

INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each original is also photographed in one exposure and is included in reduced form at the back of the book.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

UMI

A Bell & Howell Information Company
300 North Zeeb Road, Ann Arbor MI 48106-1346 USA
313/761-4700 800/521-0600

AIRCRAFT LANDING USING GPS

DEVELOPMENT AND EVALUATION OF A REAL TIME SYSTEM FOR KINEMATIC POSITIONING USING THE GLOBAL POSITIONING SYSTEM

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF AERONAUTICS AND ASTRONAUTICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

By
David Gary Lawrence
September 1996

UMI Number: 9714142

**Copyright 1996 by
Lawrence, David Gary**

All rights reserved.

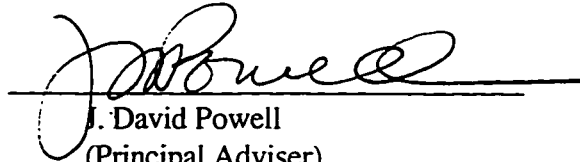
**UMI Microform 9714142
Copyright 1997, by UMI Company. All rights reserved.**

**This microform edition is protected against unauthorized
copying under Title 17, United States Code.**

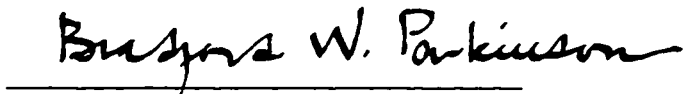
UMI
**300 North Zeeb Road
Ann Arbor, MI 48103**

© Copyright by David G. Lawrence 1996
All Rights Reserved

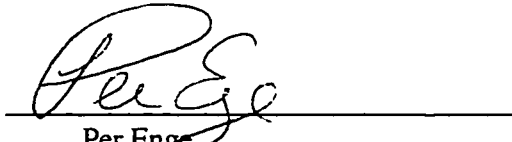
I certify that I have read this thesis and that in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.


J. David Powell
(Principal Adviser)

I certify that I have read this thesis and that in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.


Bradford W. Parkinson

I certify that I have read this thesis and that in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.


Per Eng

Approved for the University Committee on Graduate Studies:



Abstract

The advent of the Global Positioning System (GPS) is revolutionizing the field of navigation. Commercial aviation has been particularly influenced by this worldwide navigation system. From ground vehicle guidance to aircraft landing applications, GPS has the potential to impact many areas of aviation. GPS is already being used for non-precision approach guidance; current research focuses on its application to more critical regimes of flight. To this end, the following contributions were made:

- 1) Development of algorithms and a flexible software architecture capable of providing real-time position solutions accurate to the centimeter level with high integrity. This architecture was used to demonstrate 110 automatic landings of a Boeing 737.
- 2) Assessment of the navigation performance provided by two GPS-based landing systems developed at Stanford, the Integrity Beacon Landing System, and the Wide Area Augmentation System.
- 3) Preliminary evaluation of proposed enhancements to traditional techniques for GPS positioning, specifically, dual antenna positioning and pseudolite augmentation.
- 4) Introduction of a new concept for positioning using airport pseudolites.

The results of this research are promising, showing that GPS-based systems can potentially meet even the stringent requirements of a Category III (zero visibility) landing system. Although technical and logistical hurdles still exist, it is likely that GPS will soon provide aircraft guidance in all phases of flight, including automatic landing, roll-out, and taxi.

Acknowledgments

I would like to thank several individuals and organizations who made this research possible. First, I would like to thank my oral examination committee. My advisor, David Powell, introduced me to the Global Positioning System and provided me with the opportunity to work in this new and exciting field. Additionally, his skills as a teacher, advisor, and pilot are much appreciated. Bradford Parkinson has enabled this research on several levels. As the first program manager for the Air Force's Global Positioning System project, he is partly responsible for the very existence of GPS. His technical insight and management skills have been a great source of inspiration to me. I would also like to thank Per Enge for always making time to provide me with his technical expertise and guidance. I am also deeply indebted to Clark Cohen. Not only did he invent the Integrity Beacon Landing System which is a technology central to this dissertation, but his optimism and guidance have motivated me throughout this research. I am also grateful to Suzanne Pfeffer for acting as chair of my oral examination committee.

Funding for this research was provided by the Satellite Program Office of the Federal Aviation Administration. Their support of Stanford GPS research has been an enabling factor in several technological advances in areas including, but not limited to aircraft navigation. For this support, we are grateful. We would also like to thank Trimble

Navigation, NASA, and United Airlines for the roles they played in carrying out this research.

I owe a great deal to many of my fellow students at Stanford, but I would particularly like to thank two of my colleagues whose work was essential to the successful completion of this research. I benefited greatly from discussions with Boris Pervan who developed high integrity algorithms for GPS positioning. Much of what I know about GPS, I learned from Boris. Stewart Cobb designed and built the GPS pseudolites used throughout this research. His expertise in microwave design and long hours in the lab made the flight tests described in this dissertation possible.

Finally, I would like to thank my friends and family. My parents, Richard and Susan Lawrence have sacrificed a great deal to support me both financially and emotionally. I am forever in their debt, and I dedicate this dissertation to them.

Table of Contents

1. Introduction	1
1.1 Motivation	2
1.2 Required Navigation Performance	3
1.3 The Global Positioning System	5
1.4 Differential GPS	7
1.5 Kinematic DGPS	8
1.6 The Integrity Beacon Landing System	8
1.7 Wide Area DGPS	9
1.8 Previous Work	10
1.9 Contributions	12
2. DGPS Positioning	15
2.1 Carrier Phase Observable	15
2.1.1 Differential Carrier Phase	16
2.1.2 Linearized Differential Carrier Phase	17
2.1.3 Differential Code Phase	18
2.2 Integer Cycle Ambiguities	18
2.2.1 Cycle Ambiguity Resolution	19
2.2.1.1 Integer Searches	21
2.2.1.2 Carrier Smoothed Code	24
2.2.1.3 Satellite Motion	24
2.3 Position Solutions	25
3. Error Sources and Corrections	29
3.1 Receiver Noise	29
3.2 Dynamic Tracking Error	30
3.3 Multipath	31
3.4 Cycle Slips	31
3.5 Troposphere	32
3.5.1 Spacecraft	35
3.5.2 Pseudolite	35
3.6 Ionosphere	37
3.7 Ephemeris Errors	38
3.8 Differential Correction Latency	39
3.9 Circular Polarization	39
3.10 Small Corrections	40

3.10.1 Sampling Time	40
3.10.2 Spacecraft Doppler	41
3.10.3 Signal Travel Time	41
3.10.4 Linearization Error	42
4. Airborne Measurement Processing	45
4.1 Background	46
4.2 Airborne Architecture.....	47
4.2.1 Reference Phase Predictor.....	48
4.2.2 Integer Estimation	50
4.2.2.1 Estimate Initialization	52
4.2.3 Integer Measurement Updates	53
4.2.3.1 Carrier Phase Update	55
4.2.3.2 Code Phase Update	57
4.2.3.3 Integrity Beacon Update	58
4.2.4 Constellation Changes.....	59
4.2.5 RAIM	60
4.2.6 Other Features	61
4.2.6.1 Satellite Motion.....	61
4.2.6.2 Static Survey	65
4.2.6.3 Known Position Updates	66
4.3 System Performance.....	67
4.3.1 24 Hour Positioning Test.....	67
4.3.2 Airborne Performance	69
5. Navigation Performance Assessment.....	73
5.1 Integrity Beacon Landing System Performance	73
5.1.1 Static Survey Comparison	73
5.1.2 Laser Altimeter Comparison	75
5.1.3 Consistency Check	77
5.1.4 King Air Flight Test	78
5.1.5 Autoland Flight Test.....	84
5.1.6 Dynamic Tracking Experiment	92
5.1.7 Integer Residual.....	93
5.2 Wide Area Augmentation System Performance.....	95
5.2.1 Background	95
5.2.2 Integration	97
5.2.3 Flight Test Procedure	100
5.2.4 Flight Tests.....	102
5.2.4.1 Results	102
6. Dual Antenna Positioning.....	105
6.1 Background	106
6.2 Experimental Setup	107

6.2.1 Receivers	107
6.2.2 Aircraft	108
6.2.2.1 Model	108
6.2.2.2 Piper Dakota	109
6.3 Software	109
6.4 Circular Polarization Correction	111
6.4.1 Example of RHCP Correction	114
6.5 Experimental Results	116
6.5.1 Ground Tests	116
6.5.2 Flight Tests	121
7. Airport Pseudolite Augmentation.....	125
7.1 Background	126
7.1.1 Near-Far Problem	129
7.2 Experimental Hardware.....	129
7.2.1 Aircraft	129
7.2.2 Autonomous Pseudolite	130
7.2.3 Pulsed Pseudolite	131
7.2.4 Receiver.....	131
7.3 Background Analysis.....	131
7.3.1 Three Satellite Positioning	136
7.3.2 Four Satellite Fault Detection	136
7.3.3 Five Satellite Fault Isolation.....	137
7.4 Flight Tests	139
7.4.1 Ranging Test	139
7.4.1.1 Phase Profile Error.....	139
7.4.1.2 Nominal Performance.....	140
7.4.1.3 Three Satellite Positioning	141
7.4.1.4 Four Satellite Fault Detection.....	141
7.4.1.5 Five Satellite Fault Isolation.....	142
7.4.2 Near-Far Test.....	143
7.5 Stacked APLs	144
7.5.1 The Stacked APL Concept	145
7.5.2 Multipath.....	148
7.5.3 In-track APLs	148
8. Conclusion.....	153
8.1 Summary of Results and Future Work	153
8.1.1 Airborne Measurement Processing	154
8.1.2 IBLs Performance Assessment	154
8.1.3 Wide Area DGPS Integration with Local Area DGPS.....	155
8.1.4 Dual Antenna Positioning	156
8.1.5 Airport Pseudolites.....	156
8.2 Applications	157
8.3 Closing	158

List of Tables

Table 1.1: Precision Landing Minimums	4
Table 1.2: Proposed Integrity Requirements	5
Table 1.3: Proposed Continuity Requirements	5
Table 4.1: Reference Phase Predictor Error	50
Table 5.1: Navigation Sensor Error for King Air Flight Test	80
Table 5.2: Vertical TSE Statistics for King Air Flight Test	81
Table 5.3: Flight Technical Error for 737 Flight Test	89
Table 7.1: Satellite Availability and Continuity	128

List of Figures

Figure 1.1: The Global Positioning System.....	6
Figure 1.2: The Integrity Beacon Landing System	9
Figure 1.3: Wide Area Differential GPS.....	10
Figure 2.1: GPS Signal Structure	19
Figure 2.2: Constrained 2D Example.....	21
Figure 2.3: 2D Integer Space.....	21
Figure 2.4: Effect of Satellite Motion	22
Figure 3.1: Snell's Law	33
Figure 3.2: Troposphere Slab Model Geometry	35
Figure 3.3: Linearization Error.....	43
Figure 4.1: IBLS Airborne Architecture	48
Figure 4.2: Integer Estimate Initialization.....	53
Figure 4.3: Integer Update Process	55
Figure 4.4: Carrier Phase Update	56
Figure 4.5: Code Phase Update.....	58
Figure 4.6: Satellite Motion	62
Figure 4.7: 24 Hour Positioning Test.....	68
Figure 4.8: Airborne Position Error	71
Figure 5.1: Static Survey Comparison	74
Figure 5.2: Piper Dakota	75
Figure 5.3: Runway Surface Fit	76
Figure 5.4: Laser Altimeter Comparison	77
Figure 5.5: Aircraft Components for King Air Flight Test	78
Figure 5.6: Ground Components for King Air Flight Test.....	79
Figure 5.7: Navigation Sensor Error for King Air Flight Tests.....	82
Figure 5.8: Total System Error for King Air Flight Tests	83
Figure 5.9: Boeing 737 Autoland.....	84
Figure 5.10: PDOP During Autolands	86
Figure 5.11: Vertical NSE for 737 Flight Test	87
Figure 5.12: Cross Track Flight Technical Error for 737	88
Figure 5.13: Vertical Flight Technical Error for 737.....	88
Figure 5.14: Vertical FTE for PDOP < 5.5	89
Figure 5.15: 737 FTE (Approaches without IBLS resets).....	91
Figure 5.16: Normalized Residual versus Acceleration	94
Figure 5.17: Normalized Residual versus Altitude.....	94
Figure 5.18: Original WAAS Accuracy Check	97
Figure 5.19: Information Flow for Integrated System	100

Figure 5.20: Real-time WAAS Position Error	103
Figure 6.1: Model Aircraft.....	108
Figure 6.2: Geometry for Example RHCP Correction.....	115
Figure 6.3: Model Aircraft Roll Reversal.....	117
Figure 6.4: Satellite Tracking Status During Roll Reversal.....	119
Figure 6.5: Starboard Antenna Position	120
Figure 6.6: Evidence of Extreme Multipath.....	123
Figure 6.7: Az/EI Plot of Satellite Tracking Status	124
Figure 6.8: Comparison of Single and Dual Antenna Positioning	124
Figure 7.1:Autonomous Pseudolite Geometry	132
Figure 7.2: Phase Profile Error	140
Figure 7.3: Three Satellite Positioning.....	141
Figure 7.4: Four Satellite Fault Detection	142
Figure 7.5: Five Satellite Fault Isolation	143
Figure 7.6: Stacked APL Geometry.....	146
Figure 7.7: In-track “Stacked” APLs.....	149
Figure 7.8: Stacked APL Covariance Analysis	150
Figure 8.1: Research Contributions	153

CHAPTER 1

Introduction

For thousands of years, man has used the stars as a source of navigation. With the dawn of the space age, the tradition of space-based positioning continues in the form of satellite navigation. Due to its ability to provide worldwide 3D positioning, satellite navigation promises to dominate the field into the foreseeable future. At the forefront of satellite navigation systems is the Global Positioning System. The Initial Operational Capability of GPS coincides with an important decision in the development of aircraft navigation systems: the choice of the next generation of precision landing systems. Although it is clear that GPS will play a large part in the future of aircraft navigation, its limitations for this application have not been well understood. In particular, it has been unclear whether the stringent requirements of precision aircraft landing are achievable with GPS. The focus of this research has been to assess the performance of GPS for aircraft landing and develop methods to improve that performance where necessary.

1.1 Motivation

The existing system for precision approach guidance, the Instrument Landing System (ILS), is described in [Enge]. Several factors have prompted the search for a replacement for ILS.

- 1) The lateral guidance from the ILS localizer is unreliable when the aircraft is more than 10 miles out and more than 10° off course [AIM]. Therefore curved approaches cannot typically be performed with ILS [Kayton]. This shortcoming imposes undesirable procedural constraints.
- 2) The frequency spectrum associated with ILS localizer is scheduled to be reallocated in Europe and an insufficient number of ILS frequencies are available in areas with densely located airports [Skillicorn].
- 3) ILS is susceptible to interference from weather conditions, airport ground traffic, and local terrain.

Originally, the Microwave Landing System (MLS) was proposed as the replacement to ILS. With the advent of GPS however, it is possible that GPS-based systems can meet the requirements for aircraft precision landing. If ILS can be replaced by GPS, there are several potential benefits to the aviation industry. The cost of a GPS based landing system will likely be less expensive than MLS or ILS. A single GPS sensor will provide taxi, departure, en route, approach, and landing guidance. Of these applications, the Required Navigation Performance (RNP) for approach and landing is most stringent. The goal of this research was to assess the navigation performance provided by GPS and develop methods to improve that performance.

1.2 Required Navigation Performance

The Required Navigation Performance (RNP) for aircraft landing is typically specified in terms of four parameters [Kelly]:

- 1) Accuracy: The measure of the navigation output deviation from truth.
- 2) Integrity: The ability of a system to provide timely warnings to users when the system should not be used for navigation.
- 3) Continuity: The likelihood that the navigation signal-in-space supports accuracy and integrity requirements for the duration of intended operation.
- 4) Availability: The fraction of time the navigation function is usable before the approach is initiated.

Airport approaches are classified into precision and non-precision approaches. A non-precision approach is an instrument approach for which no electronic glide slope or accurate source of vertical information is provided. Precision approaches have vertical as well as horizontal guidance, and they are further classified into Categories I, II, and III based on the minimum allowable weather conditions. These minimums are defined by two parameters: the decision height (DH) and the runway visual range (RVR). The DH is defined as the height at which a missed approach must be initiated if the required visual reference to continue the approach to land has not been established [AC120-28C]. RVR represents the horizontal distance a pilot will see down the runway from the approach end [AIM]. Table 1.1 lists the required conditions for precision landing [AC120-28C].

Table 1.1: Precision Landing Minimums

	RVR (ft)	DH (ft)
Cat I	2400	≥ 200
Cat II	1200	100 \rightarrow 200
Cat IIIa	$700 \leq \text{RVR} < 1200$	none \rightarrow 100
Cat IIIb	$150 \leq \text{RVR} < 700$	none \rightarrow 50
Cat IIIc	0	none

Although the requirements for precision landing navigation performance have not been fully defined, Table 1.2 and Table 1.3 list the integrity and continuity requirements proposed in [ORD]. Integrity is specified by the probability of Hazardously Misleading Information (HMI). To prevent HMI, the system should notify the pilot within the time-to-alarm if the navigation solution is in error by more than some protection radius. Continuity Risk is the probability that the signal in space does not continue to support navigation accuracy and integrity requirements during the specified exposure time. The proposed availability is 99.9% or greater.

Table 1.2: Proposed Integrity Requirements

	Probability of HMI	Time-to-Alarm	Protection Radius
Cat I	4×10^{-8}	6 sec	TBD
Cat II	4×10^{-8}	2 sec	TBD
Cat III	4×10^{-8}	2 sec	TBD

Table 1.3: Proposed Continuity Requirements

	Continuity Risk
Cat I	5.5×10^{-5} per approach
Cat II	5.5×10^{-5} per approach
Cat IIIa	4×10^{-5} per approach
Cat IIIb/c	3×10^{-5} per approach 10^{-7} for last 30 secs

1.3 The Global Positioning System

NAVSTAR the Global Positioning System (GPS) is a satellite-based navigation system developed by the U.S. military in the 1970's. The GPS space segment consists of a nominal constellation of 24 satellites, four satellites in each of 6 orbit planes as shown in Figure 1.1.

The orbits are nearly circular with a period of 11.97 hours and an inclination of 55 degrees. The right ascensions of the ascending nodes of the orbits are evenly spaced in 60 degree increments, while the phase of the satellites within the planes is unevenly spaced to minimize the impact of a single satellite failure [Green].




Figure 1.1: The Global Positioning System

The Standard Positioning Service (SPS) signal provided to civilian users of GPS is described in detail in [SPS]. It is made up of an L-band carrier at 1575.42 MHz (referred to as the L1 carrier) modulated by a pseudorandom noise (PRN) C/A (clear acquisition) code. The satellites are distinguished from each other by their unique C/A codes which are nearly orthogonal to each other. The C/A code has a chip rate of 1.023 MHz and is repeated every millisecond. A 50 bit per second data stream is modulated with the C/A code to provide satellite ephemeris and health information. The phase of the C/A code provides a measurement of the range to the satellite. This range includes an offset due to the receiver clock and is therefore referred to as the pseudo-range. Since the receiver clock error is

common to all satellites, it represents an additional unknown to be solved for along with position. Consequently, to perform a 3D position fix with GPS traditionally requires a minimum of four satellites (one satellite phase measurement for each of the unknowns).

The Department of Defense (DOD) has stated that the positioning accuracy provided by the SPS will be 100 meters 2D RMS [SPS]. Most of this error results from an intentional degradation of the signal to prevent hostile users from achieving the full accuracy of the signal. This degradation, achieved by effectively dithering the satellite clock, is called Selective Availability (SA). Without the effect of SA, the 2D RMS errors would likely improve to better than 10 meters [Parkinson, a]. Due to geometric effects, vertical errors are typically larger than horizontal errors. Even without SA, it is clear that stand-alone GPS cannot provide the navigation accuracy required of a precision landing system.

1.4 Differential GPS

If a reference GPS receiver is placed at a known location on the ground, the bulk of the errors associated with the satellite phase measurements can be estimated. Phase corrections can be calculated and broadcast to a roving GPS user. Since most errors are highly correlated in a local area, the roving user's position solution after applying the corrections will be greatly improved. This concept is referred to as Differential GPS (DGPS) and is the basis for all proposed GPS-based aircraft landing systems.

Traditional DGPS systems use the C/A code phase measurements to arrive at position solutions. These systems provide 95% positioning accuracies on the order of a few meters.

The precision of the L1 carrier phase measurement has been used to improve the performance of DGPS. Using carrier smoothed code techniques (Section 2.2.1.2), DGPS performance improves to the meter level. Although this accuracy is arguably sufficient for the precision landing, the integrity of such a system is suspect [LIP]. Further improvements are achieved through the use of Kinematic DGPS (Section 1.5) especially when used in conjunction with the Integrity Beacon Landing System (Section 1.6).

1.5 Kinematic DGPS

Kinematic DGPS, or differential carrier phase GPS, refers to using the differentially corrected carrier phase measurements instead of code phase. Due to the short wavelength of the L1 carrier phase (about 19 cm), these measurements are extremely precise, on the order of several millimeters. Although the measurements are corrupted slightly by the errors sources outlined in Chapter 3, the potential accuracy of kinematic positioning is on the centimeter level. However, the carrier phase measurement has an integer cycle ambiguity associated with it (Section 2.2). This ambiguity arises from the fact that each cycle of the carrier phase is indistinguishable from the others; before centimeter level positioning can be achieved, the ambiguity must be resolved.

1.6 The Integrity Beacon Landing System

The Integrity Beacon Landing System (IBLS) has been developed at Stanford as a means of resolving the cycle ambiguities described in Section 1.5 with high integrity. IBLS is a kinematic DGPS system augmented by ground-based GPS transmitters called pseudolites (short for pseudo-satellites). Figure 1.2 shows the configuration of IBLS. The pseudolites are placed beneath the approach path as shown. The hemispherical “bubbles” represent the

radius outside which the pseudolite signals cannot be tracked. As the aircraft passes through the bubbles, the on-board GPS receiver perceives a doppler signature from the pseudolites. It is this signature that provides the observability required to resolve the cycle ambiguities. The observability analysis is discussed in Section 4.2.6.1 and in [Pervan, a]. When used in this capacity, pseudolites are called Integrity Beacons. After the cycle ambiguities are resolved, centimeter level positioning accuracy is available even after the aircraft exits the bubbles.




Figure 1.2: The Integrity Beacon Landing System

1.7 Wide Area DGPS

The concept of Wide Area DGPS (WADGPS) was conceived at Stanford by Parkinson and Kee [Kee]. Instead of a single reference station in the vicinity of a roving user, WADGPS uses multiple widely separated monitor stations linked to a central master station as shown in Figure 1.3. The geometric leverage provided by these widely scattered monitor stations allows the master station to solve for more than just scalar corrections for each satellite.

Satellite position errors, clock errors (including SA) and ionospheric corrections can also be found. This information, along with integrity information for each satellite is sent to a geosynchronous satellite for rebroadcast back to users anywhere in the coverage area. This process enables users to achieve accuracy on the order of a few meters anywhere in the country.




Figure 1.3: Wide Area Differential GPS

1.8 Previous Work

This research builds on a broad base of earlier work. Some important milestones in the development of GPS landing systems are discussed in this section. Others will be presented within the context of the relevant chapter.

The concept of Differential GPS was introduced by Teasley, Hoover, and Johnson [Teasley] in 1980 and extended by Beser and Parkinson [Beser] in 1982. In 1979, Counselman and Shapiro [Counselman] and MacDoran [MacDoran] independently developed the first GPS carrier-based static survey algorithms. Loomis [Loomis] and Hwang [Hwang] independently developed techniques for resolving carrier phase cycle ambiguities using satellite motion in 1989.

In the last decade there has been considerable research focused on the use of GPS for aircraft precision approach and landing applications. Klein and Parkinson (1985) [Klein] and Parkinson and Fitzgibbon (1989) [Parkinson, b] suggested the use of ground-based pseudolites to improve ranging geometry during precision approach and landing. A DGPS system developed by researchers from the Technical University of Braunschweig [Vieweg] was used to demonstrate the first automatic landing using GPS in 1991. Paielli, et al. (1993) at NASA Ames [Paielli], Van Graas, et al. (1993) of Ohio University [Van Graas, a], and Romrell, et al. (1995) at E-Systems [Romrell] independently applied differential carrier phase using on-the-fly (OTF) cycle ambiguity search techniques for a series of flight tests. Rowson, et al. (1994) [Rowson], Hundley, et al. (1995) [Hundley] of Wilcox Electric, and Van Graas, et al. (1995) [Van Graas, b] applied high-precision code receivers in DGPS architectures to perform automatic landings on transport-class aircraft.

1.9 Contributions

The goal of this research was to assess the performance of GPS-based landing systems while developing improved methods for navigating with GPS. With this goal, the following contributions were made:

- 1) Developed a real-time architecture for kinematic DGPS for the Integrity Beacon Landing System to support feasibility studies of automatic approaches and landings. This architecture was used to demonstrate 110 successful automatic landings of a Boeing 737-300 using a single-frequency GPS receiver. Implemented the architecture on a real-time operating system and demonstrated 20 Hz position solutions. (Chapter 4)
- 2) Developed a flexible integer estimator to estimate the cycle ambiguities associated with kinematic GPS without resorting to integer search techniques. This estimator accepts inputs from a variety of sources, including the results from integrity beacon processing software. Some convergence is attained before the integrity beacon pass from satellite motion and code centered carrier filtering. Receiver Autonomous Integrity Monitoring is performed before each update to the estimator. Constellations changes are handled seamlessly. (Chapter 4)
- 3) Demonstrated kinematic positioning using two antennas to improve the continuity of precise GPS positioning in banked flight. (Chapter 6)
- 4) Assessed accuracy performance of the Integrity Beacon Landing System through extensive flight testing. Several independent sources of positioning truth were used

in this analysis, including laser trackers, a laser altimeter, a runway model, and static surveys. Additional consistency checks were also performed to verify dynamic accuracies on the centimeter level. (Section 5.1)

- 5) Demonstrated kinematic positioning and RAIM augmentation using a ground-based pseudolite to improve performance of GPS positioning on final approach. Invented a new system for positioning using multiple pseudolites. (Chapter 7)
- 6) Developed an algorithm that integrated Stanford's Wide Area DGPS system with the IBLIS system to provide a terminal area position truth model for WADGPS flight tests. (Section 5.2)

CHAPTER 2

DGPS Positioning

This section will introduce the GPS observables and formulate the differential measurement equation. Emphasis is placed on the carrier phase and the associated integer cycle ambiguity. Some error sources will be neglected until Chapter 3. From the differential measurement equation, a position solution will be derived.

2.1 Carrier Phase Observable

The carrier phase measurement for the i^{th} satellite and one receiver is given by:

$$\phi_{i,r} = \rho_{i,r} + \tau_r + l_r + l_{i,s} - \tau_{i,s} + N_{i,r} + v_{i,r} + E_i \quad (2.1)$$

where:

$\phi_{i,r}$ = Carrier phase of the incoming signal relative to the receiver clock's oscillator.

$\rho_{i,r}$ = Distance between the receiver antenna and the i^{th} satellite.

τ_r = Receiver clock error (positive if clock is ahead).

l_r = Effective line bias from receive antenna to receiver phase sampler.

$l_{i,s}$ = Effective line bias from satellite signal generator to satellite transmit antenna.

τ_{i_s} = The clock error for the i^{th} satellite (including intentional clock dithering).

N_{i_r} = An integer cycle ambiguity for satellite i (Section 2.2).

v_{i_r} = Receiver phase measurement error for satellite i .

E_i = Other error sources (ionosphere, troposphere, ephemeris, etc. – see Chapter 3).

All terms are expressed in carrier cycles. For a clock error expressed in seconds to be converted to cycles, it must be multiplied by the carrier frequency.

2.1.1 Differential Carrier Phase

If two receivers measure the phase from the same satellite, the differential carrier measurement (referred to as the “single difference”) reduces to:

$$\phi_i = \Delta\rho_i + \tau + N_i + v_i + \Delta E_i \quad (2.2)$$

where:

ϕ_i = The differential carrier phase measurement (Aircraft phase - Reference phase).

$\Delta\rho_i$ = The range difference between the two receivers from the satellite.

τ = The difference between the receiver clock biases, plus the difference between receiver line biases.

N_i = The difference between the receiver integer cycle ambiguities.

v_i = The differential receiver measurement error.

ΔE_i = The difference between the other error sources as measured by the two receivers. Any common error sources cancel out.

The utility of using the differential measurement is that the satellite clock error, and most of the other error sources are eliminated. Those error sources which remain will be discussed in Chapter 3.

2.1.2 Linearized Differential Carrier Phase

The differential measurement equation may be linearized for the satellite measurements:

$$\phi_i = -\hat{e}_i \cdot \bar{x} + \tau + N_i + v_i + \Delta E_i \quad (2.3)$$

where:

\hat{e}_i = A unit vector toward the satellite.

\bar{x} = The vector from the reference station to the user.

The vector equation for several satellites can be written:

$$\phi = \begin{pmatrix} \phi_1 \\ \vdots \\ \phi_n \end{pmatrix} = \begin{bmatrix} G & I \end{bmatrix} \begin{pmatrix} \bar{x} \\ \tau \\ N_1 \\ \vdots \\ N_n \end{pmatrix} + \begin{pmatrix} \delta\phi_1 \\ \vdots \\ \delta\phi_n \end{pmatrix} \quad (2.4)$$

where:

$$G \equiv \begin{bmatrix} -\hat{e}_1^T & 1 \\ \vdots & \vdots \\ -\hat{e}_n^T & 1 \end{bmatrix} \text{ is the traditional GPS geometry matrix.}$$

$\delta\phi$ is the measurement noise and error.

n is the number of satellites.

2.1.3 Differential Code Phase

Although the differential carrier phase measurement is much more precise than code phase, differential code phase is a welcome complement due to its unambiguous measurement. This differential code phase measurement equation looks similar to Equation (2.4), without the cycle ambiguities:

$$\phi_{code} = \begin{pmatrix} \phi_1 \\ \vdots \\ \phi_n \end{pmatrix}_{code} = [G] \begin{pmatrix} \bar{x} \\ \tau \end{pmatrix} + \begin{pmatrix} \delta\phi_1 \\ \vdots \\ \delta\phi_n \end{pmatrix}_{code} \quad (2.5)$$

2.2 Integer Cycle Ambiguities

Figure 2.1 shows a diagram of the GPS carrier and code phase. Every cycle of the carrier phase is identical; when the receiver acquires the signal, it cannot tell which cycle it is on but only the fractional carrier phase. As long as it maintains continuous lock on the satellite, the receiver can accumulate the number of cycles that have passed since it acquired. However, the initial number of integer cycles in the accumulator is arbitrary giving rise to the so-called integer cycle ambiguity. This ambiguity is constant as long as the receiver tracks the satellite. In Equation (2.1), the term N_i represents this integer cycle ambiguity.

In Figure 2.1, the carrier phase is analogous to the top of the ruler; although the measurement precision is very high (on the order of half a centimeter), the gradations look identical. The code is analogous to the bottom of the ruler. The gradations are unique, so

there is no ambiguity. However, the measurement precision is less (on the order of half a meter). Once the carrier phase integer cycle ambiguity is resolved, the carrier phase measurement is as accurate as it is precise.

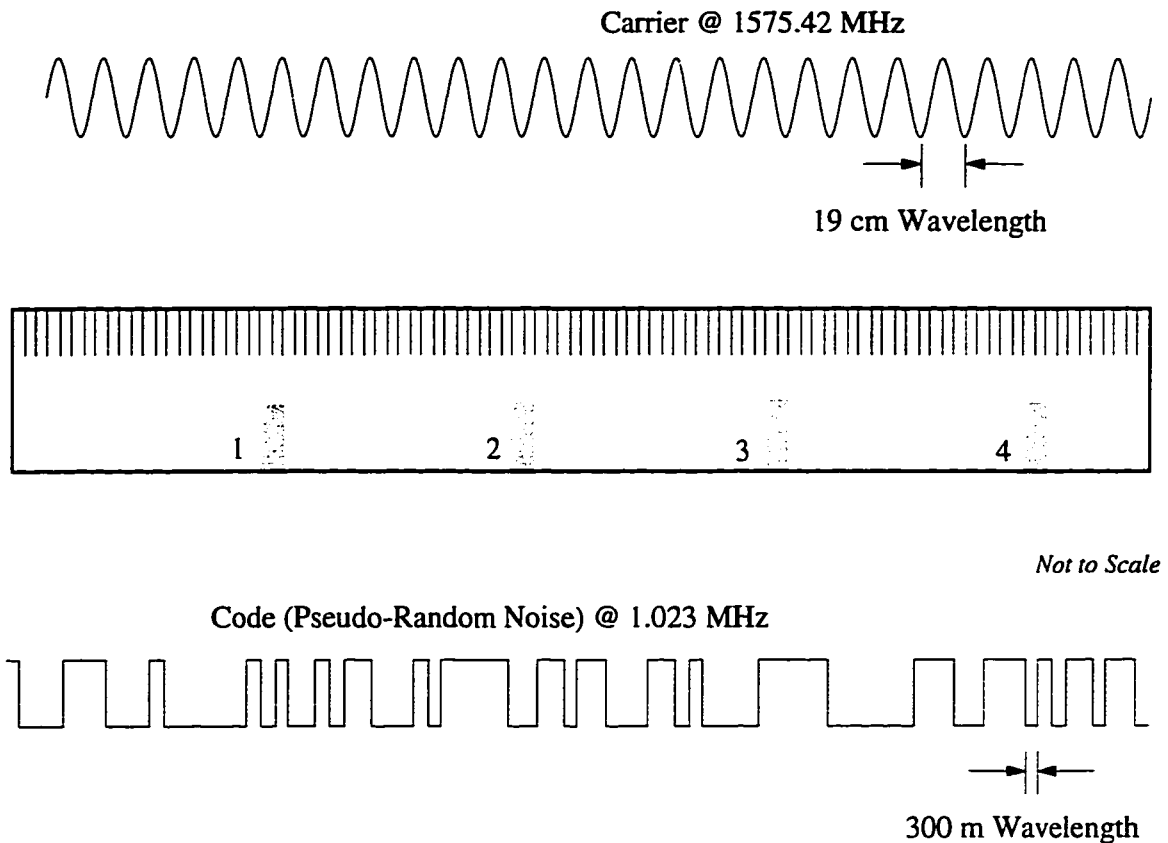


Figure 2.1: GPS Signal Structure

2.2.1 Cycle Ambiguity Resolution

Several techniques are available to resolve or at least estimate the integer cycle ambiguities. To illustrate these methods the following simple example was developed. In the simplified 2D system shown in Figure 2.2, the receiver clock biases are neglected and the DGPS user is constrained to move in one dimension (the horizontal direction). However, satellite directions are free to span two dimensions. The differential carrier phase measurement for one satellite then reduces to:

$$\phi_i = -e_i^x x + N_i + \delta\phi_i \quad (2.6)$$

where:

e_i^x is the horizontal component of the unit line-of-sight vector to satellite i .

For the case of two satellites, Equation (2.6) will produce the vector equation:

$$\begin{pmatrix} \phi_1 \\ \phi_2 \end{pmatrix} = -\begin{pmatrix} e_1^x \\ e_2^x \end{pmatrix} x + \begin{pmatrix} N_1 \\ N_2 \end{pmatrix} + \begin{pmatrix} \delta\phi_1 \\ \delta\phi_2 \end{pmatrix} \quad (2.7)$$

Similarly, Equation (2.5) will reduce to:

$$\begin{pmatrix} \phi_1 \\ \phi_2 \end{pmatrix}_{code} = -\begin{pmatrix} e_1^x \\ e_2^x \end{pmatrix} x + \begin{pmatrix} \delta\phi_1 \\ \delta\phi_2 \end{pmatrix}_{code} \quad (2.8)$$

Neglecting the noise terms, Equation (2.7) can be interpreted as a parametric equation (the parameter being x) for a line in integer space. That is, each possible value of x uniquely determines the values of the cycle ambiguities. The line in Figure 2.3 represents this locus of all possible positions in integer space. This line passes through the correct set of integers. The effect of measurement error simply adds uncertainty to the position of the line; this uncertainty is represented in the figure by the thick gray line. This graphical representation of two dimensional integer space will be used in this chapter and Chapter 4 to illustrate concepts. These concepts translate well to the higher dimensions of real-world applications.




Figure 2.2: Constrained 2D Example




Figure 2.3: 2D Integer Space

2.2.1.1 Integer Searches

One method of determining the cycle ambiguities is called an integer search. The most basic form of integer search requires redundant satellites (more satellites than states of position and clock error – five or more for real-world applications, but only two for the 2D example of Figure 2.2) and a method to bound the range of integer candidates. For example, if the position is estimated using code phase DGPS (Section 2.3), only integer combinations that produce a position consistent with that estimate are candidates.

Redundant satellites provide a measurement residual for each integer combination. This residual can be interpreted graphically as the distance from the solid line in Figure 2.3, while integer combinations are represented by the grid intersections. If the residual for one candidate set of integers is much smaller than all of the other residuals, that candidate is chosen.

As the satellite geometry changes, the orientation of the line in Figure 2.3 will change. For example, Figure 2.4 shows a hypothetical geometry in integer space at two different times. Because the integer cycle ambiguities are constant, grid points that are not consistent with *both* geometries can be rejected. This simplified example is analogous to the 3D integer search.




Figure 2.4: Effect of Satellite Motion

The integer search technique has several disadvantages for the precision landing application:

- 1) The search space must be chosen with high integrity to ensure that it contains the correct solution. However, if the search space is too large, the probability of false integer resolution increases.
- 2) Several redundant satellites with good geometry are typically required. This severely limits the availability of the technique.
- 3) Even with a limited search space volume, there can be millions of integer combinations to search. Although a method developed by Hatch [Hatch, a] can be used to reduce the number of combinations, there are still typically enough candidates that it is impossible to isolate the correct solution with high integrity and availability.
- 4) The technique relies solely on the integer property of the cycle ambiguity. It is therefore not robust to small (centimeter-level) biases which have the effect of removing this integer property. For example, a 2 cm bias on satellite 2 would translate the line in Figure 2.3 up or down by 10% of a cycle, which would change the most likely integer candidate. Such biases do exist; the troposphere, ionosphere, and multipath can all contribute to noise that is highly correlated from epoch to epoch. (Chapter 3)
- 5) If the wrong candidate is chosen, the resulting position could lie anywhere within the search space. Unless there is sufficient satellite motion, the chosen candidate is not necessarily close to the correct position.

Integer search performance can be improved by using wide-landing techniques [Paielli], [Romrell], but then requires the L2 GPS frequency which is not part of the SPS civilian signal.

2.2.1.2 Carrier Smoothed Code

Carrier smoothed code [Goad], [Hatch, b] takes advantage of the best properties of both the carrier and the code to arrive at a better phase measurement than either carrier or code alone. Namely, the precision of the carrier can be used to smooth the unambiguous but noisy code phase. Alternatively, the code phase can be used to slowly converge on the carrier phase cycle ambiguity. This latter technique will be referred to as code-centered carrier integer estimation. Depending on implementation, the two techniques may or may not yield identical results. In particular, the code centered carrier technique is often performed on the differential carrier phase (Section 4.2.3.2), while carrier smoothed code has traditionally been implemented using the undifferenced phases.

2.2.1.3 Satellite Motion

In 1989, Loomis [Loomis] and Hwang [Hwang] independently discovered that satellite motion could provide the observability necessary to resolve the cycle ambiguities in a dynamic environment without using an integer search. This observability is clearly seen by referring to Figure 2.4. The two lines represent the vector Equation (2.7) evaluated at two different times. The resulting four scalar equations can be used to explicitly solve for the horizontal position at the two times and the two cycle ambiguities. The integer property is not used; the ambiguity is allowed to be real-valued or “float.” Because the lines intersect at small angles, the integer observability is weak and the measurement error will be magnified in the estimate error. Therefore, the resulting estimate will not typically be an integer and could be wrong by considerably more than a cycle. As more time is allowed to pass, the observability gets better. However, during the span of a typical aircraft approach,

integer observability due to satellite motion alone is not sufficient to meet the RNP for Category III precision landing.

One subtle difference between the constrained 2D example and real-world applications results from the presence of the receiver clock offset. The addition of this term prevents satellite motion from resolving one component of the cycle ambiguity vector (Section 4.2.6.1). A bias common to all satellites is unobservable using satellite motion. However, any common error in the ambiguity estimates will result in an error in the estimate of the clock offset, not in the estimate of position. Therefore, “cycle ambiguity resolution” will refer to accurately estimating the differences between the integers in Equation (2.4).

IBLS is based on this idea of cycle ambiguity resolution through geometry change. As an aircraft flies over the Integrity Beacons, the geometry of the combined satellite/pseudolite constellation changes rapidly resulting in strong observability of the cycle ambiguities. Once the satellite ambiguities are resolved, centimeter-level positioning is available for the duration of the approach and landing. An estimator that incorporates the effects of satellite motion, code-centered carrier, and Integrity Beacon passes is described in Section 4.2.2.

2.3 Position Solutions

If the integer cycle ambiguities are known, Equation (2.4) can be rewritten to look like Equation (2.5) by subtracting the integers from the phase measurement. Therefore,

traditional DGPS position solutions for both code and carrier are derived from the following equation:

$$\phi = \begin{pmatrix} \phi_1 \\ \vdots \\ \phi_n \end{pmatrix} = [G] \begin{pmatrix} \bar{x} \\ \tau \end{pmatrix} + \begin{pmatrix} \delta\phi_1 \\ \vdots \\ \delta\phi_n \end{pmatrix} \quad (2.9)$$

The measurement errors, $\delta\phi_i$, are typically on the order of half a meter for code phase and half a centimeter for carrier phase. Often there are more measurements than unknowns, so a least squares solution is performed to minimize the sum of the squares residual. That is, the position/clock estimates ($\hat{x}, \hat{\tau}$) are chosen to minimize J :

$$J = \left(\phi - G \begin{pmatrix} \hat{x} \\ \hat{\tau} \end{pmatrix} \right)^T \left(\phi - G \begin{pmatrix} \hat{x} \\ \hat{\tau} \end{pmatrix} \right) \quad (2.10)$$

The resulting estimate [Gelb] is

$$\begin{pmatrix} \hat{x} \\ \hat{\tau} \end{pmatrix} = (G^T G)^{-1} G^T \phi \quad (2.11)$$

Assuming the measurement error is zero mean and uncorrelated with variance σ^2 :

$$\begin{aligned} E(\delta\phi) &= 0 \\ E(\delta\phi\delta\phi^T) &= \sigma^2 I \end{aligned} \quad (2.12)$$

then the position error is also zero mean with covariance:

$$P_{\pi} \equiv E \left(\begin{pmatrix} \tilde{x} \\ \tilde{\tau} \end{pmatrix} \begin{pmatrix} \tilde{x} & \tilde{\tau} \end{pmatrix} \right) = \sigma^2 (G^T G)^{-1} \quad (2.13)$$

If the measurement error is normally distributed, Equation (2.11) represents the maximum likelihood estimate.

The matrix $(G^T G)^{-1}$ is called the Geometric Dilution of Precision (GDOP) matrix; the square root of the trace of the GDOP matrix is the GDOP. The square root of the trace of the first three rows and columns is the Position Dilution of Precision (PDOP). The DOP is a magnification factor that transforms the variance on ranging error to the variance on positioning error. For example, the vertical position error standard deviation is $VDOP \times \sigma$, where VDOP is the square root of the third element on the diagonal of the GDOP matrix. DOP's are traditionally used to judge the geometric quality provided by the satellites being tracked. An average value for GDOP is about 2.5, so typical position errors are several times worse than the underlying measurement error. Obviously, kinematic positioning (with known cycle ambiguities) can tolerate considerably worse DOP than code DGPS due to its accurate measurement. In Chapter 4, the techniques of this section will be extended to kinematic positioning in the presence of uncertainty in the cycle ambiguities.

CHAPTER 3

Error Sources and Corrections

Significant research has been devoted to the study of error sources affecting both stand-alone and differential GPS. The intent of this chapter is not to present in-depth analysis of these error sources, but rather to introduce those errors that dominate the performance of local area DGPS systems. Additionally, small corrections that are applied to the satellite and pseudolite phase measurements for IBLs are discussed.

3.1 Receiver Noise

Thermal noise is a fundamental error source for any type of RF receiver. The impact of this broad-band white noise source can be reduced by trading sensor bandwidth for noise rejection as determined by Equation (3.1) [Carlson] which relates measurement error to the tracking loop bandwidth.

$$\sigma_r = \sqrt{\frac{B}{C/N_0}} \frac{\lambda}{2\pi} \quad (3.1)$$

where:

σ_r is the single difference carrier phase measurement error due to thermal noise.

B is the carrier phase tracking loop bandwidth (Hz).

C/N_0 is the receiver carrier to noise ratio (40 dB-Hz is a typical value).

λ is the wavelength of the carrier phase, 19 cm for L1.

Using typical bandwidths (2-10 Hz), the carrier phase error due to thermal noise is on the millimeter level. As the bandwidth is decreased, the measurement error due to thermal noise decreases at the expense of dynamic tracking error.

3.2 Dynamic Tracking Error

If the aircraft or roving user experiences dynamics with frequency content outside the bandwidth of the receiver tracking loop, phase measurement error results. Therefore, the tracking loop bandwidth should be set comfortably above the expected dynamics for the given application. For most of this research, the bandwidth was 10 Hz, well above the dynamic frequency content of even small airplanes.

For the application of aircraft landing, dynamics introduce a second concern. Any latency from the time the receiver samples the phase until the time the guidance signal is output to the user could result in Path Following Error (PFE). Such delays can be caused by processing and communication lags.

Both previous effects would go undetected by Receiver Autonomous Integrity Monitoring (RAIM) which checks the consistency of phase measurements between satellites. However, both effects were shown to be small (Section 5.1.6) using a 10 Hz receiver tracking loop bandwidth along with the real-time architecture to be discussed in Chapter 4.

3.3 Multipath

As the name suggests, multipath results from both the direct signal and reflected signals entering the antenna. As with thermal noise, both the carrier and code are affected roughly proportional to their respective wavelengths. Typical values of differential multipath error for the differential carrier phase was found to be about 5 mm RMS by Cohen [Cohen, a]. The code phase multipath error is typically less than 1 meter. Unlike thermal noise, the multipath error can be highly correlated from epoch to epoch, especially for a static user. The spectral properties of multipath depend on the distance to the reflecting source(s). For one particular static geometry Pervan [Pervan, a] found most of the frequency content of the carrier phase multipath to be less than $\frac{1}{10} Hz$. In a dynamic environment, multipath is typically whiter due to rapidly changing geometry. The implication of the time correlation of multipath is that short-term averaging will not provide the error variance attenuation expected of a white error source. This is a drawback of carrier smoothed code techniques (Section 2.2.1.2) which rely on this averaging effect to reduce the code phase error. Only the thermal noise and “fast” components of multipath experience the improvements due to averaging.

3.4 Cycle Slips

As mentioned in Section 2.2, the carrier phase tracking receivers can measure the fractional phase of the incoming carrier relative to the receiver clock and can accumulate integer cycles as long as continuous lock is maintained. If one of the preceding error sources causes the fractional phase measurement error to become comparable to a half a cycle, the potential exists for the receiver to lose count of the number of integer cycles. When this happens, the

receiver is said to have experienced a *cycle slip* on that channel. In the vast majority of cases, cycle slips are detected by the receiver at the tracking loop level. In these cases, the cycle slip does not present an integrity threat, but only a potential continuity threat. In the event that a cycle slip goes undetected at the tracking loop level, the result is a potential loss of integrity. However, Pervan [Pervan, b] showed that such problems can be detected and sometimes isolated through the use of Receiver Autonomous Integrity Monitoring (RAIM).

3.5 Troposphere

The troposphere is the layer of the atmosphere in which we live, extending from the surface to an altitude of about 10 km. The signal delay caused by the troposphere is an error source affecting the carrier and code phase measurements equally. This delay is caused by the decrease in the speed of light in the troposphere. The speed of light is determined by the index of refraction (n):

$$c = c_{\text{vacuum}} / n \quad (3.2)$$

By definition, the vacuum value of the index of refraction is one. In the troposphere, the index of refraction increases with decreasing altitude. Due to this effect, the troposphere slows and deflects the satellite signal. The deflection is governed by Snell's law:

$$\frac{\sin(\theta_1)}{\sin(\theta_2)} = \frac{n_2}{n_1} \quad (3.3)$$




Figure 3.1: Snell's Law

Figure 3.1 defines the parameters of Snell's Law. As the signal enters a medium with a higher index of refraction, the angle of incidence is decreased. In the troposphere, the satellite signal is slowly deflected as the index of refraction increases. The total delay is a combination of the increased path length and the decreased speed of light. The effect of increased path length due to curvature is negligible for elevation angles greater than ten degrees [Janes]. However, the delay due to the decreased speed of light can be quite significant. The absolute delay is usually modeled as the absolute zenith delay (the delay experience for a satellite overhead) times an obliquity factor which is a function of elevation angle. This delay can be as big as 20 meters for a ten degree elevation satellite.

The index of refraction (n) is often expressed in terms of the refractivity (N):

$$N \equiv (n - 1) \times 10^6 \quad (3.4)$$

An empirical formula for the refractivity has been found [Janes]:

$$N = K_1 \left(\frac{P_d}{T} \right) + K_2 \left(\frac{e}{T} \right) + K_3 \left(\frac{e}{T^2} \right) \quad (3.5)$$

where:

$P_d \equiv$ partial pressure of the dry atmosphere.

$e \equiv$ partial pressure of water vapor.

$T \equiv$ Temperature.

$$K_1 = 77.61 \frac{K}{mb}$$

$$K_2 = 71.6 \frac{K}{mb}$$

$$K_3 = 3.747 * 10^5 \frac{K^2}{mb}$$

Due to local variations in water vapor, temperature, and pressure, it is difficult to estimate the differential troposphere delay. A first-order model for this delay is found by neglecting the curvature of the path and using a flat earth approximation. This “slab model” is shown in Figure 3.2. Using this model, the differential delay is only a function of aircraft height, satellite elevation angle (el), and average local index of refraction:

$$\delta_{tropo} = (\bar{n} - 1)d = \frac{(\bar{n} - 1)h}{\sin(el)} \quad (3.6)$$

where \bar{n} is the average index of refraction between the aircraft and the reference station. The differential zenith delay for an aircraft at 1000 feet is about 3 inches under standard conditions. The corresponding delay for a 15 degree elevation satellite is about 1 foot. Therefore, a 10% error in the average index of refraction can cause a phase error of 15% of an L1 cycle. Although this error seems quite small, integer search techniques are susceptible to this type of error source (Section 2.2.1.1). Also, local variations in the troposphere near weather fronts and boundary layers can cause the zenith delay to vary from point to point.




Figure 3.2: Troposphere Slab Model Geometry

3.5.1 Spacecraft

Although the slab model works well over short baselines, more accurate models for differential satellite troposphere can be derived by simply applying absolute models at both the reference station and at the aircraft. Many models for absolute troposphere signal delay exist, including those described in [Black], [Chao], [Hopfield], and [Saastamoinen]. Each of these methods models the absolute delay (D_t) as the product of a zenith delay (D_t^z) times an obliquity factor mapping function (M) which is a function of elevation angle:

$$D_t = D_t^z M \quad (3.7)$$

The IBLS user software implements a modified Hopfield troposphere model at the reference station and user to account for the differential troposphere.

3.5.2 Pseudolite

Pseudolite signals also experience tropospheric delays that should be modeled. Because IBLS uses the signals from the pseudolites only when they are nearby, the slab model discussed previously was used. That is, the delay experienced by the pseudolite signal (D_t^{PL}) was approximated by Equation (3.8).

$$D_i^{PL} = (\bar{n} - 1)|\bar{r}| \quad (3.8)$$

where:

\bar{n} is the average index of refraction between the aircraft and the pseudolite.

$|\bar{r}|$ is the range from the pseudolite to the aircraft.

For long range pseudolites like the Airport Pseudolites (APLs) discussed in Chapter 7, the effect of troposphere has not yet been quantified. Preliminary results indicate errors on the order of 10 parts per million are realistic (consistent with a 3% error in the total refractivity). This estimate corresponds to a 10 cm error over a 10 km baseline¹. It may be possible to use this effect to improve the estimate of refractivity by comparing the range profile (determined using kinematic DGPS) to the pseudolite phase profile, thereby calibrating refractivity and improving both pseudolite and satellite carrier phase measurements. Further study in this area is required.

For autonomous pseudolites developed by Cohen [Cohen, b], the pseudolite signal is phase-locked to an incoming satellite signal (Section 7.2.2). The differential tropospheric delay for this satellite between the aircraft and the pseudolite should also be incorporated into the correction for autonomous pseudolites.

¹ The “stacked APL” concept (Section 7.5) is relatively insensitive to this effect because it measures the *differential* range to a pair of pseudolites, not the absolute range.

3.6 Ionosphere

The ionosphere is the layer of the Earth's atmosphere extending from approximately 100 to 400 km above the surface. The free electrons in the ionosphere create a delay in the code phase and a corresponding advance in the carrier phase [Klobuchar]. As with the troposphere, the delay is typically modeled as a zenith delay times an obliquity factor:

$$I = I^v M(el) \quad (3.9)$$

where:

I is the code phase delay due to the ionosphere for a particular satellite and user.

I^v is the ionospheric code phase delay along a path from the satellite to the Earth's center.

$M(el)$ is an obliquity factor as a function of satellite's elevation angle.

The zenith delay varies from about two meters to as large as 20 meters; the obliquity factor ranges from one at 90 degrees to about three for low elevation satellites [Spilker]. Therefore the absolute ionospheric delay can be quite large. However, the *differential* delay is typically quite small (less than a few centimeters) over tens of kilometers [Klobuchar]. Therefore, the effect of the ionosphere on local area DGPS users is quite small with a few exceptions:

- 1) As with the differential troposphere, differential ionosphere errors are slowly changing with time and position. Therefore, users relying on small phase changes for cycle ambiguity convergence from satellite motion (Section 2.2.1.3) could be affected by the ionosphere.

2) Ionospheric scintillation can cause the absolute phase to fluctuate with time [Spilker].

Therefore, users implementing carrier-smoothed-code techniques (Section 2.2.1.2) *before* differencing will experience divergence in the code and carrier over short periods of time. They must therefore be careful to implement the carrier-smoothed-code filter in such a way that the differential station and the DGPS user have been affected identically by the ionosphere; otherwise the absolute ionosphere will not cancel when differenced. Practically speaking, this means that the carrier-smoothed-code filter should have a short time constant which will limit the performance of multipath and receiver noise averaging [LARC].

3.7 Ephemeris Errors

Errors in the broadcast satellite ephemeris affect DGPS users differently from the previous ranging errors. Unlike stand-alone users, DGPS users are unaffected (to first order) by satellite position errors along a line joining the user and the satellite. However, transverse satellite position errors result in an error in the unit line-of-sight vector (\hat{e}) in Equation (2.3). This unit vector error (\tilde{e}) can be substituted into Equation (2.3) to find the equivalent phase measurement error ($\tilde{\phi}$):

$$\tilde{\phi} = -\tilde{e} \cdot \bar{x} \quad (3.10)$$

Therefore, an error in the unit line-of-sight vector manifests itself as an equivalent *ranging error that scales with the distance* between the DGPS user and the reference station in the direction of that error. The scale factor is the transverse position error of the spacecraft divided by the distance to the spacecraft. For example, a 200 m error for a satellite about 20000 km away will result in a 10 cm range error over a 10 km baseline. Although, satellite

ephemeris errors are typically less than 5 m RMS [Zumberge], this effect should be considered when designing landing systems. Due to the limited observability of such errors over short baselines, local area ground monitoring may be complicated by the potential for such errors.

3.8 Differential Correction Latency

Differential correction latency refers to the delay between the time a measurement is taken by the DGPS reference station until the time the corresponding correction is received at the aircraft. Given the two second time-to-alarm for Cat III precision approach (Table 1.2), this latency should be less than a second. Cobb discusses the potential for an autonomous pseudolite to provide differential corrections with minimum (less than a millisecond) latency [Cobb, a]. However, more traditional digital data links will experience delays on the order of half a second. These delays necessitate the use of reference phase predictors (Section 4.2.1) or the equivalent. The error in the prediction process is mainly a function of the smoothness of the satellite clock error, which is currently dominated by SA. Section 4.2.1 presents the performance results for a reference carrier phase predictor.

3.9 Circular Polarization

[SPS] specifies that the GPS signal is right hand circularly polarized. Therefore, the carrier phase measured for a GPS satellite is not only a function of antenna position, but antenna orientation as well. This effect is ignored for most applications because the phase change due to orientation is approximately the same for all satellites (and will therefore only affect the clock bias estimate). For IBLS and dual antenna positioning, the effect is larger and

should be modeled. A simple model to account for circular polarization is given in Section 6.4.

3.10 Small Corrections

For simplicity, some small, but non-negligible terms were not included in Equation (2.3). These terms, which were included in the IBLS airborne measurement processing software, are discussed in this section for completeness.

3.10.1 Sampling Time

The aircraft receiver and reference receiver often do not sample at the same time. Typical architectures allow the sample times to drift with respect to GPS time up to a certain threshold at which point a step is introduced to correct for the drift. For example, say a receiver is sampling at 1 Hz relative to its own clock which is slow by one part in a million. If it starts sampling on the GPS second, after about 15 minutes it will be sampling one millisecond after the GPS second. A common remedy to this problem is to incorporate digital logic to introduce a *leap millisecond* so the receiver is again sampling on the GPS second. This ensures that a DGPS reference receiver and roving user receiver always sample within two milliseconds of the same time. Even this small difference must be accounted for either by explicit phase interpolation or by *implicit timetagging*, a technique developed by Pervan [Pervan, a]. To perform implicit timetagging, the phase rate measured at the reference station (expressed in dimensionless units by dividing by frequency) for each satellite is subtracted from the column of ones in the GPS geometry matrix corresponding to the differential clock state. This implicitly multiplies the phase rate of each satellite by the difference in receiver sampling times, thereby converting to a common sample time.

3.10.2 Spacecraft Doppler

Spacecraft velocity along the line joining the spacecraft with the user has the same effect on a kinematic DGPS user as a frequency offset in the spacecraft clock. Both cause a doppler shift that lengthens or shortens the carrier phase wavelength. If left unmodeled, this will result in an equivalent phase error that scales with distance between the user and the DGPS reference station along the line-of-sight to the spacecraft:

$$\tilde{\phi} = (\hat{e}_i \cdot \bar{x}) \left(\frac{\lambda_i^{assumed} - \lambda_i^{actual}}{\lambda_i^{actual}} \right) \quad (3.11)$$

where:

$\lambda_i^{assumed}$ is the assumed wavelength for the i^{th} satellite (neglecting doppler).

λ_i^{actual} is the actual wavelength for the i^{th} satellite.

The resulting error is small, only a few centimeters over a 10 km baseline for a doppler of 5kHz. Nevertheless, all of the information needed to account for this effect is available, so it was incorporated into the IBLS software.

3.10.3 Signal Travel Time

From the time the GPS signal leaves the satellite until the time it is received by a terrestrial user, approximately 70 msec have elapsed. During this time, both the spacecraft and the earth have moved. In Equation (2.3) the unit line-of-sight vector (\hat{e}_i) should be expressed taking this motion into account. In particular, \hat{e}_i should represent the direction from the current user's position to the satellite's position when the signal was sent, expressed in a reference frame attached to the current orientation of the earth.

If this effect is ignored, the resulting error is equivalent to a spacecraft ephemeris error. The spacecraft move about 30 degrees per hour or about 1×10^{-5} radians during the time-of-flight of the signal from the spacecraft. The phase measurement error could therefore be as large as 10 cm over a 10 km baseline.

3.10.4 Linearization Error

The linearization performed to obtain Equation (2.3) assumes the satellite is so far away that the line of sight vector to the satellite is the same for both the DGPS reference station and the roving user. Although this approximation is true to first order, the second order terms are not negligible over baselines as short as 500 meters and increase with the square of the baseline distance normal to the line-of-sight direction. Figure 3.3 shows the geometry associated with this error. The approximate error due to linearization ($\tilde{\phi}_{lin}$) is given by Equation (3.12):

$$\tilde{\phi}_{lin} \cong \frac{d^2}{2R_{sat}} \quad (3.12)$$

where:

d is the baseline distance normal to the direction to the satellite.

R_{sat} is the distance to the satellite.

For a five kilometer baseline, this error is approximately half a meter.

Referring to Figure 3.3, the exact expression for the error due to linearization can be expressed:

$$\tilde{\phi}_{lin} = \Delta\rho - \Delta\rho_{lin} \quad (3.13)$$




Figure 3.3: Linearization Error

This error is calculated and applied as a correction in the IBLS software. Due to uncertainty in the exact position of the user, the corrected phase will also be in error. This error ($\tilde{\phi}_{lin,corr}$) will scale with d times the associated error (\tilde{d}):

$$\tilde{\phi}_{lin_{corr}} \equiv \frac{d\tilde{d}}{R_{sat}} \quad (3.14)$$

However, assuming a three meter error in \tilde{d} , consistent with the typical accuracy of code phase DGPS, the error in the corrected phase is less than three millimeters over a 20 kilometer baseline.

CHAPTER 4

Airborne Measurement Processing

Recent flight tests of the Integrity Beacon Landing System (IBLS) have demonstrated the feasibility of using GPS for Category III precision landing. To achieve these results, an airborne architecture that provided position solutions in real-time was developed. *Centimeter-level positioning accuracy was achieved using a single-frequency receiver without using integer search techniques.* This capability distinguishes the Integrity Beacon Landing System from other proposed kinematic GPS landing systems.

At the heart of the real-time architecture is a cycle ambiguity estimator. This estimator makes use of all available information to arrive at floating estimates of the integer biases associated with the GPS carrier phase measurements. The uncertainty in these estimates is stored in a covariance matrix. The estimates are updated using carrier and code phase measurements from both satellites and pseudolites.

New satellites are added to the estimate and lost satellites are removed from the estimate with ease. Given redundant satellites, the estimator will converge toward the cycle

ambiguities using satellite motion. Using seven satellites, the integer estimates typically converge to the cycle level in 15 minutes. During the pseudolite overpass, the estimates converge to the centimeter level in a matter of seconds. Receiver Autonomous Integrity Monitoring (RAIM) is performed during the pseudolite overpass to verify the consistency of the satellite and pseudolite measurements. Additionally, in all phases of flight RAIM is performed before each integer update to verify that the update is consistent with the existing integer estimates. Despite the flexibility of this architecture, it is straightforward to implement. The details of this implementation are presented in this chapter.

4.1 Background

For early flight tests of IBLS [Cohen, c], system evaluation was performed entirely in post-processing. The next step in the development of the IBLS test platform was to add a data link for differential corrections and allow position solutions to be calculated in flight [Cohen, b]. However, this original system still had several limitations:

- Positions were only available at the data uplink frequency and were delayed by the data link latency.
- Positions were only calculated when the integers were known; no positions were available before the bubble.

Since the ultimate goal of the IBLS flight tests was to demonstrate automatic approaches and landings, these limitations had to be overcome. The current version of the IBLS architecture developed in this research overcomes these limitations in an elegant yet simple

manner. Category III feasibility demonstrations using this architecture are described in [Cohen, d] and [Cohen, e].

4.2 Airborne Architecture

The current version of the IBLS architecture is shown in Figure 4.1. The inputs to the airborne software come from the aircraft GPS receiver, the data link, and an attitude source (attitude is required only to account for the lever arm from the top GPS antenna to the pseudolite receive antenna on the bottom of the aircraft). The time-tagged data from each of these inputs is stored in ring buffers. The newest data on the aircraft receiver ring buffer is used immediately to obtain a position solution. Because the reference phase is delayed and available only at the uplink frequency, its value must be predicted from the reference phase history in the data link ring buffer. Also, an estimate of the integers is required to perform a carrier phase position solution. This estimate is continuously updated using aligned data from the ring buffers. That is, when a timetag is found that is common to all of the ring buffers, the synchronous data is used to improve the estimate of the integers. This process is described in more detail in Section 4.2.2.

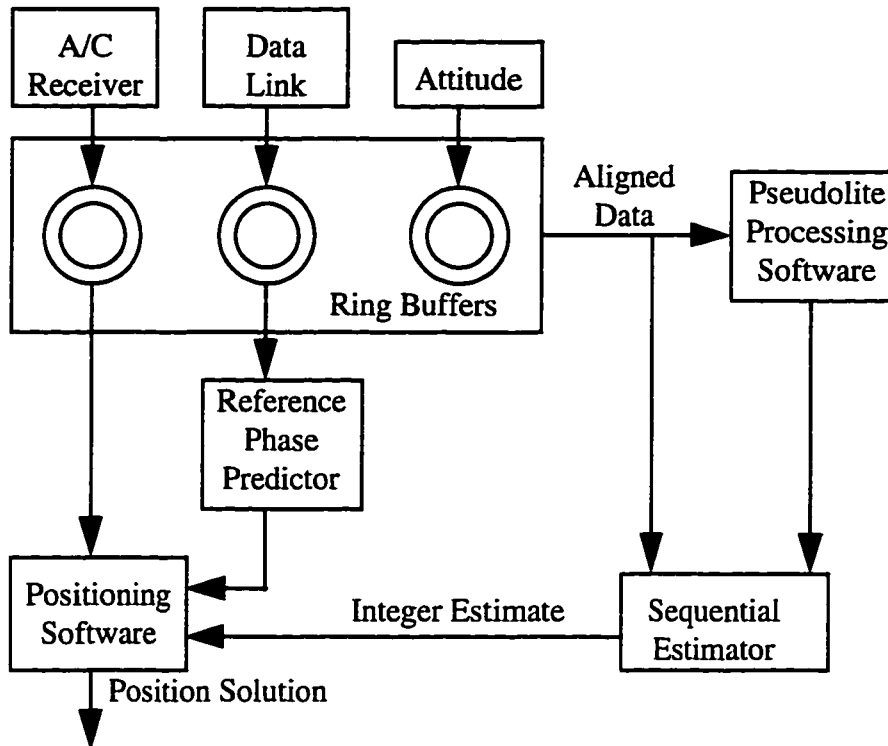


Figure 4.1: IBLS Airborne Architecture

4.2.1 Reference Phase Predictor

Instead of waiting for the data link to send up the carrier phase measured at the reference station, the reference phase can be predicted based on its history. Because the reference phase is relatively slowly changing (satellite motion and SA dominate the dynamics of the reference signal), it can be accurately predicted several seconds into the future. The goal of the reference phase predictor was to perform this prediction in a simple yet robust manner. Although it is possible to predict more than several seconds in advance, a data latency or data dropout of that magnitude would most likely be unacceptable for the precision landing application. Therefore, the requirements for the predictor were derived to:

- Predict the reference phase for several seconds with error on the order of centimeters.

- Be robust to data dropouts.
- Be easy to implement.
- Predict phase at arbitrary timetags (i.e. - not just at reference sample times).

Although several types of predictors were evaluated [Masson], a least-squares quadratic predictor was chosen because it was simple, yet it performed quite well. A quadratic function of time is fit to the previous m reference phases ($m \geq 3$). Given the timetag of an aircraft receiver measurement, the corresponding reference phase can be predicted. The performance of the predictor is given in Table 1. The table was generated by finding the difference between the predicted phase and the actual phase for different values of m and different data latencies. The data uplink frequency was one hertz. The predictor errors are quite acceptable for data latencies less than 5 seconds and m between five and eight. As expected, the prediction error starts increasing for larger values of m because older data is weighted the same as more recent data. A weighted least-squares predictor which decreased the weights with the age of the data was also evaluated. The performance improved slightly, but the improvement was not worth the additional complexity.

Table 4.1: Reference Phase Predictor Error

RMS Error (cm)	1 sec	2 sec	3 sec	5 sec	10 sec
$m=3$	1.09	2.58	4.68	10.74	36.70
$m=4$	0.75	1.53	2.55	5.57	18.50
$m=5$	0.63	1.12	1.82	3.86	13.08
$m=6$	0.56	0.95	1.52	3.24	11.32
$m=7$	0.55	0.93	1.47	3.15	11.14
$m=8$	0.56	0.95	1.52	3.24	11.43

Error in prediction for various latencies. m is the number of points (separated by 1 second) used to fit the quadratic predictor. Actual data was used.

4.2.2 Integer Estimation

To use the differential carrier phase to perform position solutions, an estimate of the integer cycle ambiguities is required. Recall from Equation (2.4) that the carrier phase measurement equation can be written:

$$\phi = \begin{bmatrix} G & I \end{bmatrix} \begin{pmatrix} \bar{x} \\ \tau \\ N \end{pmatrix} + \delta\phi \quad (4.1)$$

where ϕ ($n \times 1$) is the single-difference carrier phase measurement (expressed in L1 wavelengths), G ($n \times 4$) is the traditional GPS geometry matrix, I ($n \times n$) is the identity matrix, \bar{x} (3×1) is the position, τ (scalar) is the differential receiver clock bias, N ($n \times 1$) are the integer cycle ambiguities, $\delta\phi$ ($n \times 1$) are the measurement errors, and n is the number of satellites.

If an integer estimate, \hat{N} , and the corresponding covariance, P_N , are available, this equation may be rewritten:

$$\phi - \hat{N} = G \begin{pmatrix} x \\ \tau \end{pmatrix} + (\tilde{N} + \delta\phi) \quad (4.2)$$

where \tilde{N} is the error in the integer estimate. Assuming the measurement errors are uncorrelated with variance σ^2 , a weighted least-squares position/clock estimate may be calculated:

$$\begin{pmatrix} \hat{x} \\ \hat{\tau} \end{pmatrix} = [G^T R_e^{-1} G]^{-1} G^T R_e^{-1} (\phi - \hat{N}) \quad (4.3)$$

with covariance: $P_{x\tau} = [G^T R_e^{-1} G]^{-1}$

where $R_e \equiv \sigma^2 I + P_N$.

It is therefore convenient to keep a running estimate of the integers. An integer estimator was developed with the following goals in mind:

- Eliminate the need for mode switching from code to carrier differential GPS after the bubble pass.
- Maintain an estimate and covariance of the integers.
- Update estimates using all available information, including integrity beacon results.
- Bring satellites on and off line gracefully as they are acquired and lost.

- Allow for continuous RAIM.
- Be easy to implement.

Implementing the estimator as described here, these primary goals are met. The resulting architecture is quite flexible, and incorporates the useful features discussed in Section 4.2.6.

4.2.2.1 Estimate Initialization

When the program is first started, the integer estimates are initialized using the differential code phase measurements. Subtracting Equation (2.5) from Equation (2.4):

$$\phi - \phi_{code} = N + \delta\phi - \delta\phi_{code} \quad (4.4)$$

where ϕ_{code} is the code phase *expressed in LI cycles*. The initial integer estimate is therefore given by Equation (4.5).

$$\begin{aligned} \hat{N} &= \phi - \phi_{code} \\ P_N &= (\sigma^2 + \sigma_{code}^2)I \cong \sigma_{code}^2 I \end{aligned} \quad (4.5)$$

where:

\hat{N} is the integer estimate.

P_N is the integer estimate error covariance.

σ is the one sigma differential carrier phase measurement error.

σ_{code} is the one sigma differential code phase measurement error.

To illustrate this initialization process, Figure 4.2 shows a representation of the initial estimate for the simple example of Figure 2.2. The integer estimate is represented in integer space as an ellipse; the center of the ellipse is the estimate (\hat{N}) and the size and orientation

represent the error covariance (P_N). The ellipse is a circle in this case because it is assumed that the measurement quality for each satellite is identical and uncorrelated. More complicated measurement error models might put more error on lower elevation satellites.




Figure 4.2: Integer Estimate Initialization

4.2.3 Integer Measurement Updates

After the integer estimates are initialized in this manner, they are refined using a variety of sources. It is important to note that only the aligned measurements from Figure 4.1 are used to update the integer estimates (error in the reference phase predictor does not corrupt the estimates). In all cases, the estimate update is performed by casting the new information into the following form:

$$\begin{aligned} z &= HN + v \\ E[vv^T] &= R \end{aligned} \quad (4.6)$$

where:

z is the new measurement, which has been transformed to be a function only of the cycle ambiguities.

H is the observation matrix for the new measurement.

v is the measurement error on z .

In this form, the measurement is used in a minimum variance measurement update [Gelb]:

$$\begin{aligned} K &= P_N^- H^T [H P_N^- H^T + R]^{-1} \\ \hat{N}^+ &= \hat{N}^- + K(z - H\hat{N}^-) \\ P_N^+ &= [I - KH] P_N^- \end{aligned} \quad (4.7)$$

where the '-' and '+' in the superscripts indicate before and after the measurement update respectively. No process update is necessary, because the states being estimated are constants. The measurement update process is shown in Figure 4.3.

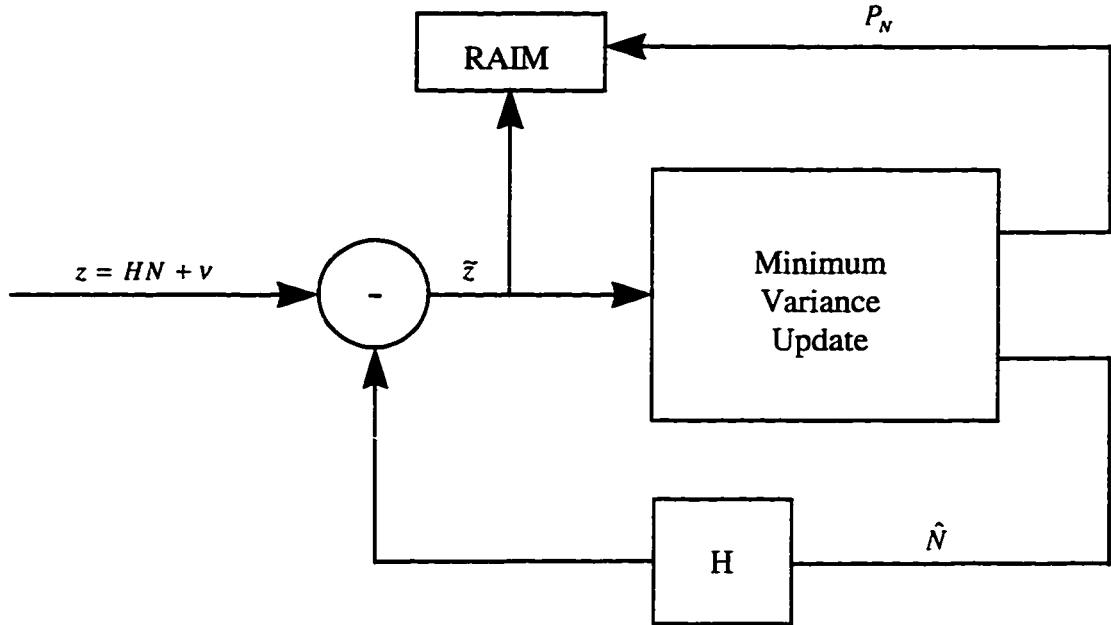


Figure 4.3: Integer Update Process

4.2.3.1 Carrier Phase Update

If there are more than four satellites, each differential carrier phase measurement can be cast into the form of Equations (4.6). Pre-multiplying Equation (4.1) by L , where the rows of L make up an orthonormal basis for the left null space of G :

$$L\phi = LN + L\delta\phi \quad (4.8)$$

Equation (4.1) has been transformed to look like Equations (4.6) with:

$$z = L\phi$$

$$H = L$$

$$R = \sigma^2 LL^T = \sigma^2 I$$

This pre-multiplication by L eliminates the position and clock terms from the measurement equation and separates integer estimation from position estimation.

For the example of Figure 2.2, Equation (2.7) is transformed from two equations for three unknowns (x, N_1, N_2) to a scalar equation with two unknowns when pre-multiplied by L . The first carrier phase update after initialization is shown in Figure 4.4. The circle represents the a priori estimate. The new information provided by Equation (4.8) can be thought of as an independent estimate of the integers. Because there is one more equation than unknown, the estimate ellipse is infinitely long as seen in the figure. Combining the two estimates using Equations (4.7), the new estimate covariance is represented by the small ellipse.




Figure 4.4: Carrier Phase Update

4.2.3.2 Code Phase Update

Just as the code phase measurements were used to initialize the integer estimates, they can be used to update the estimates. Equation(4.4) is already in the desired form of Equation (4.6) with:

$$z = \phi - \phi_{code}$$

$$H = I$$

$$R = (\sigma^2 + \sigma_{code}^2)I$$

Updating the carrier phase cycle ambiguity estimates using differential code phase measurements is called code-centered-carrier filtering and is similar to carrier smoothed code (Section 2.2.1.2). Both techniques make use of the advantages of each measurement. The carrier phase has low noise but an integer bias; the code phase has higher noise but no integer bias. One technique uses code phase measurements smoothed by the carrier, while the other uses the carrier phase centered by the code.

Figure 4.5 illustrates the code phase update process for the 2D example of Figure 2.2. In this figure, the first update after initialization is a code phase update. One of the two larger circles represents the initial estimate while the second represents the new information. The smaller ellipse represents the combined estimate. Independence of the initial measurement and the new information is assumed; code-centered-carrier techniques are susceptible to the time correlation of the code phase measurements. This correlation can be simply modeled by updating the covariance at a rate comparable to the decorrelation rate.




Figure 4.5: Code Phase Update

4.2.3.3 Integrity Beacon Update

The rapid geometry change that occurs during an integrity beacon overpass provides another update to the integer estimator. The output of the integrity beacon processing software is a high accuracy estimate of the individual satellite integer differences along with the corresponding covariances. The integers themselves are unobservable, but only the integer differences affect the position solutions. Any common bias affects only the clock solution. Pervan developed the algorithms and software contained in the pseudolite processing software module in Figure 4.1. This module implements a non-linear estimator to determine the integer differences from the integrity beacon measurements. The details of this implementation are described in [Pervan, a]. The output of this non-linear estimation

process can be used as a *linear* update to the sequential estimator. As with the other integer updates, the new information can be written in the form of Equations (4.6):

$$z = HN + v$$

$$H = \begin{bmatrix} -1 & & \\ \vdots & & I \\ -1 & & \end{bmatrix}$$

H ($n-1 \times n$) reflects the fact that only integer differences are output by the integrity beacon processing code; the integer of an arbitrary satellite is subtracted from all of the rest. The R used in the update equations is simply the integer difference covariance output by the integrity beacon processing code. R is a function of the overpass geometry. The geometry change that occurs as an aircraft flies over a pair of integrity beacons provides information comparable to several hours of satellite motion in a matter of seconds. However, due to the nonlinear nature of this geometry change (G is a function of \bar{x} for the integrity beacons), this information must be processed separately. The integrity beacon measurement processing software was developed by Pervan and is described in [Pervan, a]. The result of this process is a high accuracy, high integrity estimate of the integer differences. This estimate is used to update the overall integer estimate, as described in Section 4.2.3.

4.2.4 Constellation Changes

As an aircraft banks, it may lose the signal from some satellites and acquire the signal from others. Satellites may also be acquired or lost as they pass the receiver elevation mask angle. It is therefore desirable to remove satellites and add satellites to the integer estimate. The covariance form of the estimator was chosen primarily because it makes this task quite simple. To remove a satellite, the corresponding state is removed. The element of the

estimate and the row and column of the covariance are simply discarded. To bring a new satellite on-line, the integer estimate is initialized using the code phase measurement for that satellite:

$$\hat{N}^{new} = \phi^{new} - \phi_{code}^{new}$$

The variance for the new integer is set consistent with the code phase measurement noise; the cross covariance for the other integer states is set to zero:

$$E[\tilde{N}^{new} \tilde{N}^{new}] = (\sigma^2 + \sigma_{code}^2)$$

$$E[\tilde{N}^{new} \tilde{N}^{old}] = \bar{0}$$

After one carrier phase measurement update, the state estimate and covariance for the new satellite are consistent with the other integer estimates. Thus, if more than four satellites are in view, the integer for the new satellite is known as accurately as the others. This technique has proven to be an efficient method for handing-off satellite integer estimates.

4.2.5 RAIM

RAIM is performed before each integer update to verify that the new measurement is consistent with the existing integer estimates. If the measurement does not pass this check, the approach can be aborted. In some cases, the failure may be isolated [Pervan, b].

In preparation for an integer update, the new measurement is already in the form of Equation (4.6). The difference between the expected measurement and the actual measurement is calculated:

$$\begin{aligned} r &\equiv E[z] - z = H\hat{N} - z \\ r &= H(\hat{N} - N) - v = H\tilde{N} - v \end{aligned} \quad (4.9)$$

This residual quantity, r , is a random vector with zero mean and covariance given by:

$$P_r = HP_N H^T + R$$

A measure of the consistency of the new measurement is the weighted residual:

$$w = r^T P_r^{-1} r \quad (4.10)$$

If this weighted residual is greater than some predetermined threshold, a RAIM alert is issued. The threshold is a function of the dimension of r and the desired continuity.

4.2.6 Other Features

4.2.6.1 Satellite Motion

Loomis [Loomis] and Hwang [Hwang] pointed out that the differences between the integers are observable in a dynamic environment because of satellite motion. A welcome side-effect of this airborne architecture is that the integer differences will tend to converge using satellite motion. This convergence results from the carrier phase measurement updates described earlier. Satellite motion is automatically taken into account each time there is a carrier phase measurement update, because the L matrix in Equation (4.8) changes with

time. Figure 4.6 shows the results of satellite motion for the example of Figure 2.2. As the satellite geometry changes, the orientation of the ellipse representing the new carrier phase information changes. The combined estimate for two carrier phase measurements is represented by the smaller ellipse in Figure 4.6.




Figure 4.6: Satellite Motion

Although the implementation of the carrier phase updates are performed sequentially, the observability analysis is shown here for a batch solution. As the satellite geometry changes, the L matrix in Equation (4.8) also changes. Grouping an arbitrary satellite integer with the clock bias term, Equation (4.1) may be rewritten as in Equation (4.11).

$$\phi = \begin{bmatrix} G & \bar{I} \end{bmatrix} \begin{pmatrix} x \\ \tau + N_1 \\ N_2 - N_1 \\ \vdots \\ N_n - N_1 \end{pmatrix} + \delta\phi \quad (4.11)$$

$$\text{where } \bar{I} \equiv \begin{bmatrix} 0 & \cdots & 0 \\ & I & \end{bmatrix}$$

Pre-multiplying Equation (4.11) by the left null space of G :

$$z = L\phi = L\bar{I}N' + L\delta\phi \quad (4.12)$$

$$\text{where } N' \equiv \begin{pmatrix} N_2 - N_1 \\ \vdots \\ N_n - N_1 \end{pmatrix}$$

Several of these measurements can be stacked together for different times:

$$z_* \equiv \begin{pmatrix} z_1 \\ \vdots \\ z_k \end{pmatrix} = L_* \bar{I} N' + \begin{pmatrix} L_1 \delta\phi_1 \\ \vdots \\ L_k \delta\phi_k \end{pmatrix} \quad (4.13)$$

$$\text{where: } L_* \equiv \begin{bmatrix} L_1 \\ \vdots \\ L_k \end{bmatrix}$$

The reason that only integer differences are observable while the integers themselves are unobservable is that L_* never has rank greater than $(n-1)$. The columns of L always sum to zero, as can be seen by manipulating the definition of the left null space of G :

$$LG \equiv [0]$$

$$\begin{bmatrix} \bar{L}_{c_1} & \cdots & \bar{L}_{c_n} \end{bmatrix} \begin{bmatrix} G_x \\ \vdots \\ 1 \end{bmatrix} = [0]$$

$$\Rightarrow \sum_{j=1}^n \bar{L}_{c_j} = \bar{0}$$

Therefore, the sum of the columns of L_* are also constrained to sum to zero, and the maximum rank is $(n-1)$. However, only integer differences are required to solve for position. Given sufficient geometry change and enough redundant satellites, the matrix $L_* \bar{I}$ will have rank $(n-1)$ and Equation (4.13) may be used to estimate N' :

$$\hat{N}' = [\bar{I}^T L_*^T L_* \bar{I}]^{-1} \bar{I}^T L_*^T z_* \quad (4.14)$$

If the measurement samples are widely spaced enough in time, the noise will be uncorrelated and the estimate error covariance reduces to:

$$E[\tilde{N} \tilde{N}'^T] = D \sigma^2 \quad (4.15)$$

where $D \equiv [\bar{I}^T L_*^T L_* \bar{I}]^{-1}$.

The matrix D is similar to dilution of precision (DOP). The square root of the trace of D is analogous to PDOP and is referred to as NDOP. The quantity $(NDOP \times \sigma)$ approximates the one-sigma integer estimate error. A typical value of NDOP using seven or more satellites and two measurements separated by 15 minutes is 20. Assuming a carrier phase measurement error of 0.5 cm, the one-sigma integer estimate is typically 10 cm after 15 minutes of satellite motion. In contrast, NDOP for a 10 second integrity beacon overpass is

about 2. The information provided by the integrity beacon clearly dwarfs that provided by even long observations of satellite motion. Nevertheless, satellite motion is a welcome complement to the integrity beacon overpass.

In the sequential implementation of the carrier phase measurement update, it is not necessary to group one satellite's integer with the clock bias. Although only integer differences are observable from satellite motion, the integers themselves are initialized from code phase measurements. The implication is that one direction of the integer covariance will remain at its initial value (neglecting code phase measurement updates). Scaling problems could arise as the minimum eigenvalue decreases while the maximum eigenvalue remains the same. This issue does not present a practical limitation because adding a small amount of process noise to the covariance prevents the minimum eigenvalue from decreasing without bound.

4.2.6.2 Static Survey

For experimental purposes, it is often convenient to know the integers before the aircraft takes off. For this reason, a static survey mode was added to the estimator. When the user changes to this mode, the estimator assumes the aircraft is not moving. This static constraint allows the estimator to converge faster and with fewer satellites. When the static constraint is imposed, the integer state is augmented with the position. The implementation of this static mode is described in [Lawrence, b]. With seven satellites, convergence in static mode is typically two or three times faster than dynamic mode (with fewer satellites,

static mode is considerably faster; in the case of four satellites, dynamic mode will never converge from satellite motion). Integer estimates are often within a cycle of their true value in less than five minutes. Before the aircraft starts moving, the user can simply push a button, discarding position from the state estimate to return to dynamic mode.

4.2.6.3 Known Position Updates

If the user knows the aircraft's position, this position can be incorporated into the integer estimate. The position "measurement" can be written:

$$z \equiv L_{\tau}(\phi - G_x \hat{x}) = L_{\tau} N + G_x \tilde{x} + \delta\phi \quad (4.16)$$

which is in the form of Equation (4.6) with:

$$H = L_{\tau} \equiv null([1 \quad \dots \quad 1])^T$$

$$R = G_x P_x G_x^T + \sigma^2 I$$

where:

$$G_x \equiv \begin{bmatrix} -\hat{e}_1^T \\ \vdots \\ -\hat{e}_n^T \end{bmatrix}$$

\hat{x} is the estimated position.

\tilde{x} is the error in that estimate.

P_x is the confidence in \hat{x} expressed in covariance form.

The "measurement noise" is the uncertainty in the position. The user enters the position and a covariance matrix representing this uncertainty. An example of when this feature is useful is when the aircraft is parked at the tie-down location. Each time the aircraft is there, its

vertical position is the same to within a few centimeters while the horizontal position may be different by several meters. This uncertainty can be accurately entered into the estimate. After leaving the tie-down, the vertical position error and covariance will remain small, while the horizontal error will slowly converge. This feature was used during a flight test discussed later.

4.3 System Performance

This real-time system has been tested extensively. In July of 1994, it was used to perform 49 autocoupled approaches of an FAA King Air [Cohen, d]. The most impressive results came in October of 1994 when this system was used to perform 110 automatic landings of a United Airlines Boeing 737-300 [Cohen, e]. The following sections describe experiments designed to exercise several aspects of the system.

4.3.1 24 Hour Positioning Test

The real-time system was set up in the lab; the reference station and aircraft receivers were connected to separate antennas on the roof. To accentuate multi-path errors, ground planes were not used. Due to the antenna gain pattern and cable loss, satellites were not acquired until they reached an elevation angle of about fifteen degrees. The system was initialized and data was collected for 24 hours. The first goal of this experiment was to demonstrate the integer estimator convergence using satellite motion. Although the baseline was static, the static survey mode of the estimator was not used. The second goal was to smoothly hand-off integer estimates for 24 hours.

Figure 4.7 plots the magnitude of the 3D position error over the 24 hour period. The initial error was about three meters (off the vertical scale of the plot). Using satellite motion, the position error converged to the cycle level in about fifteen minutes. In less than an hour, it converged to less than half of a cycle where it stayed for the remainder of the test. After the first hour, the mean value of the magnitude of the position error was 2.2 cm.




Figure 4.7: 24 Hour Positioning Test

Twenty-five different satellites were used during the 24 hours. Some satellites were brought on and off line several times while their signal strength was low. Satellite integer estimates were brought on-line a total of 192 times. These hand-offs were performed seamlessly as evidenced by the position error plot of Figure 4.7. After the estimator converged on the integers for one set of satellites, the position error remained small using entirely different sets of satellites. By the end of the test, the original satellites returned to their initial positions in the sky. The implication of this periodic geometry is that the

satellites could be handed off indefinitely while retaining centimeter-level positioning accuracy.

4.3.2 Airborne Performance

To exercise the system in an airborne environment, a flight test was performed in a Piper Dakota. In contrast to the static test discussed earlier, the true position of the aircraft in flight is not known exactly. However, if a separate process knows the correct values of the rounded integers, a centimeter-level truth trajectory can be found. This trajectory can be compared with the trajectory estimate calculated using the integer estimator and reference phase predictor to evaluate the system performance.

To find the integer differences for the truth trajectory, the static survey mode of the estimator was used while the aircraft was at the tie-down. In less than fifteen minutes, the estimator converged to within a half-cycle of the correct integer differences. The integer difference were rounded to the correct values. The measurement residual was monitored for another fifteen minutes to verify that they were correct. The flight test was performed when the satellite geometry was such that six satellites could be continuously tracked, even when the aircraft banked. This ensured that the integer differences used for the truth trajectory were constant throughout the flight.

The real-time system was reset before the flight (the integer differences found from the static survey were only used to find the truth trajectory). The integers were initialized as usual, using differential code phase measurements. As mentioned earlier, the vertical position at the tie-down is well known. To demonstrate the “known position update” feature described earlier, this vertical information was incorporated into the estimate. After the position update, the vertical position was accurate to a few centimeters, while the horizontal position was wrong by several meters (as intended). During the flight, the horizontal position should converge toward the correct value, while the vertical position should remain accurate.

Six spacecraft were in view during the entire flight. However, to provide more of a challenge to the system, satellite outages were simulated in software. Every four minutes, one satellite was removed for a period of thirty seconds. During the 26 minute flight, each satellite was removed once. Therefore, each integer was taken off-line and brought back on-line.

The flight test consisted of three take-offs and landings. Figure 4.8 plots the horizontal and vertical position error and the 1σ horizontal position error bound. The times of the simulated satellite outages are marked with an “x” on the time axis. As expected, the vertical error remained small while the horizontal error converged.




Figure 4.8: Airborne Position Error

CHAPTER 5

Navigation Performance Assessment

The primary goal of this research was to assess the navigation performance of GPS-based precision approach systems and explore methods to improve that performance. This chapter discusses the accuracy assessment of two systems: IBLS and the Stanford WADGPS test-bed. The nominal accuracy of both systems is shown to exceed that necessary for their respective applications (Cat III for IBLS and Cat I for WADGPS). Chapters 6 and 7 explore possible improvements to GPS-based positioning.

5.1 Integrity Beacon Landing System Performance

IBLS has been extensively tested to verify its accuracy. Several different types of accuracy checks are available, including static survey comparison, laser comparison, and consistency checks. Additionally, the closed-loop performance for autocoupled approaches and landings is presented.

5.1.1 Static Survey Comparison

After performing cycle ambiguity resolution for a set of satellites, the same high accuracy is maintained for as long as they are continuously tracked. If the aircraft then lands and comes to a stop, its final position is also accurately known. This position solution may be checked

against an independent static survey to verify successful cycle ambiguity resolution. This experiment was performed during the first real-time flight test of the IBLs system. The integrity beacons were located on the Stanford campus. Figure 5.1 shows the trajectory followed by the aircraft as it flew through the bubbles and returned to Palo Alto Airport. After the aircraft was tied down, the IBLs position solution relative to the reference station at Stanford was noted. This same baseline was then independently calculated using established GPS static survey techniques capable of millimeter-level accuracy. The solutions agreed to within 2 cm. The implication of this result was that the cycle ambiguities were successfully resolved and that the entire trajectory of the aircraft was known to the centimeter level in real time.




Figure 5.1: Static Survey Comparison

5.1.2 Laser Altimeter Comparison

Figure 5.2 shows the Piper Dakota flight test vehicle. A laser range finder was installed in the Dakota to provide an independent measure of height. The accuracy of this laser altimeter was specified to be ± 5 cm. To compare the laser altimeter data with IBLS position fixes, an accurate model of the runway height was needed. A kinematic survey of the Palo Alto runway was performed to arrive at this model. A surface was fit to this kinematic survey data to find a precise model of runway height as a function of horizontal position. Using 15 coefficients, the surface shown in Figure 5.3 was generated. The surface matched the recorded data to 4 cm, 1σ .




Figure 5.2: Piper Dakota

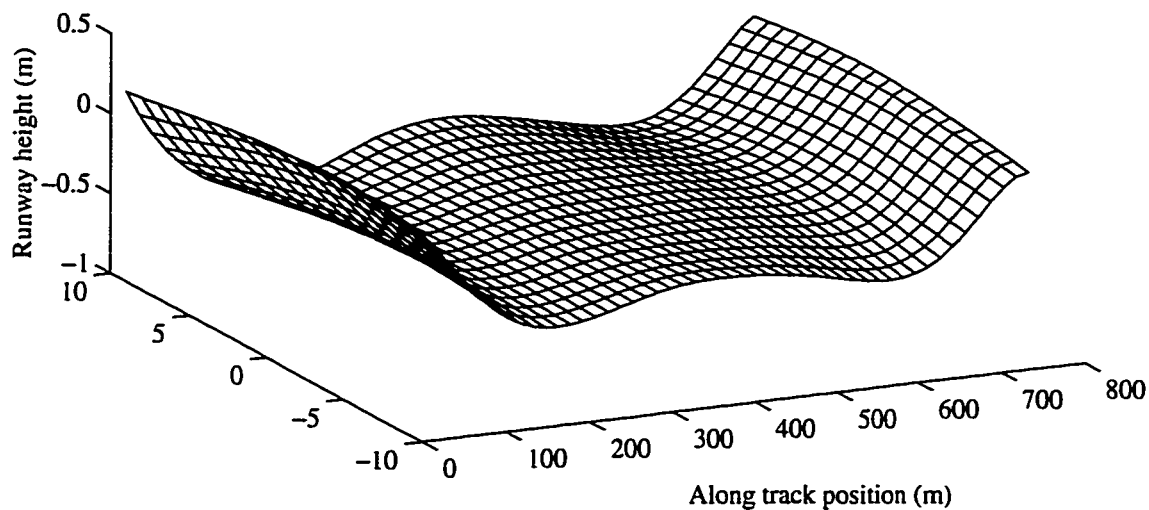


Figure 5.3: Runway Surface Fit

To synchronize the IBLS measurements with the laser measurements, the laser was set up to accept trigger pulses from the GPS receiver. Unfortunately, an undocumented lag in the laser altimeter of 0.4 seconds was discovered after the flight test. Therefore, the 1 Hz laser measurements occurred 0.4 seconds after the IBLS measurements. To correct for this timing problem, the laser data was interpolated quadratically. However, only during regions of low vertical acceleration disturbances is this interpolation effective.

Given synchronous measurements of the position of the GPS master antenna, the attitude of the aircraft, and the laser range, two independent measurements of height above the runway may be calculated. The difference between these measurements includes IBLS vertical error, laser altimeter error, and error in the runway surface fit.

Figure 5.4 shows a plot of the error for several landing passes (touch-and-go's). Note that the noise level is lowest at the middle of the runway when the aircraft was rolling on the ground. The increased noise level on the left and right of the figure results from the time synchronization problem. Therefore, the four centimeter noise level between the vertical lines is thought to be more representative of the actual total error. The contribution of IBLS vertical error to the total error is believed to be only a few centimeters.

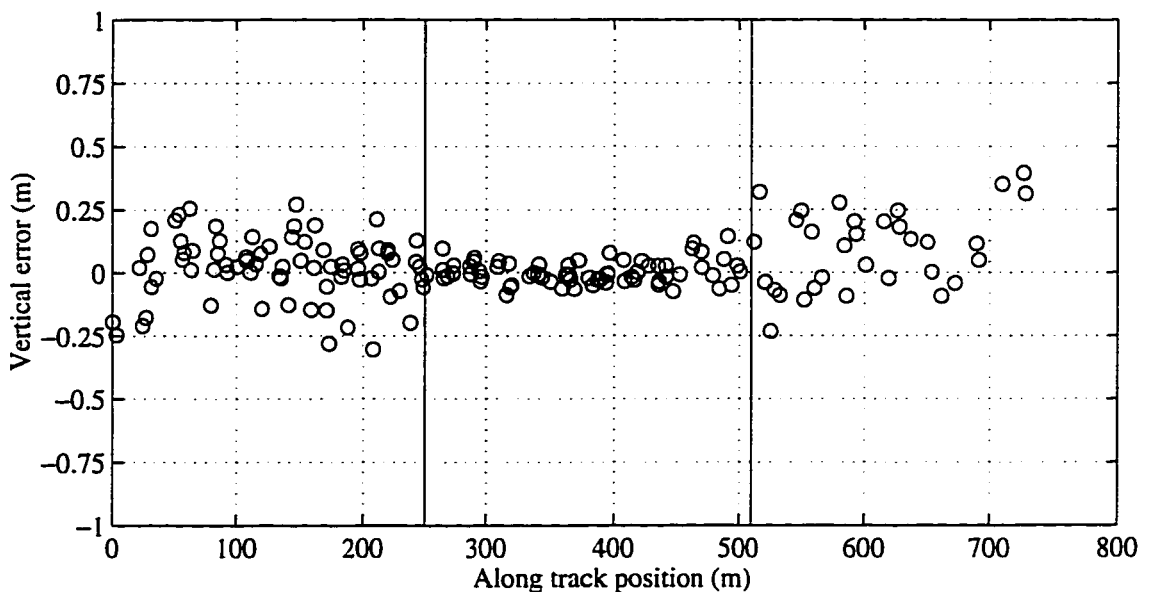


Figure 5.4: Laser Altimeter Comparison

5.1.3 Consistency Check

Another method to verify the performance of IBLS is available when the satellite signals are tracked continuously for the duration of a touch-and-go pattern. Once the integers are calculated, they should then be identical for each subsequent pass. Additionally, the fractional components of the floating integer estimates should be small. These checks have been performed during several flight tests, including those described in [Cobb, b]; IBLS consistently arrives at the same integer estimates from pass to pass. However, to more

faithfully simulate an actual instrument approach, the system is usually reset on the downwind leg. The reset involves discarding the accurate integer estimates calculated from the previous bubble pass, and estimating the new integers from code phase DGPS.

5.1.4 King Air Flight Test

To demonstrate closed-loop control using IBLS as the navigation sensor, a series of autocoupled approaches were performed using an FAA Beechcraft Super King Air 200 at the FAA Technical Center. The King Air was equipped with a Category I autopilot designed to perform autocoupled approaches down to the 200 foot Cat I decision height. The goal of the flight test was to attempt 50 coupled approaches and compile statistics on the performance of IBLS. The aircraft was equipped with a retro-reflector, and a laser tracker was available to verify the accuracy of the IBLS position solutions. Figure 5.5 shows a schematic of the aircraft hardware integration. Figure 5.6 shows the IBLS ground components.

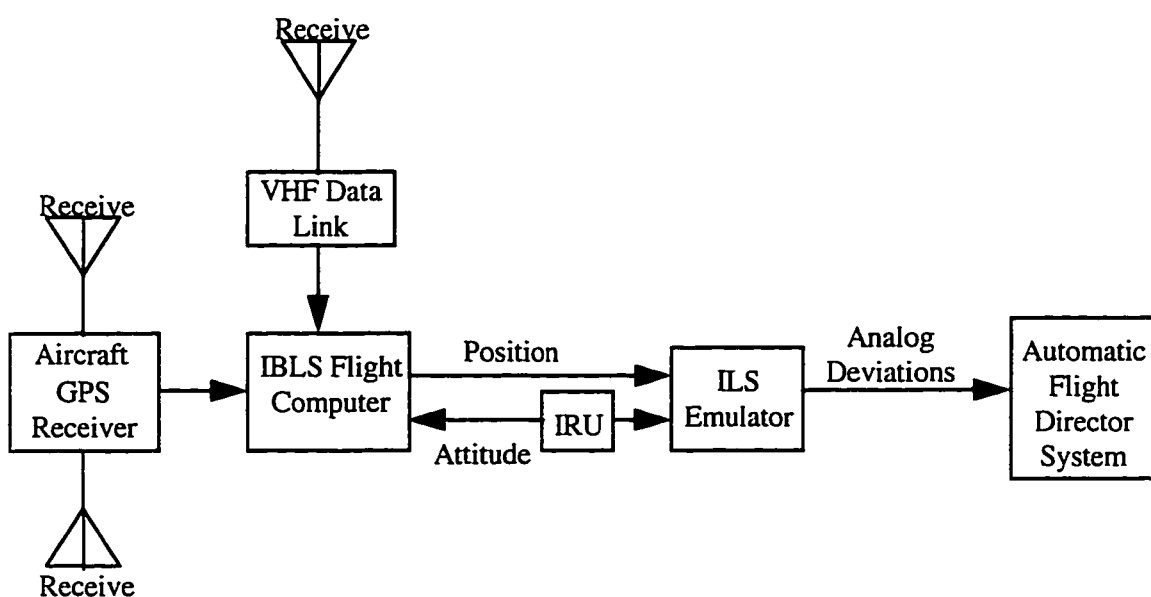


Figure 5.5: Aircraft Components for King Air Flight Test

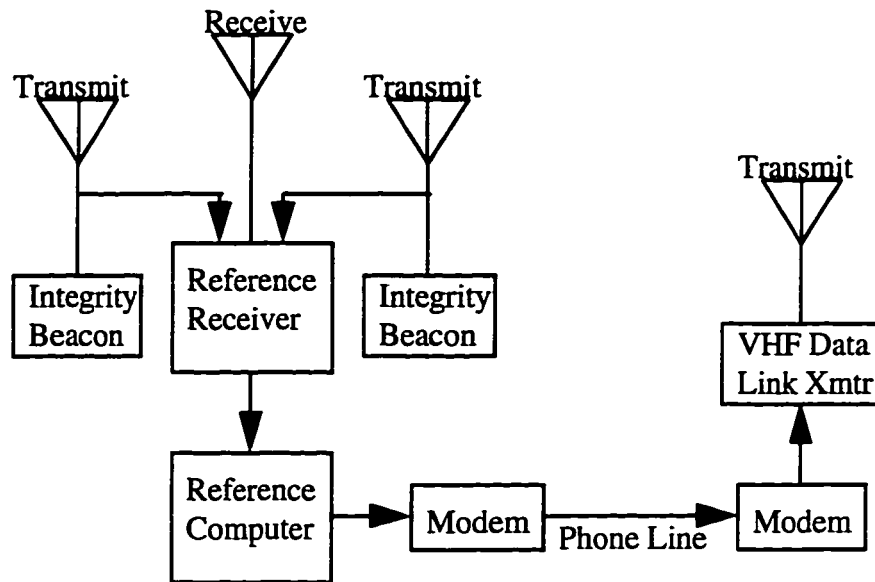


Figure 5.6: Ground Components for King Air Flight Test

After the system components were fully integrated and tested, the official test approaches began. The test pilot manually flew the aircraft through the base to final turn. The IBLS system was then reset to discard the accurate solution from the previous bubble pass and return to code DGPS accuracy. After this reset, the autopilot was engaged and used the architecture described in Chapter 4 to guide it down the approach path over the Integrity Beacons. Upon bubble exit at roughly 500 ft, the cycle ambiguities were resolved and the system achieved centimeter level positioning accuracy.

A total of 49 approaches were conducted before the aircraft had to be grounded for a repair (not related to our GPS landing system). On all 49 approaches, IBLS successfully resolved the cycle ambiguities and provided continuous guidance to the autopilot until the pilot resumed control, usually below 50 feet.

The IBLS position solutions were stored in real-time by the FAA Air Data Acquisition System and raw GPS and attitude measurements were stored on the Stanford flight computer. In post-processing the accuracy of the real time position solutions were checked against the laser tracker. Figure 5.7 and Table 5.1 show the results of this comparison. The difference between the GPS estimate and laser tracker estimate is again consistent with the laser tracker accuracy. The true NSE is believed to be typically on the order of several centimeters.

Table 5.1: Navigation Sensor Error² for King Air Flight Test

NSE (ft)	Vertical	Cross Track	Along Track
Mean (μ)	-0.78	0.18	0.38
Standard Dev. (σ)	0.64	0.53	0.76
$ \mu + 2\sigma$	2.06	1.24	1.90

Due to the nearly perfect NSE available from IBLS, the total system error (TSE) is essentially equal the flight technical error (FTE):

² Navigation Sensor Error listed here is laser position minus IBLS position.

$$TSE = NSE + FTE \equiv FTE \quad (5.1)$$

Figure 5.8 shows a plot of TSE in the cross track and vertical directions for all 49 approaches. For comparison, the proposed 95% inner tunnel boundaries [Kelly] are superimposed on the plot. The plots begin immediately after cycle ambiguity resolution; the relatively large vertical TSE at the left of the plot is due to the transition to high accuracy resulting from the integrity beacon overpass. Toward the right of the plot, the pilot had disengaged the autopilot and performed a go-around, causing the aircraft to depart from the 3 degree glideslope. Table 5.2 gives statistics for the vertical TSE. Note that the TSE using IBLS with a Cat I autopilot nearly met the specifications for a Cat III system during this flight test. The implication is that the high accuracy of IBLS could relax the constraints on FTE. This could allow less expensive autopilots to be used for Cat III approaches.

Table 5.2: Vertical TSE Statistics for King Air Flight Test

TSE (ft)	Overall	100 feet Alt	50 feet Alt
Mean (μ)	0.0	-0.8	-0.2
Standard Dev. (σ)	5.5	3.6	2.3
$ \mu + 2\sigma$	11.0	8.0	4.8


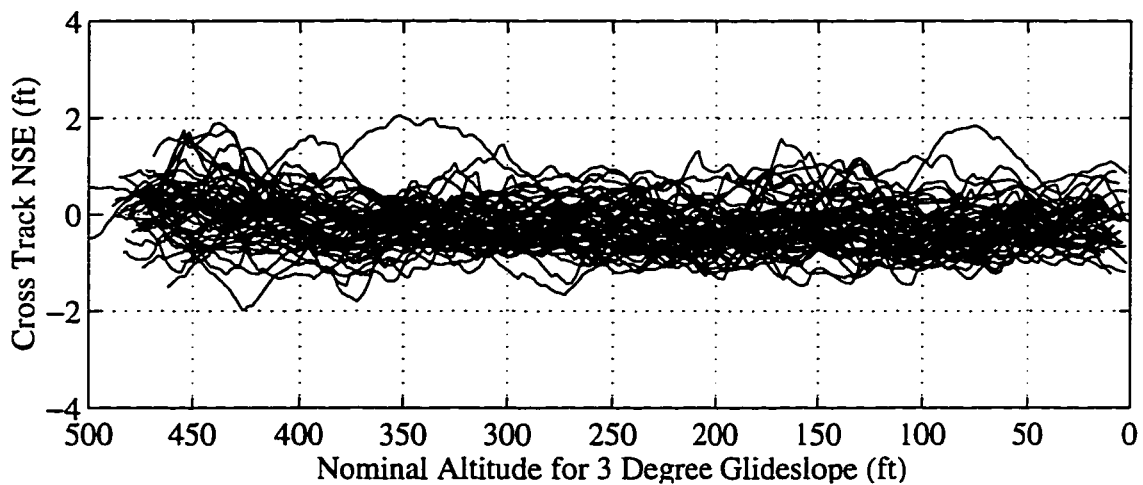
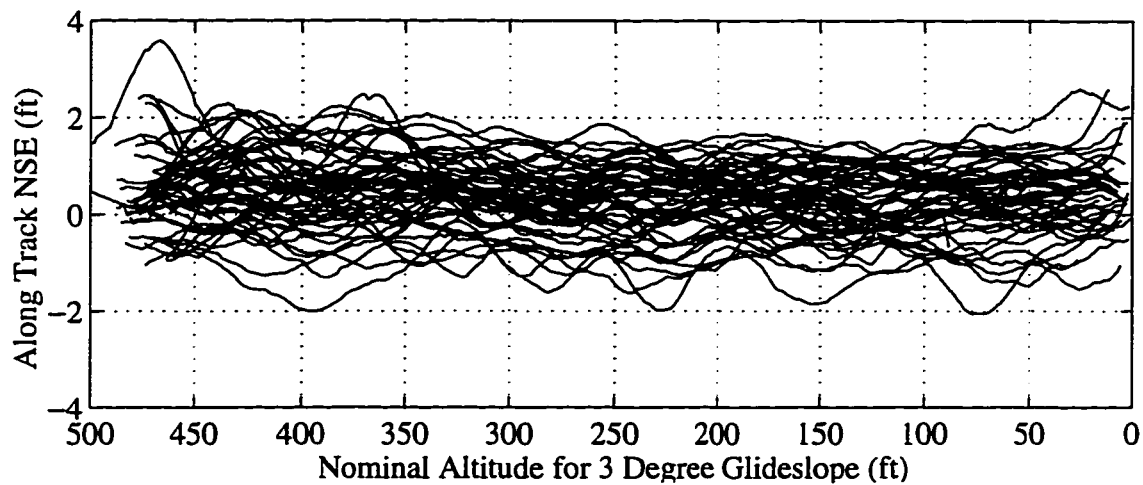


Figure 5.7: Navigation Sensor Error for King Air Flight Tests

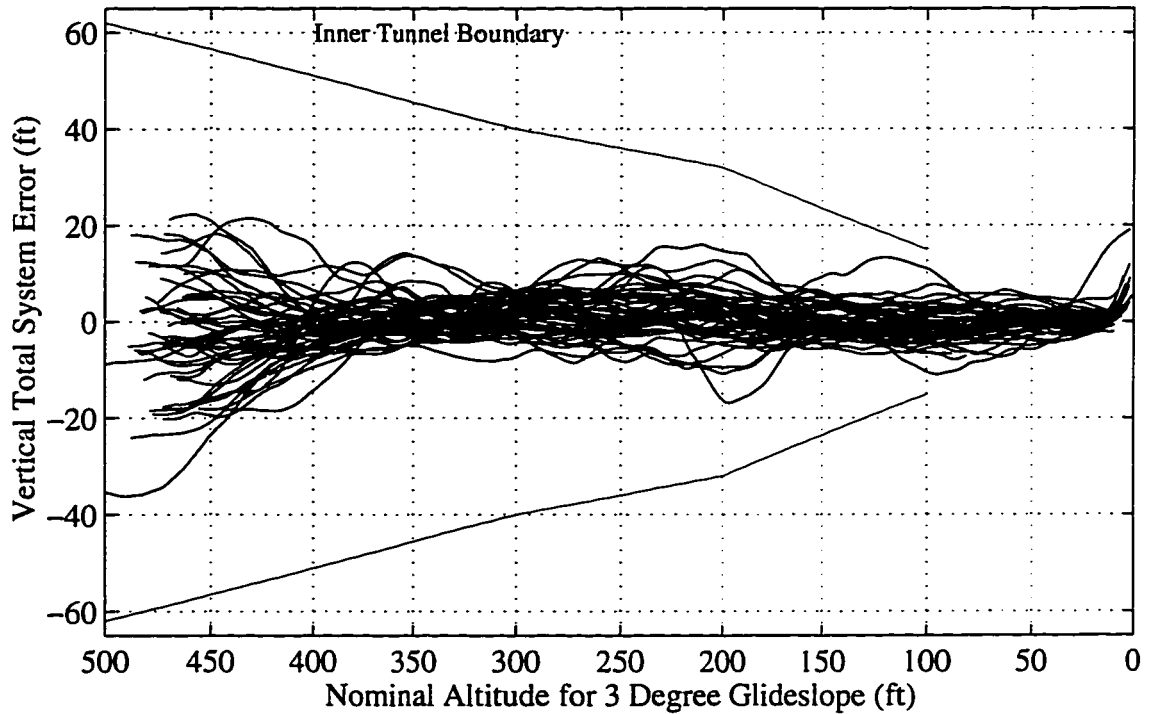
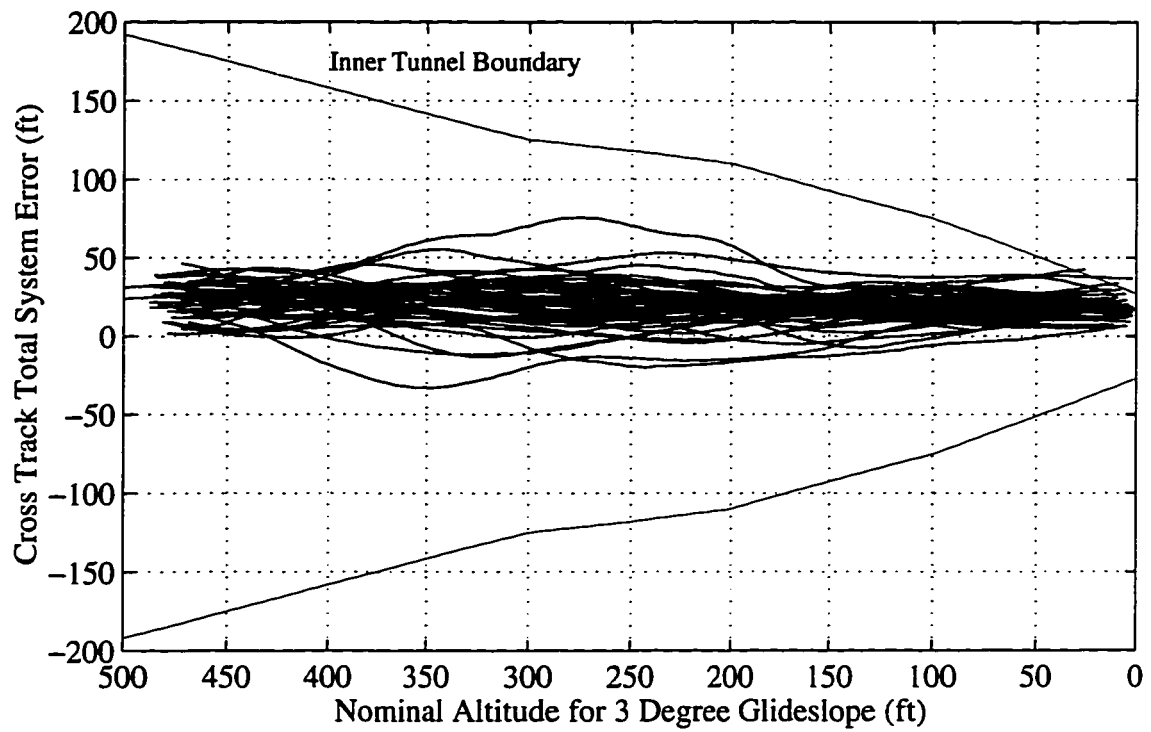


Figure 5.8: Total System Error for King Air Flight Tests

5.1.5 Autoland Flight Test

After the success of the King Air flight tests, the next goal was to demonstrate fully automated landings using IBLS. In October of 1994, IBLS was installed in a United Airlines Boeing 737-300 (Figure 5.9) equipped with a Category IIIA autopilot and a series of flight tests was performed at NASA Ames Crows Landing facility. The goal of this experiment was to perform 100 automatic landings under the guidance of GPS. IBLS data was provided to the autopilot in the form of an ILS look-alike signal. Therefore, the autopilot used a radar altimeter for vertical guidance during the last phase of landing. Although IBLS could have provided a radar altimeter look-alike signal as well, this was not done to avoid additional modifications to the aircraft (the immediate goal is to replace ILS, not the radar altimeter). A retro-reflector was installed in place of the aircraft's nose gear landing light, and a laser tracker was again used to verify the accuracy of IBLS.




Figure 5.9: Boeing 737 Autoland

The flight test procedure was essentially the same as for the King Air tests. The pilot manually intercepted the glideslope using the GPS-derived ILS display. The autopilot was engaged on long final approach and controlled the aircraft through touchdown. Category IIIA autopilots do not have roll-out guidance, so control was returned to the pilot at touchdown. Touch and goes were usually performed; full stops were performed every tenth autoland.

A total of 112 approaches were performed. Due to pilot error, one approach was aborted. A second approach was aborted when one of the four satellites being tracked was taken out of service [Cobb, c] (using the 6 channel receiver, only 4 channels are available to track satellites, while 2 track pseudolites. The next generation receiver has 9 channels and would seamlessly operate through single satellite outages in almost all cases). For this approach, IBLIS automatically disengaged the autopilot as it should have. The other 110 approaches resulted in successful automatic landings. Of those 110 approaches, 10 approaches were performed without resetting the integers calculated from the previous pass. This experiment was performed as a control to investigate the flight technical error of autopilot in the absence of the cycle ambiguity resolution transients.

Due to a satellite selection strategy that typically picked the highest four satellites, the PDOP during the autolands was quite high. Figure 5.10 plots the PDOP for all of the autolands. The mean PDOP was 6.7, with a maximum of 18.5. Due to the high PDOP, the position error before cycle ambiguity resolution was relatively high. The reason for the high

PDOPs is that the receiver used could only track four satellites. A receiver designed for IBLS use would have 9 or more channels thus allowing for tracking of two pseudolites and 7 or more satellites. In this case, PDOPs would be in the more typical range of two to four. Despite the high PDOPs however, the position estimate after the bubble pass agreed with the laser tracker to a level consistent with the laser tracker error. Figure 5.11 plots the difference between the laser tracker and IBLS; most of the difference is believed to be caused by laser error.




Figure 5.10: PDOP During Autolands




Figure 5.11: Vertical NSE for 737 Flight Test

The flight technical error for the 100 approaches when the integers were reset is plotted in Figure 5.12 and Figure 5.13, along with the inner tunnel dimensions. The large errors toward the left of the plots occurred before the autopilot was engaged. The jump near 500 feet altitude resulted from the transition to high accuracy. The transient from this jump decayed by 250 feet altitude. The upward bend at the right of the vertical plot results from the flare maneuver. The FTE statistics are given in Table 5.3 for altitudes of 100 and 50 feet.




Figure 5.12: Cross Track Flight Technical Error for 737




Figure 5.13: Vertical Flight Technical Error for 737

Table 5.3: Flight Technical Error for 737 Flight Test

FTE (ft)	Vertical (100 ft)	Cross Track (100 ft)	Vertical (50 ft)	Cross Track (50 ft)
Mean (μ)	0.3	0.7	0.3	0.3
Standard Dev. (σ)	3.6	7.2	3.3	6.9
$ \mu + 2\sigma$	7.5	15.1	6.9	14.1

It should be noted that the jumps in the FTE plots would be greatly reduced if the PDOP were lower. Figure 5.14 plots the vertical FTE for the passes where the PDOP was less than 5.5.




Figure 5.14: Vertical FTE for PDOP < 5.5

Some approaches were performed without resetting the integers calculated on the previous pass. The guidance was therefore accurate for the entire approach. Figure 5.15 plots the FTE for these passes. Using receivers with better code phase performance and more channels to improve geometry, the nominal performance should approach that shown in Figure 5.15. Even without these improvements, IBLS met the 95% tunnel specification.





Figure 5.15: 737 FTE (Approaches without IBLS resets)

5.1.6 Dynamic Tracking Experiment

Receiver tracking loop phase lags and processing delays could potentially cause errors unobservable without synchronizing the truth source with the output time of the position solution. For example, the laser tracker comparison assumed that the IBLS position solution occurred at the GPS timetag of the receiver sample. However, if a large delay existed between the receiver phase sample and the position solution time, the real-time accuracy of that solution would be worse than the laser tracker comparison would indicate. To quantify this effect for the IBLS system, video synchronization techniques were used in a dynamic environment. A star representing the position of the roving antenna was plotted on the computer monitor in a position corresponding to the real-time IBLS position estimate. The antenna was waved back and forth in front of the monitor at about 2 Hz. In the absence of GPS positioning errors, the star's position would track the antenna's position. Since the position of the star was determined from the real-time position solution, any delay in that solution would cause the star to lag the antenna. To observe such effects, a video camera recorded both the position of the antenna and the star. These positions were compared on a frame by frame basis to determine the tracking performance of the IBLS output. The results of this dynamic tracking experiment were consistent with the 10 Hz, 2nd order tracking loop of the receiver and one sixtieth of a second processing lag. For the dynamics experienced by commercial aircraft, this performance is more than sufficient to provide centimeter level accuracy (the position solution can be propagated ahead assuming constant velocity to account for short processing delays).

5.1.7 Integer Residual

Although no truth source exists to verify accuracy to the centimeter level in flight, the satellite phases can be compared against each other to verify that the measurements agree with each other to the centimeter level. Combined with the static accuracy check from Section 5.1.1 and the dynamic accuracy check from Section 5.1.6, this comparison provides compelling evidence for airborne accuracy at the centimeter level. To perform this comparison, the carrier phase residual vector is calculated:

$$r_{carrier} \equiv L(\phi - \hat{N}) \quad (5.2)$$

Since the cycle ambiguities are successfully resolved as shown in Section 5.1.1, the residual vector is related to the measurement error:

$$r_{carrier} = L\delta\phi \quad (5.3)$$

The magnitude of the carrier phase residual can be normalized to produce a scalar residual (η) with a mean square error independent of the number of satellites by dividing by the square root of the number of redundant measurements:

$$\eta \equiv \frac{|r_{carrier}|}{\sqrt{n-4}} \quad (5.4)$$

If some unmodeled error source exists only in an airborne environment, η should increase in that environment. Figure 5.16 and Figure 5.17 plot the normalized residual versus acceleration and altitude respectively for flight tests with the Piper Dakota at Palo Alto Airport. The lack of correlation in the figures is evidence that the error sources affecting the carrier phase are well understood and the accuracy in the airborne environment is

comparable to the accuracy on the ground – on the centimeter level. The large number of points at 450 meters in Figure 5.17 results from this test being conducted above pattern altitude to avoid disturbing the flow of traffic.




Figure 5.16: Normalized Residual versus Acceleration




Figure 5.17: Normalized Residual versus Altitude

5.2 Wide Area Augmentation System Performance

The Stanford University Wide Area DGPS network has provided a test bed for the development and evaluation of Wide Area Augmentation System (WAAS) algorithms. Until the research described in this section was conducted, the accuracy performance of these algorithms was assessed only for static users and users on the ground. The only truth models available relied on the user being at a known location or on a surveyed runway. To remedy this situation, the WAAS system was integrated with the Integrity Beacon Landing System. The accurate trajectories provided by IBLS are used to assess WAAS performance in an airborne environment.

The integration was achieved by porting both the IBLS user software and the WAAS user software to a real-time multi-processing operating system. Both systems now run simultaneously as separate processes on a single computer. The processes can communicate with each other for real-time comparison. They also store data to allow more detailed evaluation in post-processing. The integration of WAAS with IBLS provides more than just positioning truth for WAAS tests. The transition from en-route navigation to precision approach and landing can now be explored.

5.2.1 Background

Stanford University has developed an experimental WAAS network. The network includes three reference stations at Arcata, CA, Elko, NV, and San Diego, CA. These reference stations send GPS observables over phone lines to a master station at Stanford. The master

station generates a vector differential correction from these observables, puts the corrections into the RTCA WAAS message format, and then broadcasts the messages over a UHF data link [Walter].

As described in Chapter 1, Stanford University has also been developing IBLS as a means of augmenting GPS to provide the performance required to achieve Category III specifications. The high accuracy position provided by IBLS can be used to assess the airborne position accuracy of WAAS. However, the IBLS and WAAS systems were developed independently and could not both fit in the test aircraft at once. Previously, the dynamic accuracy of WAAS was assessed by plotting the height of the aircraft above the runway as in Figure 5.18. This figure shows the height of the aircraft during a series of touch-and-goes as determined by WAAS. When the aircraft was on the runway, this height was known and vertical errors could be observed. However, this method could not monitor the performance of WAAS for an airborne user. To solve this problem, WAAS and IBLS were integrated.




Figure 5.18: Original WAAS Accuracy Check

5.2.2 Integration

Before this research was started, the airborne software for WAAS and IBLS were based in DOS. Although both programs ran on a PC, they could not run simultaneously. To achieve the desired comparison between IBLS and WAAS position solutions, this limitation had to be overcome.

Integrating the two programs in the DOS environment would be difficult and cumbersome. Since DOS does not allow multiple programs to run at once, the WAAS and IBLS software would have to be combined into one large program. This program would almost certainly run into problems with DOS memory limitations. Another alternative would be to run the programs on two separate computers. This solution requires few software modifications but

is inefficient, does not allow for further expansion, and requires two computers to fly in the aircraft. A third alternative was sought. A more promising idea was to abandon the DOS operating system for an operating system more consistent with the needs of a GPS research lab.

A real-time, multi-processing, Unix operating system was chosen to perform the software integration. This operating system allows multiple processes or threads to share the CPU. The real-time kernel decides when each process runs using a priority-based scheduling system. Inter-process communication is available in the form of shared memory and first-in-first-out “pipes.” Semaphores (mechanisms to ensure only one process uses a shared resource at once) are available to coordinate multiple processes using the same resources. The WAAS and IBLIS programs were ported separately to this real-time operating system. Once each program worked individually, they could be run together as separate processes on the same computer.

After the WAAS and IBLIS programs were made to run in parallel, they were modified slightly to communicate with each other. Figure 5.19 shows the information flow for the integrated system. The data from the IBLIS receiver is received by the IBLIS process along with local area differential corrections. This data is used to calculate a position fix that is accurate to centimeter level (assuming the carrier phase cycle ambiguities have been resolved). Similarly, the WAAS receiver data and wide-area differential corrections are received by the WAAS process and used to calculate a wide-area DGPS position fix. These

positions are stored in memory shared by the both processes and can be compared in real time. The IBLs and WAAS position solutions are displayed side-by-side on the user display. They are also stored on disk for post-flight comparison. Raw GPS data and differential corrections are also stored to disk.

To simplify the integration, the airborne system uses two GPS receivers, one for IBLs and one for WAAS. The WAAS software was developed for a NovAtel GPSCard receiver, while the IBLs software was developed for a Trimble TANS architecture. Each receiver has advantages over the other in the application for which it is used. For example, the TANS firmware has been modified to track pseudolites. For this reason, it was decided not to attempt to use a single receiver for both WAAS and IBLs for this preliminary research. However, a single data link is used for both wide area and local area corrections.

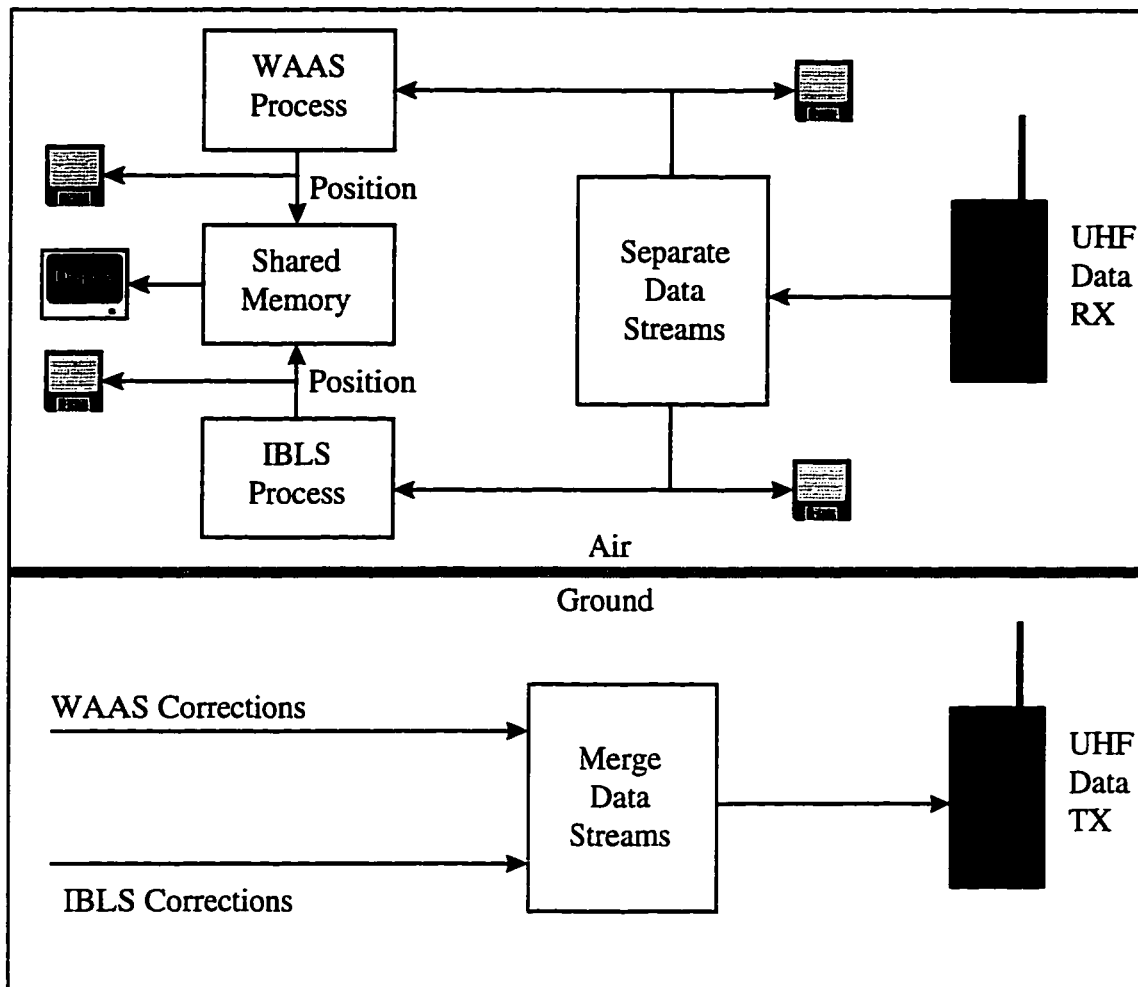


Figure 5.19: Information Flow for Integrated System

In addition to the WAAS and IBLS processes, a GPS attitude determination process may be run on the same computer. Attitude software developed by Cohen [Cohen, a] was ported to run under the new operating system. Real-time WAAS position, IBLS position, and attitude can be provided at 10 Hz or more.

5.2.3 Flight Test Procedure

Due to fundamental differences between the systems, the WAAS ground hardware is completely separate from the IBLS ground hardware. The wide area corrections are sent via

modem to the IBLs reference station where they are embedded in IBLs-like packets and transmitted to the aircraft along with IBLs corrections. The 250 bps WAAS message can easily fit into the IBLs data link. The procedure for conducting a combined WAAS/IBLS flight test is nearly identical to a WAAS flight test for the ground crew in the WAAS lab. The only difference is that the wide area corrections are sent to the IBLs lab instead of directly to the aircraft. The procedure for the ground crew in the IBLs lab is identical to an IBLs flight test³. After the ground equipment is set up, both the WAAS and IBLs reference stations are completely autonomous.

On-board the aircraft, the flight test procedure is also quite straightforward. A pallet carrying the flight test computer and GPS receivers is installed. The only addition to the IBLs hardware is the NovAtel receiver for WAAS which mounts directly into an ISA bus slot inside the computer. After providing the pallet with power and connecting the GPS and data link antennas, the airborne installation is complete. The IBLs and WAAS processes are started and begin to perform DGPS position fixes as soon as differential corrections are received. However, the IBLs positioning accuracy does not reach centimeter-level until the integer cycle ambiguities associated with kinematic GPS are resolved.

³ Recall that the IBLs airborne software has provisions to initialize cycle ambiguities without the use of pseudolites (Section 4.2.6). For the tests described in this section, pseudolites were not used.

For the purpose of this research, the integers may be resolved using one of several methods. If the initial position of the aircraft at the tie-down is known, this position may be used to solve for the integers (Section 4.2.6.3). Alternatively, if the aircraft remains stationary for several minutes, a static survey can be performed to calculate the integers (Section 4.2.6.2). If the integers are not resolved on the ground, they may be calculated in flight when the aircraft passes through a pair of Integrity Beacon bubbles (Section 4.2.3.3). Regardless of how the integers are resolved, once they are correctly determined, IBLS provides centimeter-level position truth thereafter.

5.2.4 Flight Tests

Several flight tests of the integrated IBLS/WAAS system were performed. The results of one such test are presented here. During that test, the integer cycle ambiguities for the carrier-based IBLS system were solved for at the tie-down using the known location of the aircraft. The VHF data link experienced dropouts resulting in several missed WAAS corrections in a row at times.

5.2.4.1 Results

Figure 5.20 shows the real-time position error during flight. Three sets of outliers are noticeable in the data. Two of the outliers resulted from the poor data link performance. In an operational WAAS, the data link would be required to provide significantly better performance. The third outlier resulted from a software bug that has since been corrected. Including the bad data, the overall real-time WAAS position error was 2.35 meters 3D RMS, which more than satisfies all WAAS accuracy requirements.




Figure 5.20: Real-time WAAS Position Error

The gray areas are projections of the errors into the indicated 2 dimensional planes.

CHAPTER 6

Dual Antenna Positioning

The nominal accuracy provided by GPS was shown in Chapter 5 to exceed the proposed requirements for precision landing. However, the availability of such systems is limited by the number of satellites in view. To ensure continuity, an approach should only be initiated if the satellites in view guarantee accuracy and integrity for the duration of the approach. The current GPS constellation does not appear to provide the required availability. Consequently, research was performed to explore methods which can improve the availability performance provided by GPS. This chapter explores the use of two antennas for GPS positioning. The following chapter looks at pseudolite augmentation.

Satellite availability studies commonly assume a fixed elevation mask angle. However, as an aircraft banks, relatively high satellites can be masked from the field of view of the GPS antenna. Not only are fewer satellites visible, but those that are in view are clustered in one section of the sky. This geometry typically leads to high PDOP. During steep turns, this effect can limit the availability of GPS and WAAS satellites, especially at high latitudes. However, if a second GPS antenna is installed on the aircraft, satellites masked from the main antenna may still be used. This section presents experimental results of kinematic

GPS positioning using two antennas. Experiments were first performed using a model aircraft on the ground. Flight tests were then performed on the Piper Dakota (Figure 5.2). The results of these tests show that multiple GPS antennas can be used effectively to improve the availability of GPS positioning.

6.1 Background

[Kinal] started to address the issue of satellite masking by the aircraft fuselage for the Wide Area Augmentation System (WAAS). He showed that satellites can be tracked below the local horizon of the GPS antenna, allowing the aircraft to maintain satellite tracking continuity in most maneuvers associated with final approach. This research focused on non-precision and Category I precision approaches. For the extreme Required Navigation Performance (RNP) of a Category III approach, it is useful to explore the feasibility of using two antennas for GPS positioning.

Although two antennas offer an obvious increase in satellite visibility, the optimal placement of the antennas is an open question. Ideally, there should be some overlap in the fields of view of the two antennas. That way, the second antenna can acquire the signal *before* the first antenna loses it. If this hand-off procedure from one antenna to another is robust, continuous tracking of most satellites can be guaranteed. To maximize the total coverage of the two antennas, this overlap region should be minimized. The results of [Kinal], which show that the “field of view” of an antenna may comfortably extend 10 degrees below the antenna’s horizon, imply that the optimal antenna placement is for the antennas to be pointed in opposite directions. Due to a combination of antenna patterns and

receiver signal detection thresholds, the overlap region for diametrically opposed antennas could not be consistently demonstrated here. However, positioning with diametrically opposed antennas was demonstrated as was satellite hand-offs between antennas pointed 120 degrees apart. With the appropriate hardware and software modifications, robust hand-offs should be possible with diametrically opposed antennas.

6.2 Experimental Setup

6.2.1 Receivers

A nine channel Trimble receiver was used in this research. To accept inputs from two antennas, it was modified to include a second RF section. Both RF sections run off a common local oscillator. Each channel has a switch to allow it to accept inputs from either RF section; in real time, the flight software can select antennas on a channel-by-channel basis.

To maximize continuity, satellites must be tracked by one antenna before they are lost by the other. Therefore, the same satellite must be tracked on two channels at the same time. The receiver firmware was modified to accommodate this requirement. To track a satellite on two channels, a PRN code unused by the current constellation is replaced by that satellite's PRN code, and the "aliased" satellite is tracked on one of the channels.

A similar nine channel receiver was used as the differential reference station receiver. However, the reference receiver did not require the above modifications.

The phase measurements for GPS attitude were provided by a Trimble TANS Quadrex, a six-channel multiplexing receiver. One receiver could be used for both positioning and attitude; two receivers were used here to avoid integrating all the desired functionality into one receiver. The attitude solutions were performed in the flight computer by a separate process.



Figure 6.1: Model Aircraft

6.2.2 Aircraft

6.2.2.1 Model

Figure 6.1 shows a one meter wingspan model aircraft used in the first experiments. Antennas on the nose, wings, and tail are used for attitude determination. Two additional antennas mounted behind the nose are used for positioning. Several configurations were tried for the positioning antennas; in the figure, they are placed on the starboard and port,

with 120 degrees between boresights. The model is mounted on a tripod which allows any heading and ± 90 degrees of roll.

6.2.2.2 Piper Dakota

IBLS (Section 1.6) uses a bottom GPS antenna to acquire pseudolite signals. Attitude is provided to account for the lever arm from the top antenna to the bottom antenna. With a bottom antenna and attitude already available, the IBLS test aircraft (Figure 5.2) required few changes to demonstrate GPS positioning using multiple antennas. One of the attitude antennas and the belly antenna were used for the two antenna positioning experiments.

6.3 Software

The flight software executes on a Pentium PC running the Lynx Real-Time operating system. Two processes are run simultaneously, an attitude process and a positioning process. The attitude process sends attitude solutions to the positioning process. To calculate position fixes using two antennas, attitude is required to account for the lever arm between those antennas.

As part of the positioning process, a simple algorithm is used to select the satellites and antennas for each channel. The highest seven satellites are selected and tracked on the antenna whose boresight is closest to the line-of-sight to each satellite. This selection is made in real-time at 4 Hz using GPS attitude to determine which antenna is closest. The two satellites closest to the borderline between antennas are aliased to unused PRN's and tracked a second time. These channels are used to provide tracking continuity in the event

that a borderline satellite switches antennas. The software developed by Cohen [Cohen, a] was used to solve for attitude. The attitude accuracy is sufficient to account for the lever arm correction with errors on the millimeter level.

The positioning software for this two antenna research evolved from the software described in detail in Chapter 4. Only a few changes were required to modify this software to accommodate positioning with two antennas:

- The antenna selection algorithm described above.
- A correction for the lever arm between the two antennas.
- A correction for the orientation of the antennas. This correction results from the circular polarization of the incoming signal and is described in detail in Section 6.4. It is required for this application because the two antennas may point in significantly different directions than the reference station antenna.

To apply the lever arm correction, the phases measured at the second antenna are simply projected to the first antenna:

$$\phi_{projected} = \phi + \hat{e} \cdot R_{AC}^{ENU} \bar{a} \quad (6.1)$$

where:

\hat{e} is the unit line-of-sight vector to the satellite expressed in East, North, Up.

R_{AC}^{ENU} is the rotation matrix from the aircraft coordinate frame to East, North, Up.

\vec{a} is the vector from antenna 1 to antenna 2 expressed in the aircraft coordinate frame.

After applying these corrections, all phase measurements appear to have come from one antenna; the remainder of the software is identical to the single antenna software.

The aliased satellites are treated as if they were separate satellites with identical line-of-sight vectors to the true satellite. The aliases do not switch antennas when the true satellites switch, so the satellites can be continuously tracked. Using the phase measurement of the aliased satellite, the estimator will implicitly hand off the integer estimate when the true satellite switches antennas.

6.4 Circular Polarization Correction

As specified in [SPS], the transmitted GPS signal is right-hand circularly polarized (RHCP). Therefore, GPS receive antennas are designed to receive RHCP signals. The measured carrier phase of a circularly polarized signal is a function not only of the distance between the transmit and receive phase centers, but also of the relative orientation of the antennas.

Traditionally, kinematic GPS applications do not correct for this effect. When the boresights of all of the receive antennas are parallel, the correction is common to all satellites. It therefore affects only the differential clock error or line bias, not the position or

attitude solution. However, if the angle between receive antenna boresights becomes large, a RHCP correction should be applied. One way to develop such a correction is to find a correction for each transmit and receive antenna pair. The receive antennas are assumed to be flat patch antennas; the results can be generalized for other types of antennas given their off-boresight phase characteristics. To first order, the phase characteristics of a patch antenna can be approximated by assuming the antenna is made up of two dipoles 90 degrees out of phase.

In general, the incoming signal will be elliptically polarized if the transmit antenna boresight does not point directly at the receive antenna. For terrestrial users receiving satellite signals, the ellipticity is guaranteed not to exceed 1.2 dB [SPS], so the incoming signal can be assumed to be circularly polarized. The RHCP correction is a function of the orientation of the receive antenna, the line-of-sight to the transmit antenna, and the ellipticity and orientation of the incoming signal. To develop a correction for a transmit/receive pair, two coordinate frames are defined. A right handed orthogonal coordinate frame is attached to the receive antenna with the z direction aligned with the boresight. The y direction can be arbitrarily chosen normal to z ; the x direction is then constrained. The second coordinate frame will be called the transmit frame. The transmit frame is defined such that the z axis points *opposite* the line-of-sight to the transmit antenna and the y axis points in the major axis direction of the incoming elliptically polarized signal. If the incoming signal is circularly polarized, this direction may be chosen arbitrarily. The

arbitrary terms in the absolute correction will cancel when single and double differences are performed.

The output of an idealized RHCP patch antenna can be simply modeled as the E-field component in the x direction plus the component in the y direction delayed by 90 degrees:

$$r(t) = x_r(t) + y_r\left(t - \frac{1}{4L_1}\right) \quad (6.2)$$

where:

$r(t)$ is the antenna output as a function of time (common-mode terms such as line biases are neglected in this development).

$x_r(t)$ is the E-field component in the receive antenna x direction.

$y_r(t)$ is the E-field in the receive antenna y direction.

L_1 is the carrier frequency.

Similarly, the incoming signal can be expressed in the transmit frame:

$$\begin{aligned} x_{\alpha}(t) &= \cos(2\pi L_1 t) \\ y_{\alpha}(t) &= e \sin(2\pi L_1 t) \\ \vec{E}(t) &= x_{\alpha}(t)\hat{i}_{\alpha} + y_{\alpha}(t)\hat{j}_{\alpha} \end{aligned} \quad (6.3)$$

where:

$\vec{E}(t)$ is the vector E-field at the receive antenna.

e is the ellipticity of the incoming signal.

\hat{i} is a unit vector in the x direction.

\hat{j} is a unit vector in the y direction.

The received signal is then:

$$\begin{aligned}
r(t) &= \hat{i}_r \cdot (\hat{i}_x x_x(t) + \hat{j}_x y_x(t)) \\
&+ \hat{j}_r \cdot \left(\hat{i}_x x_x \left(t - \frac{1}{4L_1} \right) + \hat{j}_x y_x \left(t - \frac{1}{4L_1} \right) \right) \\
&= \cos(2\pi L_1 t) (\hat{i}_r \cdot \hat{i}_x - e \hat{j}_r \cdot \hat{j}_x) + \sin(2\pi L_1 t) (e \hat{i}_r \cdot \hat{j}_x + \hat{i}_x \cdot \hat{j}_r) \\
&= \cos(2\pi L_1 t) (R_{11} - e R_{22}) + \sin(2\pi L_1 t) (e R_{12} + R_{21}) \\
&= \langle R_{11} - e R_{22}, e R_{12} + R_{21} \rangle \cos \left(2\pi L_1 t + \arctan \left(\frac{e R_{12} + R_{21}}{R_{11} - e R_{22}} \right) \right)
\end{aligned} \tag{6.4}$$

where R_{ij} is the (i,j) element of the rotation matrix from the transmit coordinate system to the receive coordinate system.

The phase term, $\arctan \left(\frac{e R_{12} + R_{21}}{R_{11} - e R_{22}} \right)$, represents additional delay of the received signal due to orientation. Using the sign convention such that additional delay increases phase (moving away from the spacecraft increases phase), this correction should be subtracted from the phase measurement. Care should be taken to “unwrap” the arc-tangent function.

The corrected phase, expressed in cycles can be written:

$$\phi_{corrected} = \phi - \frac{1}{2\pi} \arctan \left(\frac{e R_{12} + R_{21}}{R_{11} - e R_{22}} \right) \tag{6.5}$$

6.4.1 Example of RHCP Correction

Figure 6.2 shows a simple geometry to illustrate the RHCP correction. In general, the correction for a single transmit/receive antenna pair will contain arbitrary terms that would

be common mode and would therefore either cancel in a single difference or only affect the receiver clock bias solution. However, through judicious selection of the antenna x and y coordinate directions, these common mode terms are zero. This was done only to simplify the example.

The figure shows an aircraft flying into the page in a 30 degree right bank. The carrier phase is measured for a low-elevation satellite directly ahead of the aircraft. The coordinate frames for the transmit and receive antennas are shown on the right of the figure, along with the direction of the incoming polarized signal (assumed circular). For this geometry, the rotation matrix from transmit to receive coordinates is given by:

$$R = \begin{bmatrix} \cos(30) & -\sin(30) & 0 \\ 0 & 0 & -1 \\ \sin(30) & \cos(30) & 0 \end{bmatrix} \quad (6.6)$$




Figure 6.2: Geometry for Example RHCP Correction

The RHCP correction is therefore:

$$\Delta\phi = -\frac{1}{2\pi} \arctan\left(\frac{R_{12} + R_{21}}{R_{11} - R_{22}}\right) = -\frac{1}{2\pi} \arctan\left(\frac{-\sin(30)}{\cos(30)}\right) = \frac{1}{12} \text{ cycle} \quad (6.7)$$

For the geometry of Figure 6.2, the magnitude of this correction is about 1.6 cm.

6.5 Experimental Results

6.5.1 Ground Tests

The model aircraft shown in Figure 6.1 was used to validate the software and demonstrate satellite hand-offs from one antenna to the other. Several antenna configurations were used, including diametrically opposed antennas. For the hardware used in these tests, the greatest continuity improvement was with the antenna configuration shown in the picture. One antenna pointed up to the left and one pointed up to the right; the angle between boresights was 120 degrees.

To demonstrate the effectiveness of dual antenna positioning, the model was placed in exaggerated attitudes. Roll reversals from 90 degrees to -90 degrees were performed; kinematic positioning was maintained throughout these extreme maneuvers. Figure 6.3 plots the roll for one such maneuver. The model started level with a southward heading and remained there for about 15 minutes. During this time, satellites were tracked on both the starboard and port antennas. The model was then rolled left to a roll of nearly -90 degrees, where it remained for several minutes.




Figure 6.3: Model Aircraft Roll Reversal

The satellite tracking status for this roll to the left is shown in Figure 6.4. These figures are traditional azimuth, elevation polar plots. The radial direction represents elevation (90 degrees at the center to 0 degrees at the outside); the azimuth varies in the clockwise direction as shown. The x's and o's denote satellite positions. An 'x' means the satellite is tracked on the port antenna while an 'o' means it is tracked on the starboard antenna. Both an 'o' and an 'x' means it is simultaneously tracked on both antennas. The solid arc represents the boundary between those satellites closer to one antenna and those closer to the other. For the antenna configuration of the model aircraft, this borderline separates satellites with line-of-sight vectors having components in the starboard direction from those having components in the port direction.

In Figure 6.4a the aircraft is in its initial level configuration. Two satellites are tracked by the starboard antenna, three by the port antenna and two are tracked by both. The borderline between antennas runs north to south. In Figure 6.4b, the aircraft has started its bank to its left (east). Four satellites are now tracked on the starboard antenna, two on the port antenna

and one on both antennas. Note that one satellite has been acquired by the port antenna before being lost by the starboard antenna. This process continues until the roll reaches about -90 degrees in Figure 6.4f. At this time, all seven satellites in view are tracked only by the starboard antenna. During this bank to the left, all of those satellites were continuously tracked by at least one of the antennas. The cycles ambiguity estimates were therefore maintained throughout the bank.

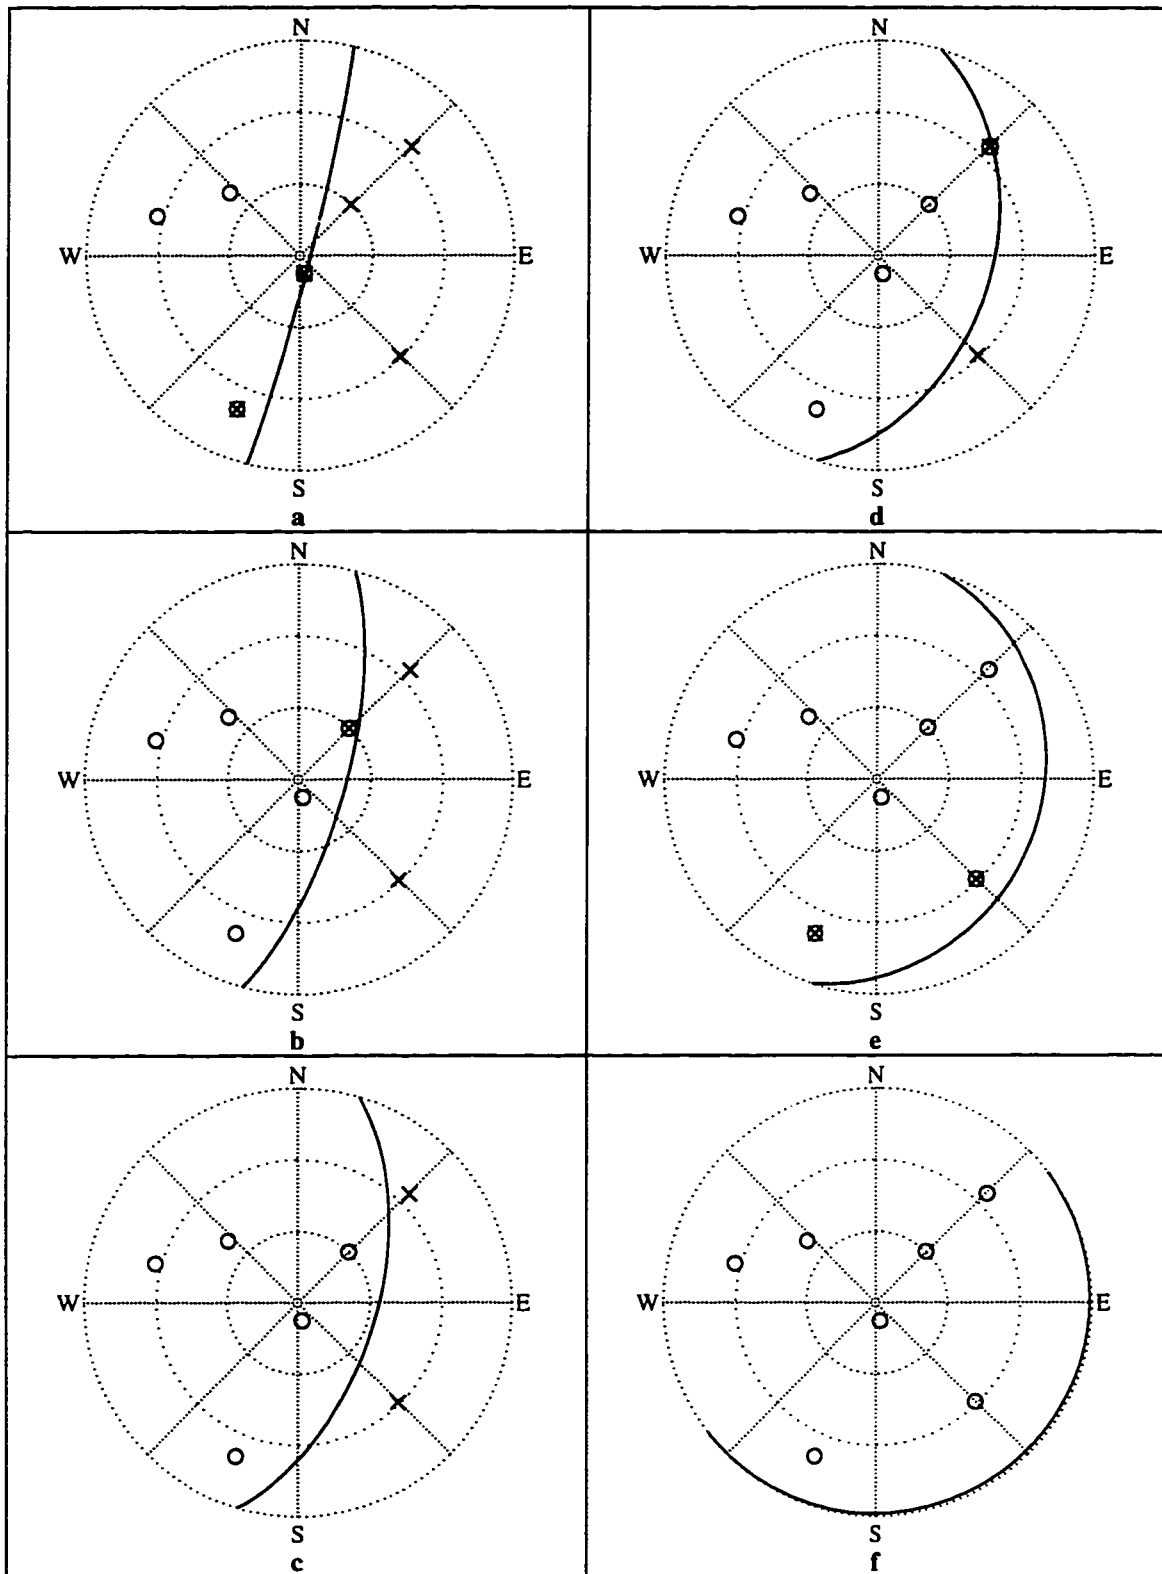


Figure 6.4: Satellite Tracking Status During Roll Reversal
x: port antenna, o: starboard antenna

The model was then returned to level “flight” and then rolled 90 degrees to the right. For the duration of this test, kinematic positioning was maintained. Interestingly, at a roll of 90 degrees, the GPS attitude solution was lost due to poor satellite visibility. Typically, GPS attitude is more robust than GPS positioning, because fewer satellites are required. However, due to the dual positioning antennas, the position solution was maintained and attitude was lost.

As the aircraft rolled about the tripod pivot, the positioning antennas traced out a nearly semi-circular arc. The kinematic position solutions for the starboard antenna are plotted in Figure 6.5. This north-looking view clearly shows the arc traced out by the antenna. A nearly independent measure of position can be derived from the attitude if the tripod pivot position and the lever arm from the pivot to the antenna are known. These positions were compared, and agreed to within 6 cm. The 1σ position error was about 2 cm.




Figure 6.5: Starboard Antenna Position

6.5.2 Flight Tests

After validating the dual antenna positioning system on the model aircraft, the system was flight tested on the Piper Dakota. The two antennas used for positioning were on the top and bottom of the Piper instead of the starboard and port.

The first flight test was conducted in the traffic pattern at Palo Alto Airport. During the crosswind and base turns, satellites were acquired on the bottom antenna. However, the dual antenna positioning solution was inconsistent. Occasionally, the Receiver Autonomous Integrity Monitoring (RAIM) software detected a high residual soon after incorporating the bottom antenna into the solution. Phase errors greater than 5 meters were discovered. Errors of that size are extremely rare for kinematic GPS; the carrier phase error is typically less than a centimeter. In post processing, it was discovered that the error was due to a reflected signal dominating the signal received by the bottom antenna. Palo Alto Airport is situated next to the San Francisco Bay; the downwind leg of the standard closed pattern is over water. Water is a good reflector, and satellite signals can reflect off the bay and into the bottom antenna. This reflected signal can enter the antenna at small angles from the boresight, while the direct signal may enter at a less favorable angle or be blocked by part of the aircraft. The phase error due to this extreme multipath scales with the altitude above the reflecting surface. The error profile can therefore be predicted from the altitude profile by computing the additional distance traveled by a reflected signal versus the direct path. Figure 6.6 shows a comparison between the predicted multipath profile and the

observed phase error profile. The correlation of the two profiles strongly suggests that the bottom antenna was receiving a reflection from the bay.

This phenomenon does not preclude using a bottom antenna for acquiring satellite signals. Care must be taken to ensure that the direct signal is dominant before using the satellite. These precautions can be taken both at the signal processing level and in the positioning software (for example, use only satellites within 100 degrees of the antenna's boresight and perform RAIM). To reduce the impact of multipath on the remaining dual antenna flight tests, the trials were conducted over land. Similar multipath problems were not encountered.

For the next flight test, a series of steep turns were performed over the Stanford campus. The cycle ambiguities were initialized using the IBLS software. One Autonomous Integrity Beacon was placed on campus for this purpose. After solving for the cycle ambiguities, the aircraft position was known to the centimeter level. The steep turns then commenced, and satellites were acquired through the bottom antenna on several occasions. During one turn, a satellite was continuously tracked for 8 seconds through the bottom antenna. Figure 6.7 shows the geometry at one time during that period. Note that the borderline between antennas now separates satellites with line-of-sight components toward the top of the plane from those with line-of-sight components toward the bottom of the plane; an 'o' means the satellite is tracked on the top antenna and an 'x' means it is tracked on the bottom antenna. In the figure, the aircraft heading is northwest and its roll is -32 degrees. Six satellites were

tracked on the top antenna and one was tracked on the bottom. To quantify the dual antenna positioning performance, this 8 second period was post-processed and single antenna position solutions were compared with dual antenna solutions. The traditional single antenna software used the six satellites tracked on the top antenna to generate a centimeter-level trajectory. The dual antenna software was then used to calculate the same trajectory using carrier phases measured from both antennas. To simulate poor satellite geometry, three of the top satellites were removed from the dual antenna solution. Therefore, the dual antenna trajectory used three top satellites and the bottom satellite. Before the bottom satellite was lost, the three top satellites were reintroduced. Figure 6.8 shows the difference between the single and dual antenna trajectories. From this comparison, it can be seen that the bottom antenna can be used to provide centimeter-level accuracy in the event that the top antenna satellite geometry is poor.




Figure 6.6: Evidence of Extreme Multipath




Figure 6.7: Az/El Plot of Satellite Tracking Status
x: bottom antenna, o: top antenna




Figure 6.8: Comparison of Single and Dual Antenna Positioning

CHAPTER 7

Airport Pseudolite Augmentation

The stringent availability and continuity requirements for precision approach and landing are difficult to achieve using *unaugmented* local area DGPS. Even more severe is the requirement on integrity. However, to check the integrity of a GPS position solution with Receiver Autonomous Integrity Monitoring (RAIM), more than four ranging sources are needed. Therefore, to meet all of the Required Navigation Performance (RNP) parameters, it may be necessary to augment the GPS constellation with additional ranging sources. Such an augmentation may be achieved by placing a pseudolite at the airport.

To show the navigation performance improvements offered by pseudolite augmentation, flight tests were performed in a Piper Dakota. A pseudolite was placed at the airport and was pulsed to eliminate the near-far problem (Section 7.1.1). The pseudolite signal was used as an additional ranging source to demonstrate the following tasks:

- Centimeter-level positioning accuracy with only three satellites.
- RAIM-based fault detection with only four satellites.

- Single-channel fault isolation with only five satellites.

These results show that a pulsed pseudolite can provide an additional kinematic ranging source to an aircraft on final approach. This low-elevation precision ranging source has the potential to greatly improve the navigation performance offered by a local area DGPS system. This section presents the results of preliminary flight tests incorporating pulsed pseudolite augmentation.

7.1 Background

Although the RNP parameters introduced in Section 1.2 are interrelated [Pervan, a], the need for additional ranging sources can be seen by considering continuity and availability. For a Category III system, the continuity objective proposed in [ORD] is one continuity interruption in ten million 80 second approaches (1×10^{-7}). The proposed requirement for availability is greater than 99.9% [ORD].

Consider a situation where only four satellites are available at the onset of an approach. If any one of those four satellites were to fail or drop below the aircraft's antenna horizon due to a large roll angle during the approach, the continuity of GPS positioning would be lost (four ranging sources are required to provide 3D positioning with GPS). To meet the above continuity requirement, the probability that any of the four satellite signals fail during an 80 second exposure time must be less than 10^{-7} . However, Cobb [Cobb, c] showed that one satellite failure mode causes Block II satellite continuity interruptions about three times per day. Using this information, the probability that four satellites will be available for a given

80 second interval is on the order of only 99.9%. This is several orders of magnitude worse than the *continuity* requirement. Neglecting the Block II anomaly, the continuity would be only 99.999% based on satellite failure models developed by Phlong and Elrod [Phlong], which predict both short-term and permanent satellite failures. Therefore, at least five satellites are required to meet the proposed continuity requirement, even without considering signal losses due to large roll angles.

In addition to the continuity issue, the high integrity required of Category III precision landing presents another problem. The integrity requirement proposed in [ORD], 4×10^{-8} per approach, may necessitate the use of RAIM. If so, there exists a second reason to require more than four ranging sources at the onset of an approach.

From the above discussion, it is clear that five or more ranging sources should be available before an approach is initiated. Pervan [Pervan, a] tabulated the expected availability of GPS satellites for San Francisco International Airport (SFO). Satellite failure models [Phlong] were simulated to generate these results. A 7.5 degree elevation mask was assumed. The availability of five or more satellites was only 99.4%, not high enough to meet the proposed requirement of 99.9%. Table 7.1 summarizes these results and indicates that augmentation is necessary. Four satellites do not provide the required continuity, while five or more satellites do not provide the required availability.

Table 7.1: Satellite Availability and Continuity

	4 satellite	≥ 5 satellite	Requirement
Availability	~99.9%	~ 99.4%	> 99.9%
Continuity Risk	$10^{-3} \rightarrow 10^{-5}$	$< 10^{-8}$	$\leq 10^{-7}$

To meet all of the proposed RNP requirements, some sort of augmentation to the GPS constellation appears necessary. Several alternatives exist to provide this augmentation. One example is the FAA's Wide Area Augmentation System (WAAS) which will provide additional ranging sources from geo-synchronous satellites. However, it would be desirable for any Local Area Augmentation System (LAAS) to be independent of WAAS for several reasons:

1. LAAS should be available as a backup if the WAAS ground segment or geo-synchronous satellites fail.
2. At high latitudes and in mountainous areas, geo-synchronous satellite visibility may be poor.
3. A solution is sought which can gain international acceptance. Dependence on the FAA's WAAS network is therefore undesirable.
4. Budgetary and schedule independence of WAAS and LAAS is sought to maximize the probability expedient transition to GPS-based landing systems.

One solution to this problem is to place pseudolites on the airport property to provide additional ranging sources. These Airport Pseudolites (APLs) are the focus of this section.

7.1.1 Near-Far Problem

As an aircraft flies toward a pseudolite, the signal level it receives from that pseudolite can vary from the threshold of detectability (less than -130 dBm) to more than 100 dB higher. Unfortunately, the GPS C/A signal format was designed to prevent mutual interference only among signals at similar power levels. The worst-case cross-correlation margin between two C/A codes is only about 21.4 dB. Strong signals received from a nearby pseudolite can overwhelm the weak satellite signals, jamming the receiver. This “near-far problem” is the major constraint on pseudolite applications [Cobb, a]. Several methods have been proposed to alleviate the near/far problem. These include frequency offsets [Brown], pulsed transmissions [Stansell] (see Section 7.2.3), and combinations of the two [van Dierendonck]. For reasons discussed in [Cobb, a], pulsed transmissions were chosen for this research.

7.2 Experimental Hardware

7.2.1 Aircraft

The Piper Dakota (Figure 5.2) used to perform this research was equipped with five GPS patch antennas, four on top for GPS attitude [Cohen, a] and position, and a downward-looking belly antenna used with the IBLS. The four top antennas are mounted on the top of the vertical tail, above the cabin, and on each wing. The tail antenna was used to receive the APL signal, because it had direct line-of-sight to the APL for the entire approach and landing. Tests were also performed using the cabin and belly antennas. As expected, the tail antenna was better able to maintain lock on the APL signal than the others.

7.2.2 Autonomous Pseudolite

Several signal formats have been suggested for pseudolites [Ndili]; the format which minimizes the impact on civilian receiver design is simply a replica of the Standard Positioning Service (SPS) format [SPS]. The pseudolite signal transmits an L1 carrier modulated with an unused PRN code. Typically, only minor receiver software modifications are required to track such a pseudolite signal.

The concept of an *autonomous pseudolite* was developed by Cohen at Stanford [Cohen, b]. Originally called an *Omni-Marker*, an autonomous pseudolite transmits a signal which is phase locked with an incoming satellite signal. This satellite will be referred to as the *Master Satellite*. The Master Satellite signal is received, the PRN code is changed, and a synchronous signal is rebroadcast. The PRN code is changed to allow the aircraft receiver to distinguish between the Master Satellite signal and the “reflected” pseudolite signal. The receive and transmit antennas can be separate, or a single antenna may be used. [Cobb, a] describes the design and implementation of an autonomous pseudolite.

One advantage of synchronizing the pseudolite with the Master Satellite is that a differential carrier phase measurement is provided without requiring a differential reference station to track the pseudolite. The pseudolite can therefore be isolated from the reference station, hence the name *autonomous pseudolite*.

7.2.3 Pulsed Pseudolite

A pulsed pseudolite refers to a pseudolite whose output is turned rapidly on and off to form a pulsed transmission. The pulses are short, about 150 μsec , with a duty cycle of less than 10%. The peak power during the pulses is high enough for a receiver to maintain signal lock on the pseudolite despite receiving it only a small fraction of the time. During the time that the pseudolite is on, the satellite signals are briefly jammed. However, the receiver tracking loops can coast through the outages with only a small decrease in satellite signal strength. There is no more “near” radius as shown experimentally in Section 7.4.2. The far radius can be extended to tens of miles by simply amplifying the pulse power. Therefore, pulsed pseudolites solve the near-far problem.

7.2.4 Receiver

The same nine channel Trimble receiver used in Chapter 6 was used for this research. The firmware was changed at Stanford to provide carrier phase output and to allow the receiver to track unused PRN codes. In lab tests, the receiver was able to track a pulsed pseudolite. However, there were some unexplained dropouts where the receiver temporarily lost the pseudolite signal. The cause of the dropouts was still under investigation at the time of this writing. More recent tests with new receivers and pseudolites did not exhibit this behavior [Lawrence, c].

7.3 Background Analysis

The single difference carrier phase measurement for an autonomous pseudolite is defined as the absolute carrier phase of the direct signal from the master satellite (Φ_{direct}) subtracted

from the absolute carrier phase of the “reflected” signal ($\Phi_{reflected}$), both measured by the roving receiver:

$$\phi_{APL} = \Phi_{reflected} - \Phi_{direct} \quad (7.1)$$

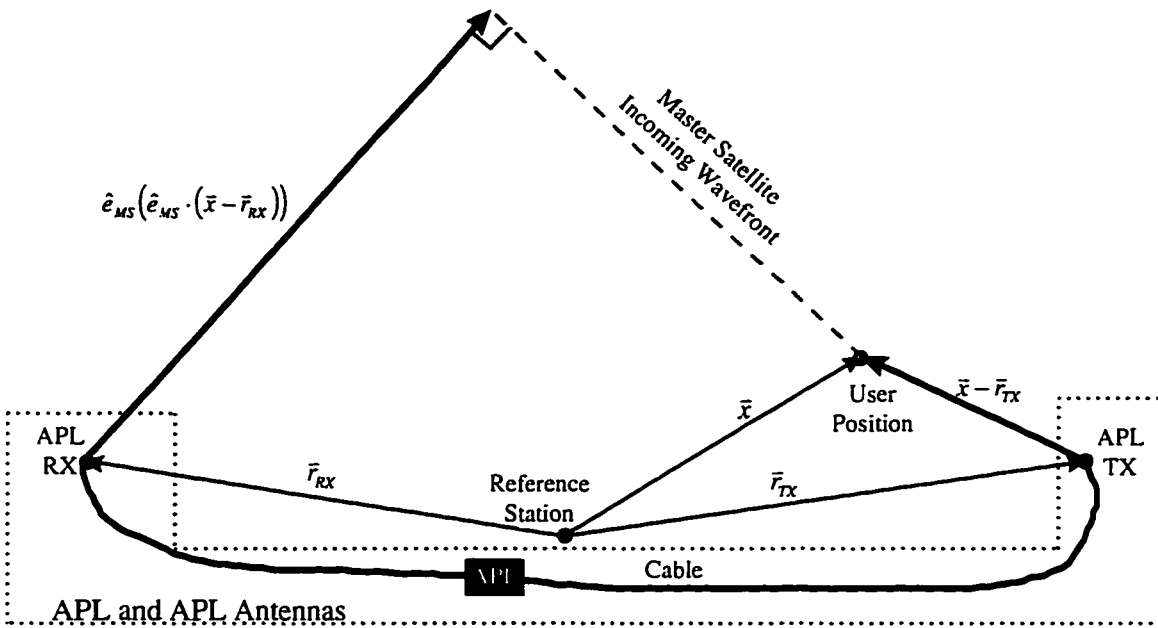


Figure 7.1:Autonomous Pseudolite Geometry

The measurement equation can be derived by referring to Figure 7.1. The carrier phase in the above equation represents the additional distance traveled by the reflected signal versus the direct signal. This additional distance is comprised of three parts:

- 1) The additional distance the signal from the master satellite travels before reaching the autonomous pseudolite receiver antenna. This distance is given by $\hat{e}_{MS} \cdot (\bar{x} - \bar{r}_{RX})$

where:

\hat{e}_{MS} is a unit vector in the direction of the master satellite.

\bar{x} is the position of the roving antenna relative to the reference station antenna.

\bar{r}_{rx} is the position of the autonomous pseudolite receive antenna relative to the reference station antenna.

2) The signal delay from the autonomous pseudolite receiver antenna, through the pseudolite electronics and the cable to the pseudolite transmit antenna. This delay is presumed constant and will be lumped into the integer cycle ambiguity.

3) The distance from the pseudolite transmit antenna to the roving antenna, $|\bar{x} - \bar{r}_{tx}|$ where \bar{r}_{tx} is the position of the autonomous pseudolite transmit antenna relative to the reference station antenna.

The measurement equation for the autonomous pseudolite is therefore:

$$\phi_{APL} = \hat{e}_{MS} \cdot (\bar{x} - \bar{r}_{rx}) + |\bar{x} - \bar{r}_{tx}| + C + \delta\phi$$

where:

C is a constant which includes the integer cycle ambiguity associated with the carrier phase.

$\delta\phi$ is the single difference phase measurement error.

To quantify the signal quality of the carrier phase received from the APL, the phase profile can be compared to the expected profile. First, the constant (C) is approximated:

$$\hat{C} = \phi_{APL} - \hat{e}_{MS} \cdot (\hat{x} - \bar{r}_{tx}) - |\hat{x} - \bar{r}_{rx}| \quad (7.2)$$

where \hat{x} is an estimate of position. If the cycle ambiguities for the satellites are known, this position estimate can be calculated with centimeter accuracy. The estimate of \hat{C} can be refined over several epochs to get an even better estimate. For all subsequent epochs, the expected phase is given by:

$$\hat{\phi}_{APL} = \hat{e}_{MS} \cdot (\hat{x} - \bar{r}_{TX}) + |\hat{x} - \bar{r}_{RX}| + \hat{C} \quad (7.3)$$

The difference between the expected phase and the measured phase will be referred to as the *phase profile error*. The error in the phase profile is a combination of noise on the position estimate and the carrier phase measurement error from the APL.

To use the additional kinematic ranging source provided by the APL, it is desirable to cast the measurement equation into a form similar to the satellite measurements. Recall that the differential carrier phase measurement equation for satellite number i is given by:

$$\phi_i = -\hat{e}_i \cdot \bar{x} + \tau + N_i + \delta\phi_i \quad (7.4)$$

Due to the proximity of the APL to the roving antenna, the measurement equation is nonlinear in position and must be linearized about some guess position, \bar{x}_0 :

$$\begin{aligned} \phi_{APL}(\bar{x}) &\equiv \phi_{APL}(\bar{x}_0) + \left. \frac{\partial \phi}{\partial \bar{x}} \right|_{\bar{x}_0} (\bar{x} - \bar{x}_0) \\ \left. \frac{\partial \phi}{\partial \bar{x}} \right|_{\bar{x}_0} &= \hat{e}_{MS} - \hat{e}_{TX}(\bar{x}_0) \end{aligned} \quad (7.5)$$

where $\hat{e}_{TX}(\bar{x}_0)$ is the unit line-of-sight vector from \bar{x}_0 towards the pseudolite transmit antenna. The linearized equation can then be written in a form similar to the satellite equation:

$$\phi_{APL} - \hat{e}_{MS} \cdot (\bar{x}_0 - \bar{r}_{RX}) - |\bar{x}_0 - \bar{r}_{TX}| \equiv (\hat{e}_{MS} - \hat{e}_{TX}(\bar{x}_0)) \cdot (\bar{x} - \bar{x}_0) + C + \delta\phi \quad (7.6)$$

The satellite and APL equations can be grouped together to form one matrix equation. To simplify notation, define $\phi^*(\bar{x}_0)$:

$$\phi^*(\bar{x}_0) = \begin{pmatrix} \phi_{satellites} \\ \phi_{APL} + \hat{e}_{MS} \cdot \bar{r}_{RX} - |\bar{x}_0 - \bar{r}_{TX}| \end{pmatrix} - \begin{pmatrix} N \\ C \end{pmatrix} - \begin{bmatrix} -\hat{e}_1^T \\ \vdots \\ -\hat{e}_n^T \\ \hat{e}_{MS}^T \end{bmatrix} \bar{x}_0 \quad (7.7)$$

The linearized matrix equation can then be written:

$$\phi^*(\bar{x}_0) = G \begin{pmatrix} \bar{x} - \bar{x}_0 \\ \tau \end{pmatrix} + \delta\phi \quad (7.8)$$

where G is the GPS geometry matrix for the system:

$$G = \begin{bmatrix} -\hat{e}_1^T & 1 \\ \vdots & \vdots \\ -\hat{e}_n^T & 1 \\ \hat{e}_{MS}^T - \hat{e}_{TX}^T & 0 \end{bmatrix}$$

The zero in the clock error column and APL row results from the direct and “reflected” signal being measured by a single receiver; there is no differential clock error for an APL measurement.

7.3.1 Three Satellite Positioning

Given the satellite cycle ambiguities, the APL measurement bias, C , and an initial position guess, the term $\phi^*(\bar{x}_0)$ can be calculated. Due to the additional ranging source provided by the APL, Equation (7.8) can be solved for position and clock error with only three satellites in view:

$$\begin{pmatrix} \bar{x} - \bar{x}_0 \\ \tau \end{pmatrix} = G^{-1} \phi^*(\bar{x}_0) \quad (7.9)$$

The guess position can be updated:

$$\bar{x}_0|_{k+1} = \bar{x}_0|_k + (\bar{x} - \bar{x}_0)|_k \quad (7.10)$$

This process is iterated through convergence to solve for position and clock bias.

7.3.2 Four Satellite Fault Detection

As with positioning, the ranging source provided by the APL allows RAIM-based fault detection to be performed with one fewer satellite. Traditionally, five satellites are required for fault detection; with an APL, only four are needed. Fault detection is performed by first calculating a position fix. The procedure for calculating position is almost identical to the three satellite case in Section 7.3.1. However, the initial position guess can be found by using the four satellites to do a traditional position solution:

$$\begin{pmatrix} \bar{x} \\ \tau \end{pmatrix} = (G_{sats}^T G_{sats})^{-1} G_{sats}^T \phi_{sats} \quad (7.11)$$

Using that initial guess, the estimate can be updated by finding the least squares solution to the over-determined linearized equation:

$$\begin{pmatrix} \bar{x} - \bar{x}_0 \\ \tau \end{pmatrix} = (G^T G)^{-1} G^T \phi^*(\bar{x}_0) \quad (7.12)$$

Again, the process is iterated through convergence. Because there are now five measurements and only four unknowns, a residual can be calculated:

$$r \equiv \phi^*(\bar{x}_f) - G \begin{pmatrix} \bar{0} \\ \tau_f \end{pmatrix} \quad (7.13)$$

where:

\bar{x}_f is the final position solution.

τ_f is the final clock error solution.

If the magnitude of this residual is greater than some fault detection threshold, a RAIM alert is issued. Due to the precision of the carrier phase, this threshold can be set quite tight (less than 10 cm) without incurring high false alarm rates.

7.3.3 Five Satellite Fault Isolation

With six ranging sources available, single-satellite faults can be isolated. One method of fault isolation is parity space analysis (see for example [Pervan, b]). The initial procedure is identical to the fault detection. The parity vector is derived from the residual vector, r :

$$p = Lr \quad (7.14)$$

where:

p is the parity vector.

The rows of L form an orthonormal basis for the left null space of G . That is, the rows of L are orthogonal to the columns of G .

Note that the residual vector lies in the left null space of G . This fact can be derived from the definition of r and \bar{x}_f . Convergence implies that \bar{x} is no longer changing:

$$\begin{pmatrix} \bar{0} \\ \tau_f \end{pmatrix} = (G^T G)^{-1} G^T \phi^*(\bar{x}_f) \quad (7.15)$$

Pre-multiplying Equation (7.13) by G^T :

$$G^T r = G^T \phi^*(\bar{x}_f) - G^T G \begin{pmatrix} \bar{0} \\ \tau_f \end{pmatrix} = G^T \phi^*(\bar{x}_f) - G^T G (G^T G)^{-1} G^T \phi^*(\bar{x}_f) = \bar{0} \quad (7.16)$$

This shows that r is in the left null space of G , the dimension of which is $n - 4$. The transformation in Equation (7.14) takes the residual vector, which has dimension ns , and transforms it to the parity vector, which has dimension $n - 4$. The vector p contains all of the information contained in r and has the same magnitude. However, the parity vector is easier to visualize and analyze due to its lower dimension.

With six ranging sources available, the parity vector has dimension two. In this two dimensional space, each single satellite (or pseudolite) failure has an associated modal direction. An example of this is illustrated in the next section by Figure 7.5, which shows these modal directions for the geometry associated with one flight test. The simplest technique for fault isolation would remove the ranging source whose modal direction is nearest the parity vector. Of course, isolation is only performed if a fault is detected. In

Figure 7.5, a fault is detected if the tip of the parity vector lies outside the detection threshold circle.

7.4 Flight Tests

A series of flight tests was performed to obtain preliminary data on the performance of Airport Pseudolites. Two of these tests were representative of the entire flight test campaign, a ranging test and a near-far test.

7.4.1 Ranging Test

The goal of this test was to quantify the kinematic ranging performance of an APL. A differential reference station was set up on the Stanford campus, 6 km from the airport. An autonomous APL was set up near the far end of the landing runway. (In practice, it would be desirable to locate an APL at a central location servicing all runways. For these preliminary tests however, it was not convenient to try to service both runways). The carrier phase cycle ambiguities were resolved on the ground using methods discussed in Chapter 4 and were maintained throughout the test. Centimeter-level positioning truth was therefore available. Three takeoffs and landings were performed (two touch-and-goes and one full-stop landing). The ranging performance of the APL was assessed for each final approach until after touchdown.

7.4.1.1 Phase Profile Error

Figure 7.2 plots the phase profile error for two of the approaches using the tail antenna. The first approach is plotted with O's and the second approach with X's. The third approach was similar to the first and is not shown in the plot.




Figure 7.2: Phase Profile Error

The phase profile error for the first approach has a 1σ noise level of 2.4 cm and a maximum of 5.6 cm. The signal was continuously tracked for the entire first approach. The second pass did not perform as well as the first. The APL signal was lost twice as seen by the discontinuities in the line with the X's. The cause of the dropouts was unknown at the time of this writing. The tail antenna had a direct line-of-sight to the APL for the entire approach, so signal blockage was not the cause. More recent flight tests using improved pseudolites and receivers have not suffered from these signal losses [Lawrence, c]. Therefore, the signal interruptions resulted from a problem with the original receiver tracking pulsed pseudolites or from hardware problems with the original pseudolites. The continued study of pseudolite signal quality is the subject of ongoing research.

7.4.1.2 Nominal Performance

Apart from the signal losses, the ranging signal provided by the APL appears suitable for tasks such as positioning, fault detection, and fault isolation. The data from the first

approach was used to demonstrate each of those tasks. In each case, the task can be performed with one less satellite than is traditionally required.

7.4.1.3 Three Satellite Positioning

During the first approach, five or six satellites were always available. The true trajectory of the aircraft was calculated using all of the satellites. Then, all but three of the satellites were artificially removed, and the APL was used to provide the fourth ranging source required for positioning. The three satellite trajectory was compared to the true trajectory and the difference is plotted in Figure 7.3. The mean position difference was only 2 cm.




Figure 7.3: Three Satellite Positioning

7.4.1.4 Four Satellite Fault Detection

To demonstrate fault detection with four satellites, all but four satellites were removed and the APL phase measurement was added to the solution. With five ranging sources, the residual was calculated as described in Section 7.3.2. The residual profile for the approach was first calculated for the actual phases measured in flight. A cycle slip of two cycles was

then simulated on an arbitrary channel midway through the approach and the residual profile was recalculated. Figure 7.4 plots both residual profiles. A fault detection threshold of 10 cm was used. This threshold is comfortably above the residual profile under normal conditions to prevent false alarms. However, the residual after the simulated fault was around 30 cm, well above the detection threshold. In this situation, the APL measurement allowed detection of a fault that would have otherwise gone unnoticed.




Figure 7.4: Four Satellite Fault Detection

7.4.1.5 Five Satellite Fault Isolation

With one additional satellite, a single-channel fault can not only be detected, but isolated as well. Five satellite fault isolation was demonstrated in a manner similar to fault detection. Under normal conditions, no fault was detected during the approach. A cycle slip of two cycles was simulated on the first channel of the receiver. Figure 7.5 shows the parity space for the given geometry. The modal directions for each of the six channels are plotted, and the “x” represents the tip of the parity vector. The fault was detected, because the parity vector lies outside the detection threshold circle. The proximity of the parity vector to the

first channel modal direction caused the isolation algorithm to correctly isolate a failure on the first channel. Centimeter-level positioning accuracy was therefore maintained throughout the approach, despite a 40 cm phase error induced on one channel.




Figure 7.5: Five Satellite Fault Isolation

7.4.2 Near-Far Test

Another flight test was performed to demonstrate the effectiveness of pulsing pseudolites in eliminating the Near-Far problem [Cobb, *a*]. A pseudolite was pulsed with a duty cycle of 10% and a peak transmit power of 25 dBm. To verify that there was no near radius, several models of GPS hand-held receivers were placed within a foot of the pseudolite transmit antenna. The receivers were able to track satellites and perform position fixes with

negligible degradation. At the same time, a receiver in the Piper Dakota was able to track the pseudolite signal up to a distance of about 20 km. Therefore, it seems that pulsed transmission does solve the near-far problem⁴.

7.5 Stacked APLs

Until now, this section has discussed only the performance of a single APL. Assuming the APL signal noise level and continuity are comparable to a satellite signal (recent flight trials indicate that this is so), a single APL provides the benefits of an additional ranging source:

- 1) Improved DOP
- 2) Improved Availability/Continuity
- 3) A small amount of geometry change for cycle ambiguity resolution (depending on where the APL is located, this geometry change may or may not be greater than that provided by a satellite during the course of a final approach).

At first thought it would seem that each additional APL would provide only these same benefits. However, the presence of multiple pseudolites provides a new observable that can improve the accuracy performance beyond the level expected from an additional ranging source. The new observable is the *difference* between the carrier phase of the pseudolites. One configuration that takes advantage of this observable to improve vertical accuracy is to

⁴ Some performance degradation was noticed with higher quality NovAtel receivers close to the pulsed APL, but design changes could obviate the problem [Cobb, a].

place APLs at the top and bottom of a tower; the method has therefore come to be called the “stacked APL” method.

7.5.1 The Stacked APL Concept

The measurement equation for the differential carrier phase between two APLs is given by Equation (7.17). The derivation of this equation assumes the line biases are calibrated for the two APL antennas such that the fractional number of cycles in the difference between line biases is known. Although this calibration is not necessary, without it the bias N will not be an integer. The following analysis assumes this calibration has been performed; [Lawrence, c] discusses the minor impact of not calibrating line biases.

$$\varphi = \left| \bar{x} - \bar{r}_{APL_2} \right| - \left| \bar{x} - \bar{r}_{APL_1} \right| + N + \delta\varphi \quad (7.17)$$

where:

φ is the carrier phase difference between the two APLs, corrected for any non-integral difference between the APL line biases.

\bar{x} is the position of the aircraft relative to the reference station.

\bar{r}_{APL_i} is the position of the i^{th} APL relative to the reference station.

N is the integer cycle ambiguity for the phase difference.

$\delta\varphi$ is the error in the measurement of φ .

The form of Equation (7.17) shows that the surfaces of constant phase difference are hyperbolic. Far away from the APLs, the hyperbolic surfaces are well approximated by straight lines as shown in Figure 7.6. The figure shows two APLs “stacked” with one at the base of a tower and one at the top. Hypothetical lines of constant phase are shown for different values of the integer cycle ambiguity (N). Note that when the aircraft is far from

the APLs the lines of constant phase are far apart. At this distance, a code phase DGPS position estimate has sufficient accuracy to determine which line of constant phase the aircraft is on, thereby resolving the cycle ambiguity. As the aircraft approaches the APLs, the lines of constant phase converge, improving the positioning geometry of the differential APL observable.




Figure 7.6: Stacked APL Geometry

The potential performance of a stacked APL system can be quantified by linearizing Equation (7.17) to find the sensitivity of φ to changes in \bar{x} as calculated in Equation (7.18)

$$\frac{d\varphi}{d\bar{x}} = \hat{e}_2 - \hat{e}_1 \equiv \Delta e \quad (7.18)$$

where \hat{e}_i is a unit vector from the aircraft toward the i^{th} APL.

The differential APL carrier phase observable therefore provides information parallel to Δe , the difference in the direction of the APLs seen from the aircraft. No information is provided normal to Δe . The dilution of precision (DOP) of the measurement, defined as

the position error parallel to Δe for a unit error $\delta\phi$, is $\frac{1}{|\Delta e|}$. For the configuration shown in Figure 7.6, $|\Delta e|$ can be approximated by:

$$|\Delta e| \cong \frac{h}{d} \quad (7.19)$$

Far from the APLs, $|\Delta e|$ is small, so the lines representing different possible cycle ambiguities are far apart (the distance between these lines is $\frac{1}{|\Delta e|}$ cycles). For example, 10000 feet from a 50 foot APL tower, the ambiguity lines are about 40 meters apart, so the cycle can be easily resolved using code phase DGPS.

DGPS accuracy performance is typically the worst in the vertical direction; however, this direction is most critical for an aircraft approach guidance system above 100 feet. Between the altitudes of 100 and 50 feet, the on-board radar altimeter begins providing a supplemental source of vertical information. Therefore, the performance of the APL will be judged based on its vertical error at an altitude of 75 feet, about 1500 feet from touchdown. At a distance of 1500 feet from a 36 foot APL tower (the height of an ILS glide slope antenna), a 1.5 cm error in ϕ will result in a position error of about 60 cm. Note that Δe is approximately vertical for the tower configuration; so this new information provides a new source of vertical positioning information accurate to about 60 cm. When combined with the existing satellite DGPS position estimates and when filtered over the entire approach, the vertical position estimate will improve.

7.5.2 Multipath

Because a stacked APL system is quite similar to an ILS glideslope, it may suffer from some of the same problems. In particular, ILS has been plagued by multipath reflected from the ground, aircraft (both taxiing and holding), and buildings. For ILS, this problem has been solved by proper sighting and design of antennas. Additionally “critical zones” have been established to prevent vehicular traffic near the ILS antennas. Stacked APLs can benefit from this heritage. Pseudolite antennas will be designed to minimize multipath by directing the bulk of the signal power up the glideslope. Also, the structure of the GPS signal is potentially more robust to multipath interference than ILS due to the shorter wavelength of the carrier and the spread-spectrum design of the signal. Therefore, the procedural burden of introducing large critical zones for APLs may not be necessary. However, further study is required in this area.

7.5.3 In-track APLs

The analysis of Section 7.5.1 lends insight for the development of other configurations which take advantage of the differential carrier phase measurement between APLs. The important parameters in this development are the size and direction of Δe . At a distance, the magnitude of Δe should be small to allow for cycle ambiguity resolution. Near the airport, $|\Delta e|$ should increase to improve the DOP of the differential measurement. The direction of Δe near the airport should be nearly vertical to improve the vertical performance of the system. A configuration that meets these requirements is shown in Figure 7.7 where an APL is placed near each side of a runway. Although the APLs are spaced longitudinally, their appearance to an aircraft on final approach will be nearly identical to the stacked configuration. The in-track configuration has the advantage of

servicing both runway ends without the need for a tower. The long baseline can provide DOP equivalent to a high tower.




Figure 7.7: In-track “Stacked” APLs

With a longitudinal baseline, the horizontal component of Δe is not negligible. The following 2D analysis shows how the APL differential carrier phase measurement can combine with a single code DGPS position solution to provide improved vertical performance. Assume the code DGPS position covariance ($P_{\hat{x}_{code}}$) is given by:

$$P_{\hat{x}_{code}} = \begin{bmatrix} \sigma_h^2 & 0 \\ 0 & \sigma_v^2 \end{bmatrix} \quad (7.20)$$

where σ_h^2 and σ_v^2 are the horizontal and vertical variances respectively. Assume also that the variance of the differential APL carrier phase measurement error is σ_ϕ^2 . This information can be shown graphically by the covariance ellipses as in Figure 7.8. There is no position information provided by the differential APL measurement in the direction normal to Δe , so the corresponding covariance ellipse is infinitely long in that direction. In the direction parallel to Δe , the width of the ellipse scales with $\frac{1}{|\Delta e|}$.




Figure 7.8: Stacked APL Covariance Analysis

Because Δe is not perfectly vertical, the APL measurement must be combined with the horizontal information provided by code DGPS to yield a vertical position estimate. The resulting variance ($\sigma_{v_{APL}}^2$) is approximated by:

$$\sigma_{v_{APL}}^2 \equiv \frac{\sigma_{\varphi}^2}{|\Delta e|^2} + \theta^2 \sigma_h^2 \quad (7.21)$$

where θ is the angle between Δe and vertical.

The combined vertical variance ($\sigma_{v_{combined}}^2$) will be better than both σ_v^2 and $\sigma_{v_{APL}}^2$ as determined by Equation (7.22).

$$\sigma_{v_{combined}}^2 \equiv \frac{1}{\frac{1}{\sigma_v^2} + \frac{1}{\sigma_{v_{APL}}^2}} \quad (7.22)$$

For the layout of Figure 7.7, when the aircraft is at a height of 75 feet, the geometry is such that:

$$\Delta e = \begin{pmatrix} 0.0043 \\ 0.0860 \end{pmatrix}$$

$$\frac{1}{|\Delta e|} = 11.6$$

$$\theta = 0.05 \text{ rad}$$

Assuming $\sigma_h = 1m$, $\sigma_v = 1.5m$, and $\sigma_\phi = 0.02m$, the resulting combined vertical standard deviation ($\sigma_{v_{combined}}$) at 75 feet is 23 centimeters, a six times improvement over code DGPS alone. Therefore, this concept is promising, but further study is required.

CHAPTER 8

Conclusion

8.1 Summary of Results and Future Work

Figure 8.1 shows the contributions made in this research and how they relate to one another.

The central contribution was the flexible airborne measurement processing system from which the others were derived. The following sections summarize the results of each area of research and offer suggestions for further work.

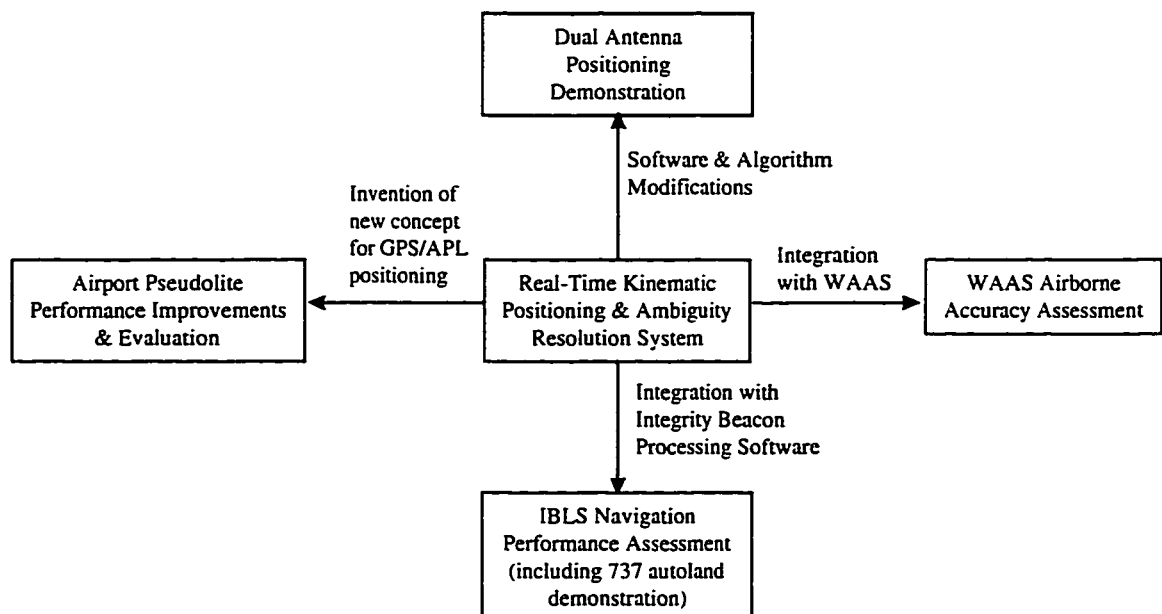


Figure 8.1: Research Contributions

8.1.1 Airborne Measurement Processing

The real-time system developed to demonstrate the Integrity Beacon Landing System was designed to be easy to implement yet provide the flexibility required of a research tool.

Distinguishing features of this architecture include:

- Only the L1 carrier and C/A code are used.
- Kinematic position solutions are provided with minimal delay in all phases of flight. Integer estimates are continuously refined.
- Several layers of RAIM provide high-integrity. The Integrity Beacon information is particularly powerful in this respect.
- Cycle ambiguities are found without using integer searches.
- Integer estimates are smoothly handed off as satellites are acquired and lost.
- Additional information such as static constraints is easily incorporated into the system.

8.1.2 IBLS Performance Assessment

This research shows the accuracy of IBLS is on the centimeter level, exceeding any realistic RNP accuracy specification. Pervan [Pervan, a] demonstrated that the integrity of IBLS can also meet the RNP specifications for Cat III instrument approach. This was achieved through analysis and computer simulation assuming prior probability failure models and validated through flight test. Plans for future work include continued validation of these integrity results through the use of a GPS simulator. The simulator will enable the study of

millions of approaches with GPS receivers in the loop and with artificially induced failures and multipath effects.

Although the accuracy and integrity of IBLs appear to meet the RNP specifications, the availability and continuity of GPS-based landing systems appear to require improvements and the practicality of the method would be improved if the pseudolites were on airport property. Methods of achieving these improvements were discussed in Chapters 6 and 7. The current research thrust has shifted to developing the most inexpensive system that will meet all of the RNP specifications. Since IBLs requires the use of pseudolites beneath the approach path, it would be more cost-effective to have an architecture that served more than one runway end. Ideally, one DGPS installation will provide Cat III performance to all runways at an airport. Currently, it is unclear whether such a system is possible. However, it is clear that such a system will shift some of the burden of integrity monitoring from RAIM-based methods to ground integrity monitoring. The development of ground monitoring algorithms is currently underway.

8.1.3 Wide Area DGPS Integration with Local Area DGPS

The Stanford Wide Area Augmentation System test-bed was integrated with the Integrity Beacon Landing System user software. This integration provides a centimeter-level positioning truth source for WAAS flight tests. Using this truth source, the Stanford WAAS network was evaluated in-flight. The RMS 3D position error for WAAS was approximately 2.35 meters over a one hour flight.

Future work includes the continued study of airborne WAAS performance. Additionally, the transition from en route to terminal area guidance will be explored.

8.1.4 Dual Antenna Positioning

Centimeter-level kinematic positioning was demonstrated in flight using dual (top and bottom) antennas. With the appropriate hardware, nearly 4π steradian visibility of satellites and pseudolites should be attainable using opposing antennas. Multiple antennas can greatly improve the availability and continuity of GPS positioning in banked turns without a large increase in expense. This improvement may be significant for applications such as curved approaches. Navigation in mountainous areas, at high latitudes, and during times of satellite failures is particularly reliant on maintaining continuous lock on all satellites in view. It is under these conditions that aircraft can benefit the most from multiple antenna positioning. However, multipath from the ground can be more severe when antennas are not oriented vertically. Care should be taken to ensure that the direct signal dominates the received signal. Ideal applications for multiple antenna positioning are systems that already have multiple antennas such as IBLS and GPS-based attitude determination systems. An area for future work is the development of a fully integrated attitude/position determination GPS receiver with more than four antenna inputs. Such a receiver would have applications in both spacecraft and aircraft navigation.

8.1.5 Airport Pseudolites

Preliminary flight tests of Airport Pseudolites (APLs) showed promising results. The APLs provided kinematic phase noise on the order of a few centimeters. Using the additional

ranging source provided by an APL, centimeter-level positioning was demonstrated with three satellites, fault detection was demonstrated with four satellites, and fault isolation was demonstrated with five satellites. Additionally, a promising new technique was developed for positioning with multiple APLs. Areas of future work involving Airport Pseudolites are listed below:

- Investigate APL signal robustness, both in flight and in the lab.
- Explore antenna location sensitivity, both transmit and receive.
- Try antennas with better low-elevation gain patterns such as quad-helices.
- Quantify the performance improvements offered by APLs.
- Flight test the “stacked APL” concept to validate the theoretically predicted performance.

8.2 Applications

The applications of GPS in general are too numerous to attempt to list here. Instead, the applications of this research already being pursued will be discussed. Apart from the original goal of using DGPS for aircraft landing, several applications are currently being developed.

One obvious application of DGPS systems is the guidance of fully automatic land and air vehicles. Many repetitive tasks that are well suited for computer control are currently being performed manually due to the lack of accurate guidance. Automated farming is one example that promises to revolutionize the industry. Current research at Stanford focuses on the autonomous control of a farm tractor with a towed implement. A similar area of

research uses kinematic GPS and pseudolites for positioning in an open-pit mine environment. Other applications include unmanned aircraft for reconnaissance, mapping, and tactical missions.

Pilot display research also benefits from DGPS. The goal of this current research is to develop an inexpensive display (affordable to General Aviation pilots) to improve the situational awareness of pilots in IFR conditions. The high accuracy and low cost of GPS positioning and attitude determination are enabling factors in the development of such systems.

8.3 Closing

The future of aircraft navigation will undoubtedly be dominated by satellite-based systems such as GPS. These systems will likely be comprised of a mixture of commercial and military satellites from several countries, forming an integrated network with enough satellites to provide guaranteed global coverage with significant redundancy. Until such a system exists, the potential performance of the current system must be understood and improved when necessary. The research presented in this dissertation has made several contributions toward this goal. The lessons learned in these early days of GPS research will form a foundation for the development of the truly global positioning systems of the future.

References

- [AC120-28C] Federal Aviation Administration, *Advisory Circular 120-28C*: “Criteria for Approval of Category III Landing Weather Minima,” Washington, D.C., 9 March 1984.
- [AIM] Federal Aviation Administration, *Federal Aviation Regulations/Aeronautical Information Manual* (FAR/AIM), November, 1995.
- [Beser] J. Beser and B. W. Parkinson, “The Application of NAVSTAR Differential GPS in the Civilian Community,” *NAVIGATION: Journal of the Institute of Navigation*, Vol. 29, No. 2, Summer 1982.
- [Black] H. D. Black, “An Easily Implemented Algorithm for Tropospheric Range Correction,” *Journal of Geophysical Research*, 38(B4), 1978.
- [Brown] A. Brown, F. Van Diggelen, T. Kelecy, “Flight Test Results Using Carrier Phase Kinematic Solutions for Aircraft Precision Approach and Landing,” ION 50th Annual Meeting, Colorado Springs, CO, June, 1994.
- [Carlson] A. B. Carlson, *Communication Systems*, McGraw Hill, Inc., New York, NY, 1986.
- [Chao] C. C. Chao, “A Model for Tropospheric Calibration from Daily Surface and Radiosonde Balloon Measurements,” JPL Technical Memorandum 391-350, Jet Propulsion Laboratory, Pasadena, California, 1972.
- [Cobb, a] H. S. Cobb, Stanford University Ph.D. Dissertation, Department of Aeronautics and Astronautics, 1996.

- [Cobb, *b*] H. S. Cobb, D. G. Lawrence, B. S. Pervan, C. E. Cohen, J. D. Powell, B. W. Parkinson, "Precision Landing Tests with Improved Integrity Beacon Pseudolites," ION GPS-95, Palm Springs, CA, September, 1995.
- [Cobb, *c*] H.S. Cobb, D. Lawrence, J. Christie, T. Walter, Y.C. Chao, J.D. Powell, B.W. Parkinson, "Observed GPS Signal Continuity Interruptions," ION GPS-95, Palm Springs, CA, September, 1995.
- [Cohen, *a*] C. E. Cohen, "Attitude Determination Using GPS," Stanford University Ph.D. Dissertation, Department of Aeronautics and Astronautics, December 1992.
- [Cohen, *b*] C. E. Cohen, B. S. Pervan, D. G. Lawrence, H. S. Cobb, J. D. Powell, B. W. Parkinson, "Real-Time Flight Test Evaluation of the GPS Marker Beacon Concept for Category III Kinematic GPS Precision Landing," ION GPS-93, Salt Lake City, Utah, September 1993.
- [Cohen, *c*] C. E. Cohen, B. S. Pervan, H. S. Cobb, D. G. Lawrence, J. D. Powell, and B. W. Parkinson, "Real-Time Cycle Ambiguity Resolution using a Pseudolite for Precision Landing of Aircraft with GPS," Second International Conference on Differential Satellite Navigation Systems (DSNS-93), March 30 - April 2, 1993, Amsterdam, The Netherlands.
- [Cohen, *d*] C. E. Cohen, D. G. Lawrence et al., "Flight Test Results of Autocoupled Approaches Using GPS and Integrity Beacons," ION GPS-94, Salt Lake City, Utah, September 1994.
- [Cohen, *e*] C. E. Cohen, D. G. Lawrence et al., "Automatic Landing of a 737 using GNSS Integrity Beacons," International Symposium on Precision Approach and Landing (ISPA 95), Braunschweig, Germany, February 1995.

- [Counselman] C. C. Counselman and I. I. Shapiro, "Miniature Interferometer Terminals for Earth Surveying," *Bulletin Geodesique*, Vol. 53, No. 2, pp. 139-163, 1979.
- [Enge] P. Enge, E. Swanson, R. Mullin, K. Ganther, A. Bommarito, R. Kelly, "Terrestrial Radio Navigation Technologies," *NAVIGATION: Journal of the Institute of Navigation*, Vol. 42, No. 1, Spring 1995.
- [Gelb] A. Gelb, ed., *Applied Optimal Estimation*, M.I.T. Press, Cambridge, MA, 1974.
- [Goad] C. C. Goad, "Optimal Filtering of Pseudoranges and Phases from Single-Frequency GPS Receivers," *NAVIGATION: Journal of the Institute of Navigation*, Vol. 37, No. 3, Fall 1990.
- [Green] G. B. Green, P. D. Massatt, and N. W. Rhodus, "The GPS 21 Primary Satellite Constellation," *NAVIGATION: Journal of the Institute of Navigation*, Vol. 36, No. 1, Spring 1989.
- [Hatch, a] R. R. Hatch, "Instantaneous Ambiguity Resolution," KIS Symposium, Banff, Canada, September 1990.
- [Hatch, b] R. R. Hatch, "The Synergism of Code and Carrier Measurements," *Proceedings of the Third International Geodetic Symposium on Satellite Doppler Positioning*, Las Cruces, New Mexico, February 1982.
- [Hopfield] H. S. Hopfield, "Tropospheric Range Error Parameters: Further Studies," *Applied Physics Laboratory Technical Memorandum TG1291*, Johns Hopkins University, Baltimore, Maryland, 1976.
- [Hundley] W. Hundley, S. Rowson, G. Courtney, V. Wullschleger, R. Velez, R. Benoist, P. O'Donnell, "FAA-Wilcox Electric Category IIIB Feasibility Demonstration Program Flight Test Results," *ION GPS-95*, Palm Springs, CA, September 1995.

- [Hwang] P. Y. C. Hwang and R. G. Brown, "GPS Navigation: Combining Pseudorange with Continuous Carrier Phase Using a Kalman Filter," *NAVIGATION: Journal of the Institute of Navigation*, Vol. 37, No. 2, Summer 1990.
- [Janes] H. W. Janes, R. B. Langley, S. P. Newby, "A Comparison of Several Models for the Prediction of Tropospheric Propagation Delay," Proceedings of the 5th International Geodetic Symposium on Satellite Positioning, Las Cruces, New Mexico, 1989.
- [Kayton] M. Kayton, ed., *Navigation: Land, Sea, Air, and Space*, IEEE Press, New York, 1990.
- [Kee] C. Kee "Wide Area Differential GPS (WADGPS)," Stanford University Ph.D. Dissertation, Department of Aeronautics and Astronautics, December 1993.
- [Kelly] R. J. Kelly and J. M. Davis, "Required Navigation Performance (RNP) for Precision Approach and Landing with GNSS Application," *NAVIGATION: Journal of the Institute of Navigation*, Vol. 41, No. 1, Spring 1994.
- [Kinal] G. Kinal, "Consideration of Practical Visibility of Geostationary Augmentation Satellites from Aircraft", ION National Technical Meeting, Anaheim, CA, January, 1995.
- [Klein] D. Klein and B. W. Parkinson, "The Use of Pseudo-Satellites for Improving GPS Performance," *NAVIGATION: Journal of the Institute of Navigation*, Vol. 31, No. 4, Winter 1984-85.
- [Klobuchar] J. A. Klobuchar, "Ionospheric Time-Delay Algorithm for Single-Frequency GPS Users," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. AES-23, No. 3, May 1987.

- [LARC] *Stanford University Report to the LAAS Architecture Review Committee (LARC)*, May 1996, Washington, D.C.
- [Lawrence, b] D. G. Lawrence, B. S. Pervan, C. E. Cohen, H. S. Cobb, J. D. Powell, and B. W. Parkinson, "Real-Time Architecture for Kinematic GPS Applied to the Integrity Beacon Landing System," ION 51st Annual Meeting, Colorado Springs, Colorado, June 1995.
- [Lawrence, c] D. G. Lawrence, B. S. Pervan, H. S. Cobb, G. Opshaug, J. D. Powell, P. Enge, B. W. Parkinson, "Performance Evaluation of On-Airport Local Area Augmentation System Architectures," to be presented at ION GPS-96, Kansas City, Missouri, September, 1996.
- [LIP] *Stanford University Report to the LAAS Integrity Panel (LIP)*, January 1996, Washington, D.C.
- [Loomis] . P. Loomis, "A Kinematic Double-Differencing Algorithm", Proceedings of the Fifth International Geodetic Symposium on Satellite Positioning, New Mexico State University, Las Cruces, NM, March, 1989.
- [MacDoran] P. MacDoran, "Satellite Emission Radio Interferometric Earth Surveying (SERIES)—GPS Geodetic System," Bulletin Geodesique, Vol. 53, No. 2, pp. 117-138, 1979.
- [Masson] A. Masson, AA 290 Graduate Research Final Report, Department of Aeronautics and Astronautics, Stanford University, Stanford, CA, April 1994.
- [Ndili] A. Ndili, "GPS Pseudolite Signal Design," ION GPS-94, Salt Lake City, UT, September, 1994

- [ORD] *Operational Requirements Document: Local Area Augmentation System*, Satellite Navigation Program Office, Federal Aviation Administration, February 28, 1995.
- [Paielli] R. A. Paielli, B. D. McNally, R. E. Bach, and D. N. Warner, "Carrier Phase Differential GPS for Approach and Landing: Algorithms and Preliminary Results," ION GPS-93, Salt Lake City, Utah, September 1993.
- [Parkinson, a] B. W. Parkinson, "Introduction and Heritage of NAVSTAR," contribution to *The Global Positioning System: Theory and Applications*, B. W. Parkinson, J. J. Spilker, Jr., P. Axelrad, P. Enge, editors, AIAA Series, 1996.
- [Parkinson, b] B. W. Parkinson and K. T. Fitzgibbon, "Aircraft Automatic Landing Systems using GPS," *Journal of Navigation*, Vol. 42, No. 1, January 1989.
- [Pervan, a] B. S. Pervan, "Navigation Integrity for Aircraft Precision Landing Using the Global Positioning System," Stanford University Ph.D. Dissertation, Department of Aeronautics and Astronautics, March 1996.
- [Pervan, b] B. S. Pervan, C. E. Cohen, B. W. Parkinson, "Parity Space Methods for Autonomous Fault Detection and Isolation using GPS Carrier Phase," IEEE Position, Location, and Navigation Symposium (PLANS '96), Atlanta, GA, April 1996.
- [Phlong] W. S. Phlong and B. D. Elrod, "Availability Characteristics of GPS and Augmentation Alternatives," NAVIGATION: *Journal of the Institute of Navigation*, Vol. 40, No. 4, Winter 1993-94.
- [Romrell] G. Romrell, R. Brown, G. Johnson, and D. Kaufman, "FAA/FEDSIM – E-Systems Cat IIIB Feasibility Demonstration Flight Test Preliminary Results," ION GPS-95, Palm Springs, CA, September 1995.

- [Rowson] S. Rowson, G. Courtney, and R. Hueschen, "Performance of Category IIB Automatic Landings Using C/A Code Tracking Differential GPS," ION National Technical Meeting, January 1994, San Diego, CA.
- [Skillicorn] G. N. Skillicorn, "The FAA's Program to Sustain and Improve ILS Capability In a Rapidly Changing Environment," International Symposium on Precision Approach and Landing (ISPA-95), Braunschweig, Germany, February 21-24, 1995.
- [Saastamoinen] J. Saastamoinen, "Contributions to the Theory of Atmospheric Refraction," *Bulletin Geodesique*, 105-107, 1973.
- [Spilker] J. J. Spilker, B. W. Parkinson, "GPS Operation and Design," contribution to *The Global Positioning System: Theory and Applications*, B. W. Parkinson, J. J. Spilker, Jr., P. Axelrad, P. Enge, editors, AIAA Series, in publication, 1996.
- [SPS] "Global Positioning System Standard Positioning Service Signal Specification (Second Edition)," June 2, 1995.
- [Stansell] T.A. Stansell. "RTCM SC-104 Recommended Pseudolite Signal Specification," *NAVIGATION: Journal of the Institute of Navigation*, Vol. 33, No. 1, Spring 1986.
- [Teasley] S. P. Teasley, W. M. Hoover, and C. R. Johnson, "Differential GPS Navigation," Position, Navigation, and Location Symposium (PLANS 80), Atlantic City, New Jersey, December 1980.
- [Van Dierendonck] A.J. van Dierendonck. "The Role of Pseudolites in the Implementation of Differential GPS" IEEE Position, Location, and Navigation Symposium, April 1990.

- [Van Graas, a] F. van Graas, D. Diggle, R. Hueschen, "Interferometric GPS Flight Reference/Autoland System," ION GPS-93, Salt Lake City, Utah, September 1993.
- [Van Graas, b] F. van Graas, D. Diggle, M. Uijt, V. Wullschleger, R. Velez, D. Lamb, M. Dimeo, G. Kuehl, and R. Hilb, "FAA/Ohio University/United Parcel Service DGPS Autoland Flight Test Demonstration," ION GPS-95, Palm Springs, CA, September 1995.
- [Vieweg] S. Vieweg, I. G. Schaenzer, "Error Estimation in a Hybrid Flight Navigation System Using Realtime-Differential GPS and Inertial Measurement Units," First International Symposium Real Time Differential Applications of the Global Positioning System, Braunschweig, Germany, September, 1991.
- [Walter] T. Walter, C. Kee, Y. C. Chao, Y. J. Tsai, et. al., "Flight Trials of the Wide Area Augmentation System (WAAS)," ION GPS-94, Salt Lake City, Utah, September, 1994.
- [Zumberge] J. F. Zumberge and W. I. Bertiger, "Ephemeris and Clock Navigation Message Accuracy," contribution to *The Global Positioning System: Theory and Applications*, B. W. Parkinson, J. J. Spilker, Jr., P. Axelrad, P. Enge, editors, AIAA Series, in publication, 1996.