Zurich Switzerland.

Martin, M.K., and Yakos, M.D. (1994), *Modular MIGITS - A Low Cost INS-GPS System For Guidance, Navigation, and Control (GN&C)*, Rockwell International Corporation.

McRuer, D. T. (1973), *Aircraft Dynamics and Automatic Control*, Princeton University Press

McRuer, D.T., and Krendel, E.S. (1974), *Mathematical Models of Human Pilot Behavior*, AGARDograph No. 188.

Nagel, D.C. (1988), *Human Error in Aviation Operations*, Human Factors in Aviation, ed. Wiener and Nagel, Academic Press.

Novatel (1995), *GPSCard^{TM} Command Description Manual*.

Novatel(1996), *3151R GPSCard^{TM} Performance Series Specifications*.

Parkinson, B.W., and Spilker J.J. ed. (1996a), *Global Positioning System: Theory and Applications*, AIAA Progress in Astronautics and Aeronautics Vol. 163.

Parkinson, B.W. (1996b), *GPS Error Analysis*, GPS: Theory and Application, ed. Parkinson and Spilker, AIAA Progress in Astronautics and Aeronautics Vol. 163, pp. 469-484.

Phillips, R.E., and Schmidt, G.T. (1996), *GPS/INS Integration, System Implications and Innovative Applications of Satellite Navigation*, AGARD MSP, LS-207.

Powell, J.D., and Katz, P. (1975), *Sample Rate Selection for Aircraft Digital Control*, AIAA Journal, Vol. 13, No. 8, August, pp. 975-979.

Presidential Decision Directive (PDD) (1996), *Fact Sheet - U.S. Global Positioning System Policy*, The White House, Office of Science and Technology Policy, National Security Council, Department of Transportation Press-Release 62-96, March 29.

Rockwell (1993), *C-MIGITS Message Definitions User's Guide*, Preliminary Version, Rockwell International, Collins Avionics and Communications Division, May 12.

Roskam, J. (1995), *Airplane Flight Dynamics and Automatic Flight Controls*, Part II, DARcoporation.

Schmidt, G.T. (1997), *High Integrity Global Precision Navigation Systems*, Draper Laboratory P-3554, January.

Steinmetz, G.G. (1980), *Simulation Development and Evaluation of an Improved Longitudinal Velocity-Vector Control-Wheel Steering Mode and Electronic Display Format*, NASA Technical Paper 1664, Langley Research Center, Hampton, VA.

Steinmetz, G.G. (1986), *Development and Evaluation of an Airplane Electronic Display Format Aligned With the Inertial Velocity Vector*, NASA Technical Paper 2648, Langley Research Center, Hampton, VA.

Stengel, R.F., and Berry, P.W. (1977), *Stability and Control of Maneuvering High-Performance Aircraft*, NASA CR-2788, April.

Stengel, R.F., and Miller, G.E. (1980), *Flight Tests of a Microprocessor Control System*, AIAA Journal of Guidance and Control, Vol. 3, No. 6, Nov.-Dec., pp. 494-500.

Stevens, B.L., and Lewis, F.L. (1992), *Aircraft Control and Simulation*, John Wiley & Sons.

Tazartes, D., Buchler, R., Tipton, H., and Grethel, R. (1995), *Synergistic Interferometric GPS-INS*, ION National Technical Meeting, Anaheim, CA, January 18-20, pp. 657-671.

Theunissen, E. (1997), *Integrated Design of a Man-Machine Interface for 4-D Navigation*, Ph.D. Thesis, Delft University Press, Delft, The Netherlands.

van Dierendonck, A.J. (1996), *GPS Receivers*, Understanding GPS: Principles and Applications, ed. Kaplan, Artech House Publishers.

van Graas, F., and Braasch (1996), M. S., *Selective Availability*, GPS: Theory and Application, ed. Parkinson and Spilker, AIAA Progress in Astronautics and Aeronautics Vol. 163, pp. 601-622.

Walker, E.M. (1998), *Experimental Evaluation of Trajectory Guidance Systems Using Single Antenna GPS*, Final Research Report 16.622, Dec. 8.

Ward, P. (1996), *Satellite Signal Acquisition and Tracking*, Understanding GPS: Principles and Applications, ed. Kaplan, Artech House Publishers.

Yang, L.C., and Hansman R.J. (1995), *Application of the Analytic Hierarchy Process for Making Subjective Comparisons Between Multiple Automation/Display Options*, 6th IFAC / IFIP / IFORS / IEA Symposium, Cambridge, MA, June 27-29.

Zhuang, W. (1996), *Performance Analysis of GPS Carrier Phase Observable*, IEEE Trans. on Aerospace and Electronic System, Vol. 32, No. 2, April.

Appendix A

Proof of Roll Synthesis Using Equations of Motion

The following proof serves as a complementary illustration to the geometric argument of the roll synthesis given in Section 3.1.2. The proof uses the force equations of the aircraft equations of motion to show that the Eq. (3.8) - Eq. (3.19) yield the roll angle in wind axes.

Unlike in Section 3.1.2, where vectors are expressed in NED components, vectors in the following treatment are expressed in wind axes components. With the exception of Eq. (3.14), this change does not alter the meaning of the Eq. (3.8) - Eq. (3.19). To reduce the complexity of the notation used, no special notation will be introduced to mark the coordinate change.† The vectors $\mathbf{v}_a$, $\mathbf{a}_g$ and $\mathbf{g}$ expressed in wind axes are then

$$\mathbf{v}_a = \begin{bmatrix} v_a \\ 0 \\ 0 \end{bmatrix}, \qquad \mathbf{a}_a = \mathbf{a}_g = \begin{bmatrix} a_{wx} \\ a_{wy} \\ a_{wz} \end{bmatrix} \tag{A.1}$$

$$\mathbf{g} = g_o \cdot \begin{bmatrix} \sin\theta_w \\ \cos\theta_w \cdot \sin\phi_w \\ \cos\theta_w \cdot \cos\phi_w \end{bmatrix} \tag{A.2}$$

Recalling that the wind x_w-axis points in the direction of $\mathbf{v}_a$, the normal components of $\mathbf{a}_g$ and $\mathbf{g}$ are simply

$$\mathbf{a}_g^n = \begin{bmatrix} 0 \\ a_{wy} \\ a_{wz} \end{bmatrix}, \qquad \mathbf{g}^n = g_o \cdot \begin{bmatrix} 0 \\ \cos\theta_w \cdot \sin\phi_w \\ \cos\theta_w \cdot \cos\phi_w \end{bmatrix} \tag{A.3}$$

† The vector $\mathbf{a}_g$ denotes the aircraft acceleration *with respect to* the inertial NED frame $\mathbf{F}_{NED}$. $\mathbf{a}_g$ may, however, be *expressed* in components of *other* frames, such as in wind axes components.

The lift vector **l** is then

$$\mathbf{l} = \mathbf{a}_a^n - \mathbf{g}^n = \begin{bmatrix} 0 \\ a_{wy} - g_o \cos\theta_w \cdot \sin\phi_w \\ a_{wz} - g_o \cos\theta_w \cdot \cos\phi_w \end{bmatrix} \tag{A.4}$$

The force equations of the aircraft equations of motion expressed in wind axes are (Frost 1983)

$$\begin{aligned} ma_{wx} &= m\dot{v}_a + m(q_w w_z - r_w w_y) = T_{xw} - D - mg_o \sin\theta_w \\ ma_{wy} &= m(r_w(v_a + w_x) - p_w w_z) = -C + mg_o \cos\theta_w \sin\phi_w \\ ma_{wz} &= m(p_w w_y - q_w(v_a + w_x)) = T_{zw} - L + mg_o \cos\theta_w \cos\phi_w \end{aligned} \tag{A.5}$$

where T_{xw} and T_{zw} are trust forces in wind x_w- and z_w-axes, respectively, D, C, L are aerodynamic drag, side and lift forces, respectively, $\mathbf{w} = (w_x, w_y, w_z)$ is the wind vector, m is the aircraft mass, and $\omega_w = (p_w, q_w, r_w)$ is the angular velocity relative to $\mathbf{F}_{NED}$ in wind axes components.

Because of the assumption of coordinated flight the sideforce is zero, i.e. C = 0, and the second equation in Eq. (A.5) gives

$$a_{wy} = g_o \cos\theta_w \sin\phi_w \tag{A.6}$$

and by substituting Eq. (A.6) in Eq. (A.4)

$$\mathbf{l} = \mathbf{a}_a^n - \mathbf{g}^n = \begin{bmatrix} 0 \\ 0 \\ a_{wz} - g_o \cos\theta_w \cdot \cos\phi_w \end{bmatrix} \tag{A.7}$$

Eq. (3.17) yields for the reference vector **p** and its magnitude

$$\begin{aligned} \mathbf{p} &= \mathbf{g}^n \times \mathbf{v}_a = \begin{bmatrix} 0 \\ v_a g_o \cos\theta_w \cos\phi_w \\ -v_a g_o \cos\theta_w \sin\phi_w \end{bmatrix} \\ |\mathbf{p}| &= v_a g_o \cos\theta_w \end{aligned} \tag{A.8}$$

Eq. (3.19) finally gives

$$\phi_w = \operatorname{asin}\left(\frac{\mathbf{l}\cdot\mathbf{p}}{|\mathbf{l}|\cdot|\mathbf{p}|}\right) = \operatorname{asin}\left(\frac{(-v_a g_o \cos\theta_w \sin\phi_w)\cdot(a_{wz} - g_o\cos\theta_w\cdot\cos\phi_w)}{\left|v_a g_o \cos\theta_w\right|\cdot\left|a_{wz} - g_o\cos\theta_w\cdot\cos\phi_w\right|}\right) \tag{A.9}$$

For $-\pi/2 \le \theta_w \le \pi/2$, i.e. for the typical flight regime of conventional aircraft (excluding inverted flight), the expression in the first magnitude bars of the denominator in Eq. (A.9) is always positive and the magnitude bars can be omitted. Similarly, for $-\pi/2 \le \phi_w \le \pi/2$, excluding again inverted flight, the lift vector is always pointing along the negative z_w-axis (i.e. in the 'up' direction) and the expression in the second magnitude bars of the denominator is always negative. Eq. (A.9) finally yields

$$\begin{aligned}\phi_w &= \operatorname{asin}\left(\frac{(-v_a g_o \cos\theta_w \sin\phi_w)\cdot(a_{wz} - g_o\cos\theta_w\cdot\cos\phi_w)}{(v_a g_o\cos\theta_w)\cdot(-1)\cdot(a_{wz} - g_o\cos\theta_w\cdot\cos\phi_w))}\right)\\ &= \operatorname{asin}(\sin\phi_w)\\ &= \phi_w\end{aligned} \tag{A.10}$$

This concludes the proof.

Appendix B

Preliminary Simulator Study of the Pseudo-Attitude Display and the Required Display Update Rate

The work presented in this section was performed by Henderson (1996, 1997) in close collaboration with the author of this thesis.† The study is included here in abbreviated form. For details, the reader is referred to aforementioned references.

This study was performed early on in the research work presented in this thesis. At that time, the efforts were mainly concentrated on synthesizing velocity vector based attitude in body axes as described in Section 3.1.2 and Section 3.1.3. For the same reason, no appropriate GPS model was included in the simulator study.

Section B.1 outlines the motivation and the objectives of the simulator study. Section B.2 describes the experimental setup and Section B.3 discusses the results.

B.1 Motivation and Objectives

Before a pseudo-attitude system could be implemented which allowed the pilot to close control loop around it, the following issues had to be addressed: (1) The use and the display of velocity vector based attitude information to close the pilot's control loop was novel and, thus, untested. It was not known whether this non-traditional attitude information and its depiction would allow for adequate pilot performance and aircraft controllability. (2) In contrast to conventional gyro based attitude indicators which provide nearly continuous attitude information, GPS receivers at the time of publication, typically output at frequencies between 1 and 10 Hz, thereby setting the sampling frequency of the control loop. A higher sampling frequency and display update rate is associated with higher equipment cost, as mentioned previously. Therefore, the minimum update rate sufficient for pilots to safely fly the aircraft had to be determined.

---

† This research work was part of an undergraduate summer project jointly supervised by the author of this thesis and Dr. R. John Hansman.

The objectives of the preliminary simulator study were, thus, twofold.

- First, the effectiveness of velocity vector based attitude and attitude displays to close the pilot's control loop had to be investigated.
- Second, the minimum update rate necessary for pilots to maintain adequate controllability had to be determined.

The investigation included the objective measurement of pilot performance and the subjective assessment of pilot preference. In the following sections, the experimental design and the results are briefly outlined.

B.2 Experimental Design

Using the MIT Advanced Cockpit Simulator (ACS), active instrument-rated (IFR) pilots flew a series of tasks using a number of different displays. The tasks, performed under IFR conditions, were the same for each subject with the display order randomized among the subjects to minimize learning effects. For each task and display, the root mean square deviation of selected variables was used to quantify performance. Cooper-Harper data was also collected to obtain a subjective assessment of pilot preference for each display.

For this study, the dynamic model of a Cessna 182 (C182), available on the ACS, was used. The C182 resembles in dimension and weight a Piper Arrow aircraft which was the aircraft type used for flight tests described in the thesis.

The displays tested included:

- a Full Panel: This display simulated the conventional cockpit arrangement found in a C182.

- ‘Needle, Ball and Airspeed’ (NBA): This display froze the traditional attitude indicator, thus simulated the cockpit arrangement resulting from a failure of the attitude indicator. Using this display, the pilot had to infer aircraft attitude information from the remaining sensors.

- Velocity Vector (VV) based Attitude: This display is similar to the full panel arrangement, with the attitude indicator driven by attitude in body axes synthesized from velocity information, as outlined in Eq. (3.5) - Eq. (3.19) and Eq. (3.20)-Eq. (3.23). The known lift curve was used to determine α. Also, the turn coordinator was removed to force the subjects to use the attitude indicator.†

- Velocity Vector (VV) based Roll/Flight Path Angle (Pseudo-Attitude display‡): This display was similar to the foregoing. However, the aircraft symbol of the attitude indicator was referenced to the flight path angle, as computed in Eq. (3.7). The roll angle in body axes was computed as in the foregoing display.

- Combined Velocity Vector (VV) based Attitude/Flight Path Angle Display: This display format was similar to the third display. An additional horizontal bar on the attitude indicator indicated flight path angle.

All displays were updated at 6 Hz. This was the highest update rate obtainable from the Silicon Graphics computer driving the ACS. In addition, to investigate the minimum update rate necessary, the Velocity Vector based Attitude display format was shown at update rates of 3 Hz, 2 Hz, and 1Hz. These displays were selected from the full test matrix (i.e. all display formats at all update rates) to achieve the objectives of the experiment while minimizing the total testing time.

---

† To emulate the effects of a GPS receiver, the same display was also tested with simulated measurement noise and a Kalman filter applied to it. Since the setup did not properly represent GPS receiver characteristics, it is left out of the current discussion.

‡ This display was in essence the pseudo-attitude display shown in Figure 3.9. The sole difference was that roll in body axes, synthesized from velocity information and angle of attack, was displayed instead of pseudo-roll. However, it was shown in Section 3.1.5 that the two roll angles are approximately equal.

The protocol included preflight briefing, training on the different displays as well as the actual tasks. The tasks included recovery from unusual attitudes, straight and level flight at 2000 ft. and 030 heading, a full 360 deg level turn to the left with 30 deg bank angle, intercepting glideslope and localizer and performing an ILS approach. All the tasks were flown under simulated IFR conditions and simulated turbulence. The performance metrics observed for the different tasks were: RMS deviations from the required altitude and heading for the straight and level flight task, RMS deviations from the required altitude and bank angle for the 360 deg turn[†], and RMS deviations from the path referenced by the glideslope and localizer. No performance data for the recovery from unusual attitudes was taken. Immediately after completing the set of tasks with each display, subjects assigned a Cooper-Harper rating to the display.

B.3 Results and Discussion

Eight IFR pilot subjects participated the simulator study[‡]. They averaged 3650 total flight hours and 1788 total instrument hours. In the following objective and subjective results are presented and discussed.

Objective Results

Two-sided paired t-tests were used to compare the average of the subjects' performance metrics of one display to those of another display. For a detailed listing of the statistically significant differences between displays for each task, the reader is referred to (Henderson 1997). Based on the comparisons, the following observations are noteworthy:

- In many instances, the velocity vector based displays yielded significantly better performance (95% or more significance level) than the NBA display, although there are a few display comparisons which are inconclusive with regard to some performance metrics.

† The test run with the NBA display omitted this task, since no attitude is available on this display.
‡ A ninth subject pilot was disqualified for failing to complete the assigned tasks.

- In most instances, no statistically significant difference between the full panel display and the velocity vector based displays could be observed.

- There were no significant differences in performance between the velocity vector based attitude display at 6 Hz, 3 Hz, and 2 Hz update rate. However, for some tasks a significant difference did exist between these update rates and a 1 Hz update rate.

Subjective Results

Table B.1 shows the Cooper-Harper Pilot Opinion Rating for the different displays and update rates (Henderson 1996). The Cooper-Harper scale assigns numerical values between 1 and 10 to each configuration tested, with 1 indicating highly desirable flying qualities and 10 indicating deficiencies great enough to cause loss of control.

Displays	**Cooper-Harper Scale[a]**									
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Full Panel	3	1	1	2						1
NBA				1	1			1	3	2
VV Pseudo-Attitude	3		4							1
VV Combo Display	2		4	1		1				
VV Attitude	2		3	1	1				1	
VV Attitude 3 Hz			2	2	1			1	1	1
VV Attitude 2 Hz				3	1			3		1
VV Attitude 1 Hz								1	5	2

Table B.1: Cooper-Harper Ratings for the Different Displays

a. The Cooper-Harper scale ranges from 1 to 10, where 1 is the highest rating.

The full panel and all the velocity vector based displays (at 6 Hz) received Cooper-Harper ratings between 1 and 3 from the majority of pilot subjects indicating satisfactory performance without improvement. In addition, pilots made several positive comments about the VV based roll/flight path angle display, both during and after the experiment.

The NBA display, on the other hand, received ratings between 8 and 10 from the majority of pilot subjects indicating major deficiencies which require improvements.

The velocity VV based attitude display updated at 1 Hz received ratings from all pilots between 8 and ten. The same display at 2 Hz and 3 Hz update rate received moderate ratings distributed between 3 and 10.

Discussion

Because the simulation tested these displays in only a few different situations, a broad conclusion cannot be made about any of the displays. However, the study shows that the pilots rated the velocity vector based displays similarly as the full panel, and that they preferred those displays over the NBA display. Also, some of the objective data - while not conclusive in all cases - suggest a similar trend. The VV based displays' superior performance to the NBA arrangement may, thus, justify the advantage of having an alternate attitude indicator based on GPS.

In regard to the update rate, some of the objective results - while not conclusive in all cases - suggest that no significant difference exists between update rates of 2 Hz, 3 Hz and 6 Hz, but that a significant difference to a 1 Hz update rate exists. The Cooper-Harper ratings, however, show a clear pilot preference for a 6 Hz update rate. The latter is also in good correspondence with Ackermann's rule of thumb for the choice of the sampling rate for adequate controllability, as stated in Section 3.2.3. The rule calls for a sampling rate of 10 times the highest frequency of the aircraft. In the case of the C182, the highest eigenfrequencies are approximately in the 3 - 4 rad/sec range (0.5 - 0.65 Hz).

It is important to add that the study did not include the characteristics of a typical GPS receiver (latency and bandwidth) in the simulation. Thus, the insights gained from this study may not automatically be applicable to the case where these characteristics are significant. However, the insights were useful to establish preliminary guidelines for the implementation of an actual GPS velocity vector based attitude system.

Appendix C

Topics Related to the Global Positioning System (GPS)

C.1 GPS Signal Structure

GPS satellites transmit on two carrier frequencies, L1 and L2. For civilian applications, the L1 frequency is used primarily and its signal structure is shown in Figure C.1. The L1 signal consists of a carrier signal which is modulated by the Coarse/ Acquisition (C/A) and Precision (P(Y)) spread spectrum signals. The carrier frequency is 1575.42 MHz and is a multiple (154) of the satellite nominal reference frequency f_o.†

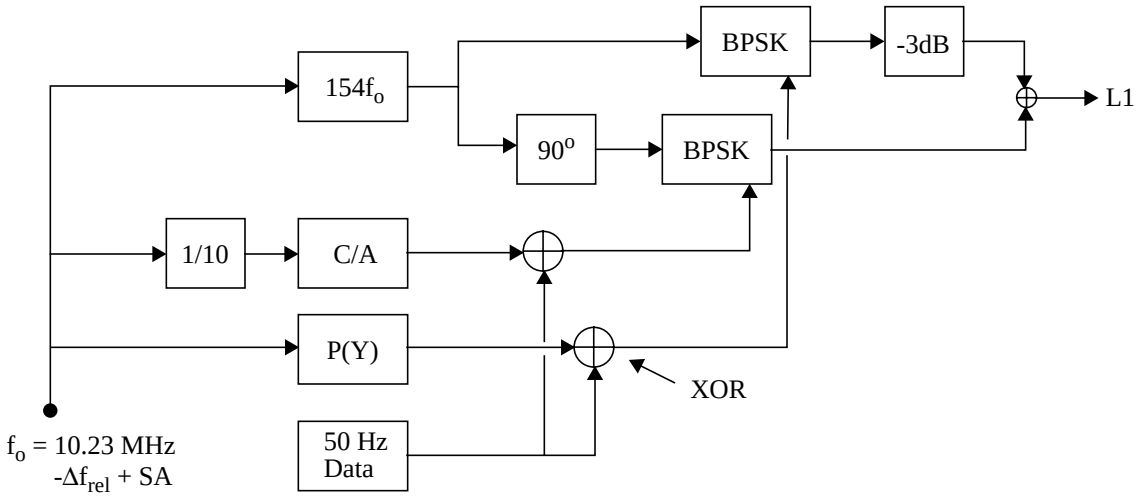

Figure C.1: GPS L1 Signal Structure

The C/A code is a *ranging code* allowing the receiver to determine the range to the satellite. The C/A code consists of a pseudo-random noise (PRN) sequence of zeros and ones which is repeated with a period of about 1 ms. Every satellite has a different C/A code which is orthogonal to the C/A codes of other satellites. Because of this property, all the satellites can transmit on the same carrier frequency using Code Division Multiple Access (CDMA) techniques. A second, more accurate ranging code is the P(Y) code with

† The satellite nominal reference frequency f_o has to be compensated for relativistic effects. This is done by subtracting an offset Δf_{rel} = 0.00457 Hz from f_o.

a frequency ten times higher than the C/A code and a period of one week, allowing for ten times better ranging accuracy. This code is not available for civilian use and is typically encrypted due to its military purpose.

A 50 Hz data stream consisting of navigation messages is modulated onto the C/A and P(Y) code using an exclusive-or (XOR) process. The navigation messages contain information regarding all the satellite's orbital parameters (almanac), clock characteristics and other pertinent information. The composite C/A and P(Y) code are modulated onto the carrier frequency using Binary Phase Shift Keying (BPSK). The P(Y) code is thereby modulated in-phase quadrature with the C/A code.

Both the C/A and P(Y) codes, as well as the L1 and L2 carrier frequencies, are subjected to the dither frequency of Selective Availability (SA). This SA phase modulation effect creates an error in both the pseudo-range and delta range measurements. SA is meant as an intentional degradation of the GPS accuracy by the Department of Defense. SA is encrypted and can only be removed by US military or other authorized users. In addition, SA includes an offset error into the satellites's almanac data broadcasted as part of the navigation data.

The L2 carrier frequency is modulated by the P(Y) code only. Its civilian use is limited mainly to ionospheric delay measurements. L2 signal characteristics will therefore not be treated here.

C.2 Carrier Phase-Locked Loop

Figure C.2 shows a generic carrier phase-locked loop (PLL) (Zhuang 1996). The carrier PLL tracks the incoming baseband signal r by eliminating any phase difference e, between the incoming carrier signal and the carrier replica generated by a local voltage controlled oscillator (VCO), in a closed loop manner. The objective of the carrier PLL is to keep the phase error between the incoming carrier signal and the replica at zero. Any Doppler shift due to satellite or user motion will cause the phase of the incoming signal to be advanced or delayed with respect to the phase of the local carrier. The difference is

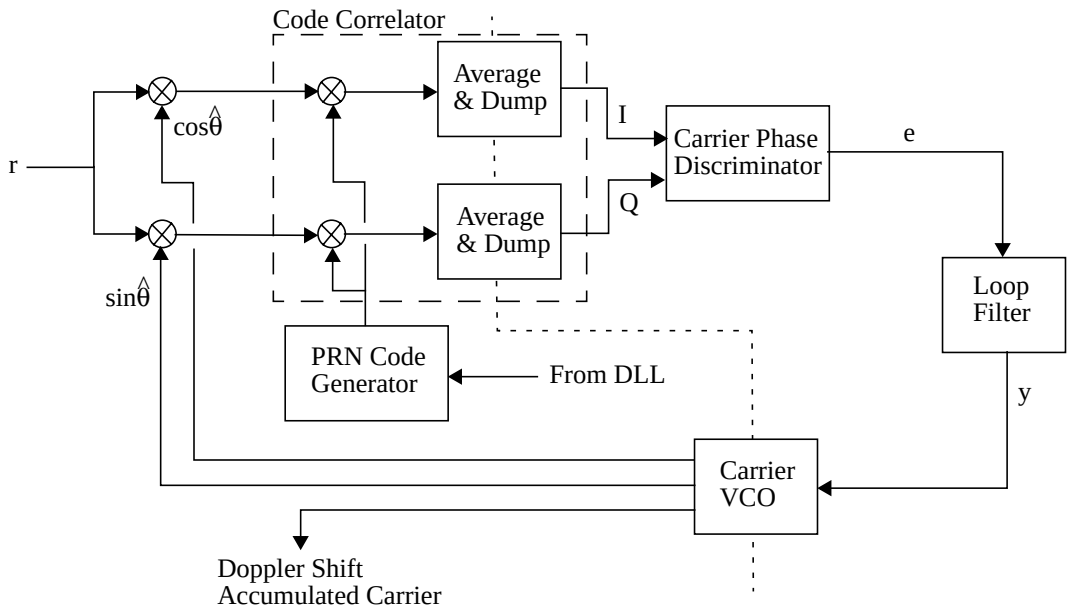

Figure C.2: Carrier Phase-Locked Loop

measured by a carrier phase discriminator and is fed into the VCO to bring the local carrier replica $\hat{\theta}$ in phase alignment with the incoming carrier phase θ. To reduce the effects of thermal noise, the phase difference e is low-pass filtered and the filter output y is fed into the VCO. The phase and Doppler frequency shift of the incoming signal can then be obtained from those of the local VCO output.

Before the difference e of the incoming and the local carrier phase is determined in the carrier phase discriminator, a number of processing steps are performed on the baseband signal as shown in detail in (Zhuang 1996). First, the VCO generates local in-phase $(\cos\hat{\theta})$ and quadrature $(\sin\hat{\theta})$ reference signals to remove the carrier signal by mixing them with the incoming baseband signal. Next, the PRN code of the signals is removed in the in-phase and quadrature code correlator. The signals are hereby multiplied by the aligned (i.e. prompt) replica PRN code as supplied by DLL of the same channel, and filtered by an average (or integrate) and dump detector. The integration time is sometimes called correlation interval or predetection integration time. The correlations yield an in-

phase signal I and quadrature signal Q which contain the desired phase difference information. Finally, the phase difference can be obtained in the carrier phase discriminator.

Typically, the predetection integration time ranges from 1 to 20ms, but does not exceed 20 ms in order to not integrate over a data bit transition of the 50 Hz navigation data stream. Thus, the predetection bandwidth is limited to 50 Hz.

C.3 Single-Point Solution

In the single-point solution, the navigation estimate is the least squares solution to the linearized measurement equation made at a single time. The navigation states are given by the state vector **x**(k) as

$$\mathbf{x}(k) = [\mathbf{r}\ \mathbf{v}\ cb_u\ \delta_u] \tag{C.1}$$

where the elements in the vector are the three dimensional position and velocity vectors, the receiver clock offset and frequency bias, in that order. To obtain a linear measurement model, Eq. (4.2) and Eq. (4.5) must be linearized and in relation to the states to be estimated. The linearization is performed about the current best estimate of the state

$$\hat{\mathbf{x}} = [\hat{\mathbf{r}}\ \hat{\mathbf{v}}\ c\hat{b}_u\ \hat{\delta}_u]^T \tag{C.2}$$

Hence, given an a priori best estimate of the states, predicted pseudo-ranges $\hat{\rho}_i$ and delta ranges $\hat{\delta}_i$ can be calculated as

$$\begin{aligned} \hat{\rho}_i &= \left|\mathbf{r}_i - \hat{\mathbf{r}}_u\right| + c \cdot \hat{b}_u + \hat{\varepsilon}_i \\ \hat{\delta}_i &= (\mathbf{v}_i - \hat{\mathbf{v}}_u) \cdot \hat{\mathbf{1}}_i + \hat{\delta}_u + \hat{\zeta}_i \end{aligned} \tag{C.3}$$

where $\mathbf{r}_i$ is the satellite position, $\hat{\mathbf{1}}_i$ is the estimated line of sight vector from the user to the satellite *i*, and $\hat{\varepsilon}_i$, $\hat{\zeta}_i$ are the estimates of the range error and range error rate, respectively.

The difference between the actual measurements for satellites 1..n, and the predicted values can then be modeled as linearly related to the error in the states (Axelrad 1996). That is,

$$\begin{bmatrix} \rho_1 \\ \vdots \\ \rho_n \\ \delta_1 \\ \vdots \\ \delta_n \end{bmatrix} - \begin{bmatrix} \hat{\rho}_1 \\ \vdots \\ \hat{\rho}_n \\ \hat{\delta}_1 \\ \vdots \\ \hat{\delta}_n \end{bmatrix} = \begin{bmatrix} -\mathbf{1}_1^T & 0 & 1 & 0 \\ \vdots & \vdots & \vdots & \vdots \\ -\mathbf{1}_n^T & 0 & 1 & 0 \\ 0 & -\mathbf{1}_1^T & 0 & 1 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & -\mathbf{1}_n^T & 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta\mathbf{r} \\ \Delta\mathbf{v} \\ c\,\Delta b_u \\ \Delta\delta_u \end{bmatrix} + \begin{bmatrix} \Delta\varepsilon_1 \\ \vdots \\ \Delta\varepsilon_n \\ \Delta\zeta_1 \\ \vdots \\ \Delta\zeta_n \end{bmatrix} \tag{C.4}$$

or in more compact form

$$\begin{bmatrix} \rho \\ \delta \end{bmatrix} - \begin{bmatrix} \hat{\rho} \\ \hat{\delta} \end{bmatrix} = \mathbf{G}\,\Delta\mathbf{x} + \Delta\varepsilon \tag{C.5}$$

where the first matrix on the right hand side is the is measurement connection matrix **G**, and is frequently referred to as the geometry matrix because it contains the line of sight vectors from the user to the satellites. The vector $\Delta\varepsilon$ models the noise in the measurements and its elements are assumed to be zero mean. Eq. (C.5) is to be solved for $\Delta\mathbf{x}$**,** the correction to the a priori state estimate. The least squares solution is then given by

$$\Delta\hat{\mathbf{x}} = (\mathbf{G}^T\mathbf{G})^{-1}\mathbf{G}^T \begin{bmatrix} \rho \\ \delta \end{bmatrix} \tag{C.6}$$

and has to be subtracted from the a priori state estimate $\hat{\mathbf{x}}(k)$ to improve the state estimate.

C.4 Kalman Filter Equations

Extended Kalman Filter (EKF)

With the four matrixes Φ, **Q**, **G** and **R** known, the EKF is in principle specified. Initial estimates of the state $\hat{\mathbf{x}}_o$ and the state error covariance matrix $\mathbf{P}_o$ have to be established for its initialization.

Using the current best estimate, the predicted measurement vector $\hat{\mathbf{z}}_k = \mathbf{h}(\hat{\mathbf{x}}_k^-)$ is computed from Eq. (4.20) for all satellites in view. The measurement connection matrix **G** is then constructed by linearization. The EKF forms an updated state estimate $\hat{\mathbf{x}}_k^+$ as a linear blend of the current measurement information and the previous estimate, projected forward to the current time using the dynamic model. The relative weighting in the blend, **K**, is determined from the a priori error covariance of the states $\mathbf{P}^-$, the measurement covariance matrix **R** and current measurement connection matrix **G**. After updating the state error covariance, the state estimate and its error covariance are propagated to the next measurement time using the assumed process dynamics. Estimates of the state and the covariance after the measurement update are indicated by a superscript "+"; estimates of the state and the covariance propagated ahead are indicated by a superscript "-". The subscript "k" denotes the k-th iteration.

Issues, such as numerical instability and filter divergence, have to be considered when implementing an EKF and means to mitigate them have been developed. Also, the selection of the **Q** and **R** matrices has a significant impact on the convergence, bandwidth and accuracy of the filter solutions. They should be carefully adjusted or *tuned* to achieve best possible performance.

Traditional Kalman Filter

The equations of the traditional discrete Kalman Filter are equivalent to the equations of the EKF with the first step (Computation of the **G** Matrix) omitted, **G** replaced by the measurement connection matrix **H,** and Eq. (C.7) replaced by $\hat{\mathbf{z}}_k = \mathbf{H}\hat{\mathbf{x}}_k^-$.

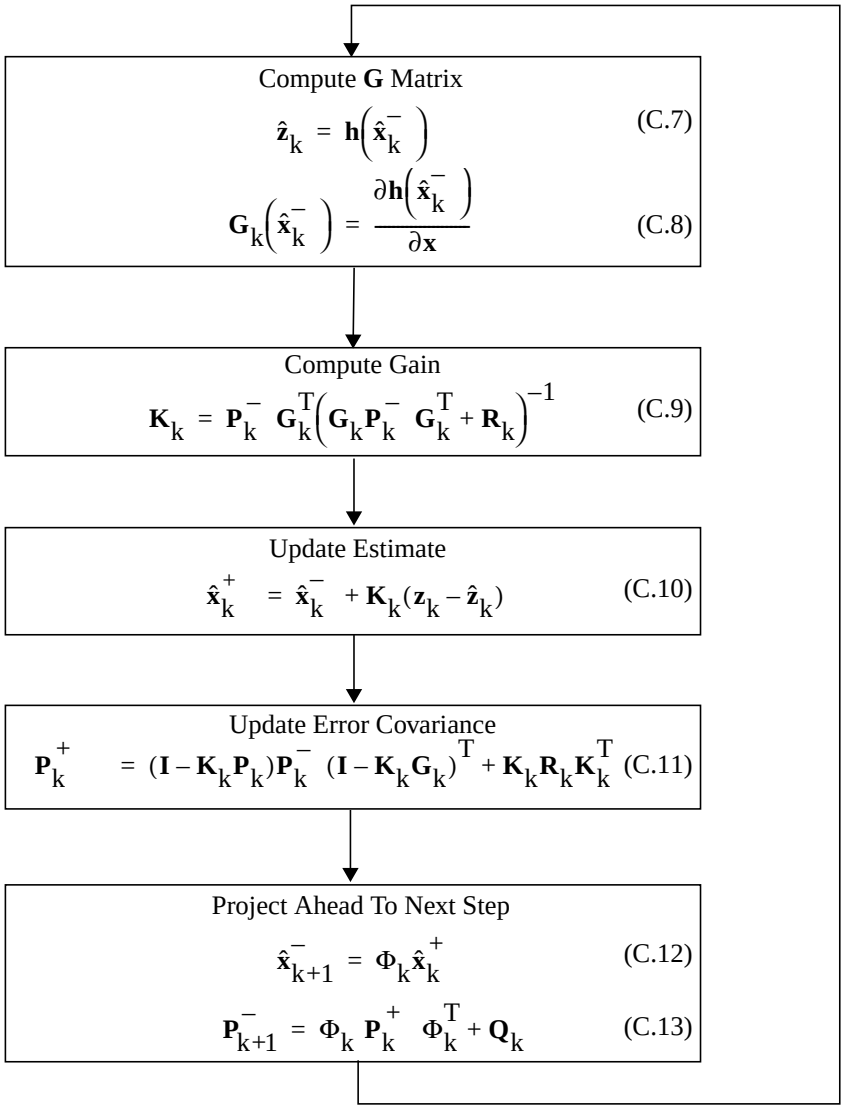

Figure C.3: Discrete Extended Kalman Filter

Appendix D

Transfer Function of Acceleration Estimating Kalman Filter

Given is a system with state vector **x**, input vector **u**, dynamic uncertainty **v**, output vector **y**, measurement noise **w** and the matrices **A**, **B**, **C**, i.e.

$$\begin{aligned} \dot{\mathbf{x}} &= \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} + \mathbf{F}\mathbf{v} \\ \mathbf{y} &= \mathbf{C}\mathbf{x} + \mathbf{w} \end{aligned} \tag{D.1}$$

where

$$\begin{aligned} E[\mathbf{v}(t)\mathbf{v}'(\tau)] &= \mathbf{V}(t)\delta(t-\tau) \\ E[\mathbf{w}(t)\mathbf{w}'(\tau)] &= \mathbf{W}(t)\delta(t-\tau) \\ E[\mathbf{w}(t)\mathbf{v}'(\tau)] &= \mathbf{0} \end{aligned} \tag{D.2}$$

The optimum state estimator is then (Friedland 1986)

$$\dot{\hat{\mathbf{x}}} = \mathbf{A}\hat{\mathbf{x}} + \mathbf{B}\mathbf{u} + \mathbf{K}(\mathbf{y} - \mathbf{C}\hat{\mathbf{x}}) \tag{D.3}$$

where the gain **K** and the error covariance **P** are given by

$$\begin{aligned} \mathbf{K} &= \mathbf{P}\mathbf{C}^{\mathrm{T}}\mathbf{W}^{-1} \\ \dot{\mathbf{P}} &= \mathbf{A}\mathbf{P} + \mathbf{P}\mathbf{A}^{\mathrm{T}} - \mathbf{P}\mathbf{C}^{\mathrm{T}}\mathbf{W}^{-1}\mathbf{C}\mathbf{P} + \mathbf{F}\mathbf{V}\mathbf{F}^{\mathrm{T}} \end{aligned} \tag{D.4}$$

The second equation in Eq. (D.4) is a Riccati equation. A steady state solution may exist, given by the solution of the algebraic Riccati equation (Friedland 1986)

$$\mathbf{0} = \mathbf{A}\mathbf{P} + \mathbf{P}\mathbf{A}^{\mathrm{T}} - \mathbf{P}\mathbf{C}^{\mathrm{T}}\mathbf{W}^{-1}\mathbf{C}\mathbf{P} + \mathbf{F}\mathbf{V}\mathbf{F}^{\mathrm{T}} \tag{D.5}$$

The solution is unique and positive definite if either:

- the system is asymptotically stable, or
- the system [**A**,**C**] is observable and the system [**A**, $\mathbf{FV}^{1/2}$] is controllable.

For the case of a triple integrator plant used to estimate acceleration from velocity measurements and described in Section 4.3.2 the states to be estimated are velocity, acceleration and jerk and the measurement is velocity information y. The continuous model is given as

$$\mathbf{A} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}, \qquad \mathbf{F} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \qquad \mathbf{C} = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}, \qquad \mathrm{B} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \tag{D.6}$$

with the state vector **x** defined as

$$\mathbf{x} = \begin{bmatrix} \text{velocity} \\ \text{acceleration} \\ \text{jerk} \end{bmatrix} \tag{D.7}$$

For the triple integrator plant, **V** and **W** are scalars and are therefore denoted V and W, respectively, in the following treatment. It can be shown that the system meets the controllability and observability condition, and that therefore a unique and positive definite solution exists. With **P** given as

$$\mathbf{P} = \begin{bmatrix} P_1 & P_2 & P_3 \\ P_2 & P_4 & P_5 \\ P_3 & P_5 & P_6 \end{bmatrix} \tag{D.8}$$

Eq. (D.5) yields

$$\mathbf{0} = \begin{bmatrix} P_2 & P_4 & P_5 \\ P_3 & P_5 & P_6 \\ 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} P_2 & P_3 & 0 \\ P_4 & P_5 & 0 \\ P_5 & P_6 & 0 \end{bmatrix} - \begin{bmatrix} P_1^2 & P_1P_2 & P_1P_3 \\ P_1P_2 & P_2^2 & P_2P_3 \\ P_1P_3 & P_2P_3 & P_3^2 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & V \end{bmatrix} \tag{D.9}$$

Solving Eq. (D.9) results in

$$\mathbf{P} = \begin{bmatrix} 2V^{1/6}W^{5/6} & 2V^{1/3}W^{1/3} & V^{1/2}W^{1/2} \\ 2V^{1/3}W^{2/3} & 3V^{1/2}W^{1/2} & 2V^{2/3}W^{1/3} \\ V^{1/2}W^{1/2} & 2V^{2/3}W^{1/3} & 2V^{5/6}W^{1/6} \end{bmatrix}$$
$$\mathbf{K} = \begin{bmatrix} 2(V/W)^{1/6} \\ 2(V/W)^{1/3} \\ (V/W)^{1/2} \end{bmatrix} = \begin{bmatrix} 2\Omega \\ 2\Omega^2 \\ \Omega^3 \end{bmatrix} \quad \text{where} \quad \Omega = (V/W)^{1/6} \tag{D.10}$$

To determine the desired transfer function, Eq. (D.3) has to be transformed into the Laplace domain, hence

$$\mathbf{x}(s) = (s\mathbf{I} - (\mathbf{A} - \mathbf{K}\mathbf{C}))^{-1}\mathbf{K}y(s) \tag{D.11}$$

Evaluating Eq. (D.11) for the acceleration state yields the desired transfer function

$$\frac{x_2(s)}{y(s)} = \frac{2\Omega^2 s^2 + \Omega^3 s}{s^3 + 2\Omega s^2 + 2\Omega^2 s + \Omega^3} \tag{D.12}$$

It is well known that the poles of the transfer function correspond to a Butterworth configuration.

The parameters V and W determine the natural frequency Ω of the transfer function as shown in Eq. (D.10). The corresponding parameters for the discrete-time case of the triple integrator plant are the white noise spectral amplitude S (or the process noise covariance matrix **Q**) and the measurement variance R and are characterized in Section 4.4. The relationship between the continuous-time and discrete-time covariance parameters are given to first-order by (Brown 1997)

$$\mathbf{Q} \approx \mathbf{F}V\mathbf{F}^{\mathrm{T}}\Delta t$$
$$R \approx W/\Delta t \tag{D.13}$$

From Eq. (D.13) it follows that

$$S = V \tag{D.14}$$

The discrete-time parameter values for the North, East, and Down directions described in Section 4.4 are

$$\begin{aligned} S_N = S_E &= 0.001 && m^2/s^7 \\ S_D &= 0.0005 && m^2/s^7 \\ R_N = R_E &= (0.01)^2 && m^2/s^2 \\ R_D &= (0.1)^2 && m^2/s^2 \end{aligned} \tag{D.15}$$

and yield the following filter bandwidths

$$\begin{aligned} \Omega_N = \Omega_E &= (S_N/R_N\Delta t)^{1/6} = 2.15 && rad/sec \\ \Omega_D &= (S_D/R_D\Delta t)^{1/6} = 0.89 && rad/sec \end{aligned} \tag{D.16}$$

Appendix E

Linearized Aircraft Model

The reference flight condition for the linearization of the aircraft model is given by:

- Rectilinear, symmetric flight with wings level
- Vertical flight path angle of γ_o = -3 deg
- Constant aircraft speed of U_o = 85 knots in North direction
- No steady-state wind: $\mathbf{w} = 0$
- Initial altitude of 1500 ft.

and corresponds to an aircraft on an approach flight path. The selected flight condition decouples the linearized longitudinal and lateral aircraft models completely. The linearized, longitudinal and lateral aircraft models are then given by

$$\dot{\mathbf{x}}_{Lon} = \mathbf{A}_{Lon}\mathbf{x}_{Lon} + \mathbf{B}_{Lon}\delta_e + \mathbf{G}_{Lon}\mathbf{g}_{Lon} \tag{E.1}$$

$$\dot{\mathbf{x}}_{Lat} = \mathbf{A}_{Lat}\mathbf{x}_{Lat} + \mathbf{B}_{Lat}\mathbf{u}_{Lat} + \mathbf{G}_{Lat}\mathbf{g}_{Lat} \tag{E.2}$$

The matrices $\mathbf{A}_{Lon}$, $\mathbf{B}_{Lon}$, $\mathbf{G}_{Lon}$, $\mathbf{A}_{Lat}$, $\mathbf{B}_{Lat}$, and $\mathbf{G}_{Lat}$ contain the stability derivatives.† They are evaluated for a Cessna 182 at 1500 ft. altitude and 85 knots approach speed on -3 deg glidepath (Roskam 1995)as:

$$\mathbf{A}_{Lon} = \begin{bmatrix} -0.1366 & 0.1976 & 0 & -32.1409 & 0 & 0 \\ -0.7649 & -1.5511 & 138.7747 & 1.6677 & 0 & 0 \\ 0.0073 & -0.0483 & -4.2080 & -0.0159 & 0 & 0 \\ 0 & 0 & 1.0000 & 0 & 0 & 0 \\ 0 & -1.0000 & 0 & 143.4600 & 0 & 0 \\ 0.0523 & 0.9986 & 0 & -143.2634 & 0 & 0 \end{bmatrix} \tag{E.3}$$

† The air flow is assumed to be quasi-steady. Only the derivatives of w and v are considered in order to account for downwash and sidewash. However, the effects of gusts on the downwash are neglected.

$$\mathbf{B}_{\text{Lon}} = \begin{bmatrix} 0 \\ -21.0659 \\ -15.0515 \\ 0 \\ 0 \\ 0 \end{bmatrix} \qquad \mathbf{G}_{\text{Lon}} = \begin{bmatrix} 0.1366 & -0.1976 & 0 \\ 0.7649 & 1.5511 & 3.2630 \\ -0.0073 & 0.0483 & 2.8546 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \tag{E.4}$$

$$\mathbf{A}_{\text{Lat}} = \begin{bmatrix} -0.1045 & -1.3224 & -142.2121 & 32.1409 & 0 & 0 \\ -0.1044 & -9.5836 & 3.9557 & 0 & 0 & 0 \\ 0.0364 & -0.8976 & -1.0762 & 0 & 0 & 0 \\ 0 & 1.000 & -0.0524 & 0 & 0 & 0 \\ 1.000 & 0 & 0 & 0 & 0 & 143.2634 \\ 0 & 0 & 1.0014 & 0 & 0 & 0 \end{bmatrix} \tag{E.5}$$

$$\mathbf{B}_{\text{Lat}} = \begin{bmatrix} 0 & 9.2530 \\ 35.4075 & 2.2729 \\ -5.8571 & -4.5009 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \qquad \mathbf{G}_{\text{Lat}} = \begin{bmatrix} 0.1045 & 1.3224 & -1.2479 \\ 0.1044 & 9.5836 & -3.9557 \\ -0.0364 & 0.8976 & 1.0762 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \tag{E.6}$$

Appendix F

Linearization of Pseudo-Attitude

In this section pseudo-attitude is linearized around the nominal flight condition specified in Chapter 5 and in Appendix D. The linearization yields pseudo-attitude in stability axes and allows pseudo-attitude to be related to the linearized aircraft states. The stability axes reference frame is a special body axes reference frame obtained by choosing the body x_b-axis to point into the relative wind for the nominal flight condition.

Flight Path Angle

Flight path angle γ is determined from the velocity vector $\mathbf{v}_g = [v_{gN} \; v_{gE} \; v_{gD}]^T$ in the NED frame as

$$\gamma = \text{atan}\left(\frac{-v_{gD}}{\sqrt{v_{gN}^2 + v_{gE}^2}}\right) \tag{F.1}$$

Performing a Taylor expansion about a reference flight condition $\mathbf{v}_{go}$ and retaining only the linear term yields

$$\gamma(\mathbf{v}_g) - \gamma(\mathbf{v}_{go}) = \Delta\gamma = \left.\frac{d\gamma}{d\mathbf{v}_g}\right|_{\mathbf{v}_{go}} \cdot \Delta\mathbf{v}_g \tag{F.2}$$

where the vector $\Delta\mathbf{v}_g = [\Delta v_{gN} \; \Delta v_{gE} \; \Delta v_{gD}]^T$ represents the velocity disturbance state in the NED frame from the reference flight condition. For the reference flight condition given by (see Chapter 5)

$$\mathbf{v}_{go} = \begin{bmatrix} U_o \cos\gamma_o \\ 0 \\ -U_o \sin\gamma_o \end{bmatrix} \tag{F.3}$$

the derivative in Eq. (F.2) and the linearized flight path angle evaluate to

$$\left.\frac{d\gamma}{d\mathbf{v}_g}\right|_{\mathbf{v}_{go}} = \left[-\frac{\sin\gamma_o}{U_o} \quad 0 \quad -\frac{\cos\gamma_o}{U_o}\right] \tag{F.4}$$

$$\Delta\gamma = -\frac{\sin\gamma_o}{U_o}\cdot\Delta v_{gN} - \frac{\cos\gamma_o}{U_o}\cdot\Delta v_{gD} \tag{F.5}$$

Using the following kinematic relations (Etkin 1996)

$$\begin{aligned}\Delta v_{gN} &= \cos\gamma_o\Delta u + \sin\gamma_o\Delta w - U_o\sin\gamma_o\Delta\theta\\ \Delta v_{gD} &= -\sin\gamma_o\Delta u + \cos\gamma_o\Delta w - U_o\cos\gamma_o\Delta\theta\end{aligned} \tag{F.6}$$

where Δu and Δw are the velocity disturbance in the stability x_s- and z_s-axis directions, respectively, Eq. (F.5) can be rewritten as†

$$\begin{aligned}\Delta\gamma &= \Delta\theta - \frac{w}{U_o}\\ &= \Delta\theta - \Delta\alpha\end{aligned} \tag{F.7}$$

Pseudo-Roll

Unlike in Section 3.1.4, where vectors are expressed in NED components, vectors in the following treatment are expressed in body axes components. With the exception of the definition of the Gravity vector **g**, this change does not alter the meaning of the Eq. (3.31) - Eq. (3.37). To reduce the complexity of the notation used, no special notation will be introduced to mark the coordinate change.‡ The velocity vector $\mathbf{v}_g$ is then expressed in body axes components as

$$\mathbf{v}_g = [v_{gx} \; v_{gy} \; v_{gz}]^T \tag{F.8}$$

† For simplicity, the Δs have been omitted in Chapter 5.

‡ The vector $\mathbf{v}_g$ denotes the aircraft velocity *with respect to* the inertial NED frame $\mathbf{F}_{NED}$. $\mathbf{v}_g$ may, however, be *expressed* in components of *other* frames, such as in body or stability axes components.

To simplify the linearization, a specific force vector **f** is introduced as the difference between the aircraft acceleration $\mathbf{a}_g$ and the gravity vector **g**, that is

$$\mathbf{f} = \mathbf{a}_g - \mathbf{g} = [f_x \;\; f_y \;\; f_z]^T \tag{F.9}$$

where **g** is expressed in body axes as

$$\mathbf{g} = \begin{bmatrix} -g_o \sin\theta \\ g_o \sin\phi\cos\theta \\ g_o \cos\phi\cos\theta \end{bmatrix} \tag{F.10}$$

The pseudo-roll synthesis expressed in body axes is then given by

$$\begin{aligned} \tilde{\mathbf{l}} &= \mathbf{f} - \frac{\mathbf{f}\cdot\mathbf{v}_g}{|\mathbf{v}_g|^2}\cdot\mathbf{v}_g = [\tilde{l}_x \;\; \tilde{l}_y \;\; \tilde{l}_z]^T \\ \tilde{\mathbf{p}} &= \mathbf{g}\times\mathbf{v}_g = [\tilde{p}_x \;\; \tilde{p}_y \;\; \tilde{p}_z]^T \\ \tilde{\phi} &= \operatorname{asin}\left(\frac{\tilde{\mathbf{l}}\cdot\tilde{\mathbf{p}}}{|\tilde{\mathbf{l}}|\cdot|\tilde{\mathbf{p}}|}\right) = f(\mathbf{v}_g, \mathbf{f}, \theta, \phi) \end{aligned} \tag{F.11}$$

where pseudo-roll $\tilde{\phi}$ is a function of the velocity $\mathbf{v}_g$, the specific force vector **f**, pitch angle θ and traditional roll angle ϕ. These variables are combined in a vector **r** as

$$\mathbf{r} = [\mathbf{v}_g \;\; \mathbf{f} \;\; \theta \;\; \phi]^T \tag{F.12}$$

Performing a Taylor expansion of the last equation of Eq. (F.11) about a reference flight condition $\mathbf{r}_o$ and retaining only the linear term yields

$$\tilde{\phi}(\mathbf{r}) - \tilde{\phi}(\mathbf{r}_o) = \Delta\tilde{\phi} = \left.\frac{d\tilde{\phi}}{d\mathbf{r}}\right|_{\mathbf{r}_o}\cdot\Delta\mathbf{r} \tag{F.13}$$

where the vector $\Delta\mathbf{r} = [\Delta\mathbf{v}_g \;\; \Delta\mathbf{f} \;\; \Delta\theta \;\; \Delta\phi]^T$ represents the disturbance state from the reference flight condition.

For the reference flight condition $\mathbf{r}_o$ and because of the choice of stability axes, one obtains

$$\theta_o = \gamma_o, \qquad \phi_o = 0$$

$$\mathbf{f}_o = -\mathbf{g}_o = \begin{bmatrix} g_o \sin\gamma_o \\ 0 \\ -g_o \cos\gamma_o \end{bmatrix} \tag{F.14}$$

$$\mathbf{v}_{go} = \begin{bmatrix} U_o \\ 0 \\ 0 \end{bmatrix}$$

and thus

$$\tilde{\mathbf{l}}_o = \begin{bmatrix} 0 \\ 0 \\ -g_o \cos\gamma_o \end{bmatrix}$$

$$\tilde{\mathbf{p}}_o = \begin{bmatrix} 0 \\ U_o g_o \cos\gamma_o \\ 0 \end{bmatrix} \tag{F.15}$$

$$\tilde{\phi}_o = 0$$

That is, for the reference flight condition the pseudo-roll angle is zero, as expected. For the following treatment, it is helpful to introduce an auxiliary variable h defined as

$$h = \frac{\tilde{\mathbf{l}} \cdot \tilde{\mathbf{p}}}{|\tilde{\mathbf{l}}| \cdot |\tilde{\mathbf{p}}|} \tag{F.16}$$

where $h_o = 0$ for the reference flight condition.

Taking the derivative of $\tilde{\phi}$ with respect to $\mathbf{r}$ as shown in Eq. (F.13) yields

$$\left.\frac{d\tilde{\phi}}{d\mathbf{r}}\right|_{\mathbf{r}_o} = \left.\frac{d\tilde{\phi}}{dh}\right|_{h(\mathbf{r}_o)} \cdot \left[\left.\frac{\partial h}{\partial \tilde{\mathbf{l}}}\right|_{\tilde{\mathbf{l}}_o} \cdot \left.\frac{d\tilde{\mathbf{l}}}{d\mathbf{r}}\right|_{\mathbf{r}_o} + \left.\frac{\partial h}{\partial \tilde{\mathbf{p}}}\right|_{\tilde{\mathbf{p}}_o} \cdot \left.\frac{d\tilde{\mathbf{p}}}{d\mathbf{r}}\right|_{\mathbf{r}_o} \right] \tag{F.17}$$

Evaluating the partial derivatives on the left hand side of Eq. (F.17) individually gives:

$$\left.\frac{d\tilde{\phi}}{dh}\right|_{h(\mathbf{r}_o)} = \frac{1}{\sqrt{1-h_o^2}} = 1 \tag{F.18}$$

$$\begin{aligned}\left.\frac{\partial h}{\partial \tilde{\mathbf{l}}}\right|_{\tilde{\mathbf{l}}_o} &= \left[\frac{\partial h}{\partial \tilde{l}_x} \quad \frac{\partial h}{\partial \tilde{l}_y} \quad \frac{\partial h}{\partial \tilde{l}_z}\right] \\ &= \left.\frac{\partial}{\partial \tilde{\mathbf{l}}}\left(\frac{\tilde{l}_x\tilde{p}_x + \tilde{l}_y\tilde{p}_y + \tilde{l}_z\tilde{p}_z}{(\tilde{l}_x^2 + \tilde{l}_y^2 + \tilde{l}_z^2)^{1/2}\cdot|\tilde{\mathbf{p}}|}\right)\right|_{\tilde{\mathbf{l}}_o} \\ &= \left[0 \quad \frac{1}{g_o\cos\gamma_o} \quad 0\right]\end{aligned} \tag{F.19}$$

The expression $\left.d\tilde{\mathbf{l}}/d\mathbf{r}\right|_{\mathbf{r}_o}$ is defined as

$$\left.\frac{d\tilde{\mathbf{l}}}{d\mathbf{r}}\right|_{\mathbf{r}_o} = \begin{bmatrix} \frac{\partial\tilde{l}_x}{\partial v_{gx}} & \frac{\partial\tilde{l}_x}{\partial v_{gy}} & \frac{\partial\tilde{l}_x}{\partial v_{gz}} & \frac{\partial\tilde{l}_x}{\partial f_x} & \frac{\partial\tilde{l}_x}{\partial f_y} & \frac{\partial\tilde{l}_x}{\partial f_z} & \frac{\partial\tilde{l}_x}{\partial\theta} & \frac{\partial\tilde{l}_x}{\partial\phi} \\ \frac{\partial\tilde{l}_y}{\partial v_{gx}} & \frac{\partial\tilde{l}_y}{\partial v_{gy}} & \frac{\partial\tilde{l}_y}{\partial v_{gz}} & \frac{\partial\tilde{l}_y}{\partial f_x} & \frac{\partial\tilde{l}_y}{\partial f_y} & \frac{\partial\tilde{l}_y}{\partial f_z} & \frac{\partial\tilde{l}_y}{\partial\theta} & \frac{\partial\tilde{l}_y}{\partial\phi} \\ \frac{\partial\tilde{l}_z}{\partial v_{gx}} & \frac{\partial\tilde{l}_z}{\partial v_{gy}} & \frac{\partial\tilde{l}_z}{\partial v_{gz}} & \frac{\partial\tilde{l}_z}{\partial f_x} & \frac{\partial\tilde{l}_z}{\partial f_y} & \frac{\partial\tilde{l}_z}{\partial f_z} & \frac{\partial\tilde{l}_z}{\partial\theta} & \frac{\partial\tilde{l}_z}{\partial\phi} \end{bmatrix} \tag{F.20}$$

Due to the zeros in the vector of Eq. (F.19), only the second line in the matrix of Eq. (F.20) is evaluated,i.e.

$$\begin{aligned}\left.\frac{d\tilde{l}_y}{d\mathbf{r}}\right|_{\mathbf{r}_o} &= \left.\frac{d}{d\mathbf{r}}\left(f_y - \frac{f_x v_{gx} + f_y v_{gy} + f_z v_{gz}}{v_{gx}^2 + v_{gy}^2 + v_{gz}^2}\cdot v_{gy}\right)\right|_{\mathbf{r}_o} \\ &= \left[0 \quad -\frac{g_o\sin\gamma_o}{U_o} \quad 0 \quad 0 \quad 1 \quad 0 \quad 0 \quad 0\right]\end{aligned} \tag{F.21}$$

The fourth term in Eq. (F.17) is given by

$$\left.\frac{\partial h}{\partial \tilde{\mathbf{p}}}\right|_{\tilde{\mathbf{p}}_o} = \left.\frac{\partial}{\partial \tilde{\mathbf{p}}}\left(\frac{\tilde{l}_x p_x + \tilde{l}_y p_y + \tilde{l}_z p_y}{|\tilde{\mathbf{l}}| \cdot (\tilde{p}_x^2 + \tilde{p}_y^2 + \tilde{p}_z^2)^{1/2}}\right)\right|_{\tilde{\mathbf{p}}_o} \tag{F.22}$$

$$= \left[0 \quad 0 \quad -\frac{1}{U_o g_o \cos\gamma_o}\right]$$

Due to the zeros in the vector of Eq. (F.22), only the third line in the matrix $d\tilde{\mathbf{p}}/d\mathbf{r}|_{\mathbf{r}_o}$ is evaluated to

$$\left.\frac{d\tilde{p}_3}{d\mathbf{r}}\right|_{\mathbf{r}_o} = \frac{d}{d\mathbf{r}}(-g_o \sin\theta \cdot v_{gy} - g_o \sin\phi \cos\theta \cdot v_{gx}) \tag{F.23}$$

$$= [0 \;\; -g_o \sin\gamma_o \;\; 0 \;\; 0 \;\; 0 \;\; 0 \;\; 0 \;\; -U_o g_o \cos\gamma_o]$$

Evaluating Eq. (F.17) using the expressions in Eq. (F.18) - Eq. (F.19) and Eq. (F.21) - Eq. (F.23) yields

$$\left.\frac{d\tilde{\phi}}{d\mathbf{r}}\right|_{\mathbf{r}_o} = \left[0 \;\; 0 \;\; 0 \;\; 0 \;\; \frac{1}{g_o \cos\gamma_o} \;\; 0 \;\; 0 \;\; 1\right] \tag{F.24}$$

and linearized pseudo-roll is thus †

$$\Delta\tilde{\phi} = \left.\frac{d\tilde{\phi}}{d\mathbf{r}}\right|_{\mathbf{r}_o} \cdot \Delta\mathbf{r} = \Delta\phi + \frac{\Delta f_y}{g_o \cos\gamma_o} \tag{F.25}$$

† For simplicity, the Δs have been omitted in Chapter 5.

Appendix G

Organization of Flight Test Data

The full configuration flight test system generated three data acquisition files on the hard drive of the laptop computer: 1) NovatelGPS.dat containing the data from the Novatel GPS receiver, 2) Migits.dat containing the data from the MIGITS GPS/INS unit used as a reference during the tests, and SNAV.dat containing the data from the air data unit. The portable configuration flight test system generated only the file NovatelGPS.dat.

NovatelGPS.dat

The data stored in this file were from the following three data logs transmitted at 10 Hz update rate by the Novatel 3151R GPS receiver (Novatel 1995):

- VLHB: Velocity, Latency, and Direction over Ground Log
- POSB: Computed Position Log
- DOPB: Dilution of Precision Log

The data was stored in column format in the order shown in Table G.1. Some of the data acquired was reformatted before stored in the file. In order to synchronize the Novatel GPS data with the data from the other sensors, the arrival time of the VLHB message with respect to the start of the data acquisition was recorded.†

Migits.dat

The data stored in this file were from the following three data messages transmitted by the GPS/INS MIGITS unit (Rockwell 1993):

- Message 3500: System Status at 1 Hz update rate
- Message 3501: System Foreground Navigation Solution at 10 Hz update rate
- Message 3502: System Delta-Velocity and Delta-Theta at 10 Hz update rate

† The determination of the exact arrival time of the other logs was less critical (within 0.1 seconds of the VLHB log) since position and DOP data changed typically little within 0.1 seconds.

Column	Data	Data Log	Units
1	Time since start of data acquisition[a]	-	[milliseconds]
2	GPS Seconds-of-Week count	VLHB	[seconds]
3	Velocity latency	VLHB	[seconds]
4	North velocity	VLHB	[meters/seconds]
5	East velocity	VLHB	[meters/seconds]
6	Down velocity	VLHB	[meters/seconds]
7	Solution status	VLHB	see Table G.2
8	Velocity status	VLHB	see Table G.3
9	Latitude	POSB	[degrees.degrees]
10	Longitude	POSB	[degrees.degrees]
11	Height	POSB	[meters MSL][b]
12	HDOP	DOPB	-
13	VDOP	DOPB	-
14	Number of satellites tracked	DOPB	-

Table G.1: Column Format of NovatelGPS.dat

a. Used to synchronize the data with data from other sensors
b. Referenced to WGS84

Value	Description
0	Solution computed
1	Insufficient observations

Table G.2: Solution Status

Value	Description
3	Velocity from single point computations
4	Old velocity from single point computations (higher latency)
5	Invalid velocity

Table G.3: Velocity Status

The data was stored in column format in the order shown in Table G.4. In order to synchronize the MIGITS data with the data from the other sensors, the arrival times of Messages 3501 and 3502 with respect to the start of the data acquisition were recorded.†

Column	Data	Message ID	Unit
1	Time since start of data acquisition[a]	-	[milliseconds]
2	North velocity	3501	[meters/second]
3	East velocity	3501	[meters/second]
4	Up velocity	3501	[meters/second]
5	Pitch angle	3501	[degrees]
6	Roll angle	3501	[degrees]
7	Heading	3501	[degrees]
8	Latitude	3501	[degrees.degrees]
9	Longitude	3501	[degrees.degrees]
10	Height	3501	[meters MSL][b]
11	Time since start of data acquisition	-	[milliseconds]
12	x-velocity increment	3502	[meters/second][c]
13	y-velocity increment	3502	[meters/second][c]
14	z-velocity increment	3502	[meters/second][c]
15	Roll angle increment	3502	[degrees][c]
16	Pitch angle increment	3502	[degrees][c]
17	Heading angle increment	3502	[degrees][c]
18	Current mode	3500	See Table G.5
19	System status	3500	See (Rockwell 1993)
20	Velocity variance	3500	[meters^2/second^2]

Table G.4: Column Format of Migits.dat

a. Used to synchronize the data with data from other sensors
b. Referenced to WGS84
c. To get accelerations and roll, pitch and yaw rates, the data must be multiplied by 10.

Due to an unexpected behavior of the software or operating system, the data in the NovatelGPS.dat and Migits.dat files occasionally had spikes. The spikes could easily be recognized by their values: the sign of a variable suddenly changed or the value of a variable dropped instantaneously to zero. If necessary, the spikes were removed by replacing them with the average of the previous and the following value.

† The determination of the exact arrival time of Message 3500 was less critical for this application (within 1 second of Message 3501).

Value	Description
1	Test
2	Initialization
3	Coarse leveling
4	Fine align
5	INS only
6	GPS/INS
7	GPS only

Table G.5: Current Mode

SNAV.dat

The data stored in this file was transmitted at 0.5 Hz update rate by the SNAV air data computer (Cambridge Aero Instruments 1988)

The data was stored in column format in the order shown in Table G.6. In order to synchronize the SNAV data with the data from the other sensors, the arrival time of the data with respect to the start of the data acquisition was recorded.

Column	Data	Unit
1	Time since start of data acquisition[a]	[milliseconds]
2	True Airspeed	[knots]
3	Heading (Not Used)	-
4	Altitude	[feet]
5	Temperature	[deg Celsius]
6	Equivalent Airspeed	[knots]

Table G.6: Column Format of SNAV.dat

a. Used to synchronize this data with data from other sensors

Appendix H

Cooper-Harper and AHP

Modified Cooper-Harper Scale

Figure H.1 shows the modified Cooper-Harper rating scale used in the flight demonstrations. It allows the pilot to assess the systems characteristics subjectively. The scale ranges from 1 to 10, where 1 is the highest (i.e. most desirable) rating. A Cooper-Harper rating is obtained by guiding the pilot through a series of three questions. These identify one of four rating blocks shown in Figure H.1.Within this block, the pilot chooses the final rating.

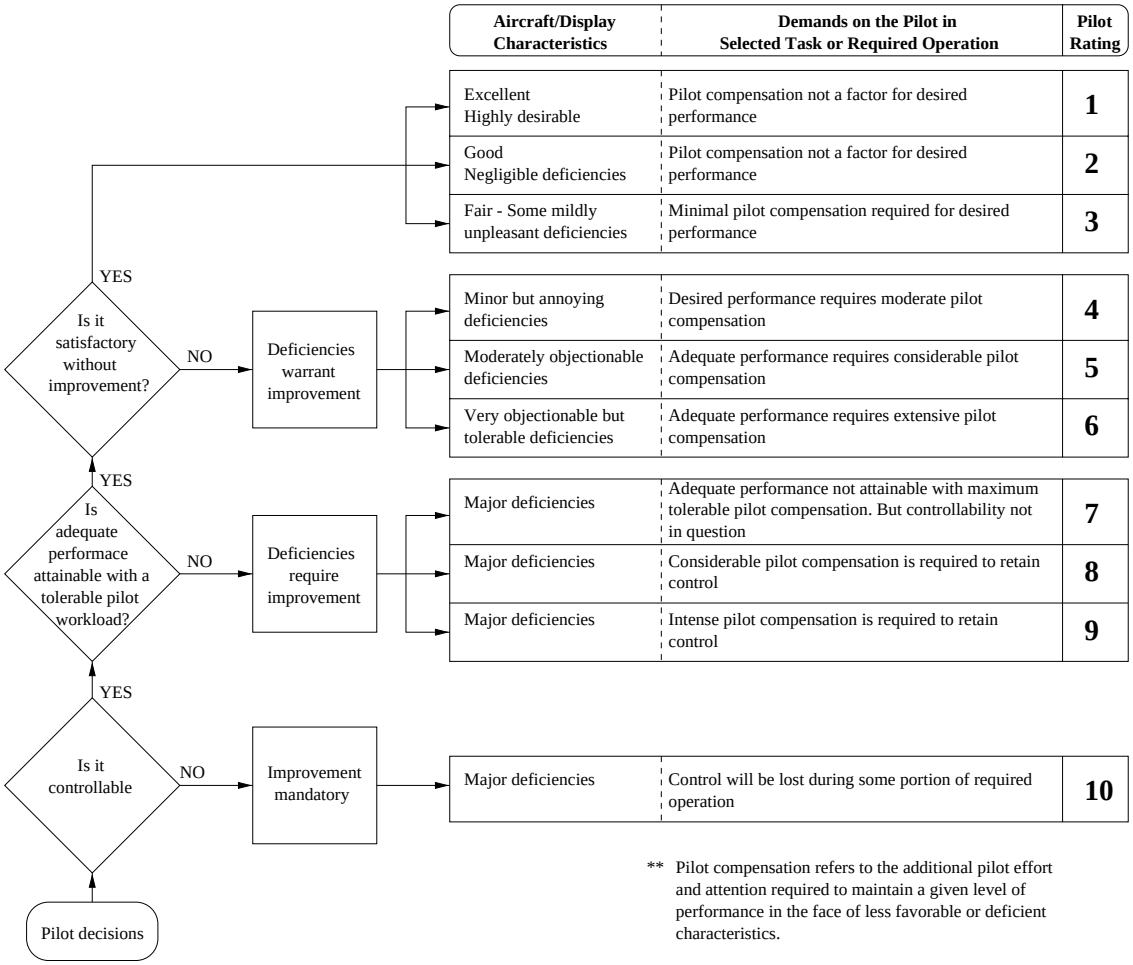

Figure H.1: Modified Cooper-Harper Scale

Analytical Hierarchy Process (AHP)

The AHP is a means of ranking multiple options through a series of paired comparisons that are recombined to produce an overall weighted ranking (Yang 1995). This method captures information regarding the relative size of the interval between the option ranked.† Figure H.2 shows the AHP dominance scale used for paired comparison of guidance options. A description of the AHP can be found in aforementioned reference.

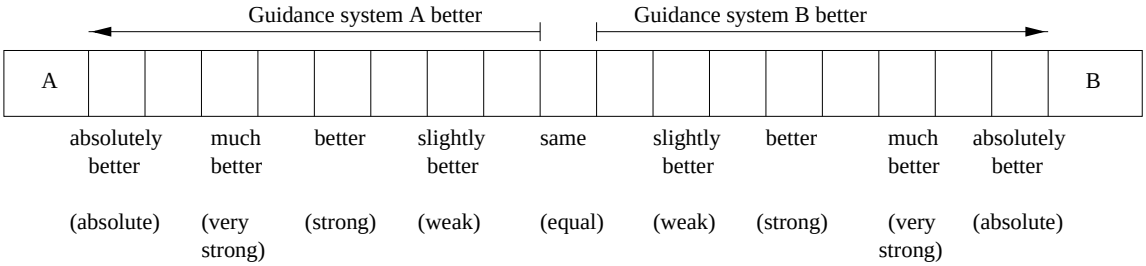

Figure H.2: Analytical Hierarchy Process (AHP) Dominance Scale

† This information is not available if the options under evaluation are only ranked in order of preference.