

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.



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

Order Number 9309582

**Attitude determination using GPS: Development of an all
solid-state guidance, navigation, and control sensor for air and
space vehicles based on the global positioning system**

Cohen, Clark Emerson, Ph.D.

Stanford University, 1993

Copyright ©1992 by Cohen, Clark Emerson. All rights reserved.

U·M·I

**300 N. Zeeb Rd.
Ann Arbor, MI 48106**

ATTITUDE DETERMINATION USING GPS

*DEVELOPMENT OF AN ALL SOLID-STATE
GUIDANCE, NAVIGATION, AND CONTROL SENSOR
FOR AIR AND SPACE VEHICLES
BASED ON 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
Clark Emerson Cohen

December, 1992

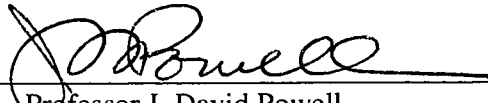
© Copyright by Clark Cohen 1992
All Rights Reserved

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



Professor Bradford W. Parkinson
(Principal Advisor)

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



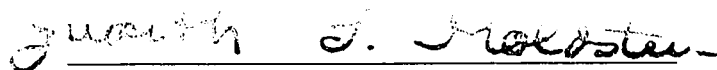
Professor J. David Powell

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



Professor Daniel B. DeBra

Approved for the University Committee
on Graduate Studies:



Abstract

The Global Positioning System (GPS) is poised to revolutionize aerospace Guidance, Navigation, and Control. While traditionally thought of as a navigation system, *GPS has now been demonstrated to be a formidable sensor for attitude determination* as well. In fact, GPS may now be considered a 10 state sensor, providing continuous real-time Position, Velocity, Time, and Attitude (PVTa) for all axes.

Attitude determination is derived from extremely precise (5mm, 1σ) differential range measurements among multiple GPS antennas mounted on a flight vehicle. While physics does not set the ultimate achievable limits of performance, *experimental testing is demonstrating pointing accuracies below 0.1 deg (1σ) with a usable bandwidth of several tens of Hz.*

This dissertation describes the design and development of a new and unique GPS receiver with full PVTa capability. It chronicles the key innovations which have made this technology possible and addresses the fundamentals and performance of GPS-based attitude determination. A prototype receiver based on *antenna multiplexing* is presented which offers prospects of significantly reducing the size, weight, power, and cost of the hardware. Other conceptual designs are presented for achieving the highest possible

performance. Methods of resolving the cycle ambiguities of the GPS carrier based on motion are developed and tested in an effort to drive the performance of attitude determination as close as possible to perfection.

Maiden space flights of the new receiver are scheduled for launch into low Earth orbit. Test data is presented from extensive flights on a single-engine Piper Dakota that has been used to explore aircraft applications of attitude determination. Among these, attitude and precise position sensing using GPS can be applied to the estimation of the aircraft dynamic model in flight. The wealth of dynamic response data obtained through GPS can be used to estimate aircraft stability derivatives and, surprisingly, aeroelastic effects with millimeter-level clarity.

Acknowledgments

Special thanks go to my advisor, Professor Bradford W. Parkinson, not only for providing me the freedom and latitude to carry out this work but for being a mentor and source of inspiration through his own great accomplishments, which I deeply admire and respect.

Professor J. David Powell is gratefully acknowledged for providing his Piper Dakota, skillful piloting, and technical insight and vision for the flight tests presented herein. His sound advice, clarity of teaching, and professional enthusiasm are always appreciated.

For its constructive criticism of this manuscript, I wish to thank my reading committee, including Professor Daniel B. DeBra whose intellect and quick wit are matched only by his transcontinental physical stamina. I would also like to extend my appreciation to my oral examination committee, including Professors Ronald N. Bracewell and Gene H. Golub, whose teaching I have greatly valued and enjoyed.

The Department of Aeronautics and Astronautics, Stanford University, its Faculty, Staff, and my fellow students and friends are gratefully acknowledged for establishing the most stimulating environment imaginable for the pursuit of graduate studies.

In developing the receiver graphical display and preparing for the flight tests, the efforts of Stu Cobb are appreciated enormously.

The research presented herein was made possible by the National Aeronautics and Space Administration under Contract NAS8-36125.

The many contributions of Trimble Navigation, Ltd. are enormously appreciated. It has been a privilege to work with the bright, enthusiastic, and energetic individuals who comprise the company, and I have a thorough respect for their capabilities, innovations, and achievements. On a tangible level, I greatly appreciate the company's donation of receiver hardware and the opportunity to build upon the foundations it has laid. But on a much higher plane, the experience has been a tremendous source of inspiration, tempered with fundamental insight into the triumphs and trials of entrepreneurship.

I wish to give special thanks to the United States Air Force for making possible the first test flights of the new receiver in space. The thoughtful efforts of many individuals from the Aerospace Corporation, Defense Systems, Incorporated, Onizuka Air Force Base, and Orbital Sciences Corporation have all contributed to making these flights a reality.

I owe my parents, Robert and Carolyn, a great personal debt for the investment they have made in me. This dissertation is dedicated to them. Of my soul mate Eliza Knox, I shall always adore her friendship and humor.

Table of Contents

Chapter 1: Introduction and Overview	1
Definition of the Performance Envelope.....	1
Antenna Multiplexing GPS Receiver Design.....	1
Matrix Methods for Cycle Ambiguity Resolution	1
Generalization of Wahba's Problem	2
Multipath Calibration	2
Estimation of Aircraft Dynamics and Aeroelasticity	2
1.1 Overview	2
1.2 History of Attitude Determination Using GPS	3
1.3 Performance of Attitude Determination using GPS	4
1.3.1 Principal Error Sources	4
Multipath	4
Troposphere.....	5
Carrier to Noise Ratio	5
Receiver-Specific Errors	6
1.3.2 The Performance Envelope	6
1.4 Development of an Attitude Receiver for Aerospace	9
System Architecture	10

Top-Level Processing Design	13
Chapter 2: Receiver Hardware and Signal Processing.....	15
Using Multiple, Independent Receivers.....	16
Use of a Common Reference Oscillator.....	16
The Common Local Oscillator	16
Antenna Multiplexing	18
2.1 Architectural Foundations of the Receiver.....	18
2.1.1 GPS Signal Structure.....	20
2.1.2 Carrier Phase Tracking.....	21
2.1.3 Differential Phase Signal Processing	23
2.2 The Continuous (Non-Multiplexed) Receiver Design	24
2.2.1 Crosstalk.....	24
2.2.2 Line Bias	27
2.2.3 State Space Tracking Loop Implementation	27
2.3 The Antenna Multiplexing Receiver	34
2.3.1 Fundamental Implementation.....	35
2.3.2 Impact Of Oscillator Stability	36
2.3.3 Selection Of Switch Rate As A Design Parameter	39
Receiver Time And Received Signal Time.....	41
Detecting Bit Transitions	43
Reading Data	44
2.3.4 Gain Selection	45
2.2.5 Absolute Carrier Phase Tracking	46
Analysis of Unmodeled Timing Errors	48
2.3.6 Half Cycle Ambiguity from Data Modulation	49
2.3.7 Aliasing	49

Sample Period Aliasing.....	50
Multiplexing Period Aliasing	51
Disturbance Aliasing.....	53
2.3.8 Acquisition	56
2.4 Considerations for Launch and Low Earth Orbit.....	60
Oscillator Vibration and g Sensitivity.....	60
Time Tagging	61
User Spacecraft Almanac for Search and Acquisition	61
Steady-State Acceleration	63
Code Phase Bias Rate.....	63
2.5 Carrier Phase on a Spinning/Oscillating Platform	63
Chapter 3: Cycle Ambiguity Resolution	67
Baseline Length Constraint.....	68
Integer Searches	68
Motion-Based Methods.....	70
3.1 Platform Motion	71
3.1.1 Displacement Vectors	72
3.1.2 Computing the Initial Guess for Large Angle Motion	73
3.1.3 Small Angle Motion: The “Eigenaxis” Method.....	78
Coplanar Baseline Case.....	83
3.2 Refining the Initial Guess.....	83
3.2.1 Estimating Platform Structural Flexibility	84
3.2.2 Linearized Observation Matrix	85
3.2.3 Sparse Cholesky Factorization	87
3.3 Static Platform Resolution via GPS Satellite Motion	88
Initial Guess.....	89

Iterative Refinement.....	89
3.4 Quasi-Static Integer Resolution	90
3.4.1 Cycle Ambiguity Resolution for Spacecraft	91
Initial Guess.....	91
Solution Refinement.....	92
3.5 Fourier Cycle Resolution Method	93
3.6 Self-Survey.....	96
Sequential Estimation.....	96
Line Bias Estimation	97
Platform Attitude During Self-Survey	98
Chapter 4: Operational Attitude Solutions	99
4.1 Non-Linear Least-Squares Solution	99
Estimation of Aircraft Wing Flexure	101
Angular Velocity Estimation.....	102
4.2 Wahba's Problem	103
SVD Solution	106
Quaternion Solution	107
The Rank of the B Matrix	108
Fast Solution Results.....	109
4.3 Integrity Checking.....	110
Setting INTEGER_VALID	110
Clearing INTEGER_VALID.....	111
Chapter 5: Antennas and Multipath	113
The Infinite Ground Plane Experiment	114
5.1 Multipath Calibration for Small Satellites	115

5.1.1 Theory Of Calibration	116
A Model of Multipath	117
Spherical Harmonic Model	118
5.1.2 Experimental Results.....	118
Experimental Antenna Configuration	119
GPS Satellite Sky Coverage.....	120
Experimental Data.....	121
Extrapolation of Results to Flight Vehicle Performance	126
5.1.3 Practical Implications.....	126
5.2 Antenna Pattern Shaping.....	127
5.2.1 Helical Antenna Experiment.....	129
5.2.2 Screen Wire Fence Pattern Shaping Experiment	129
5.2.3 Antenna Pattern Shaping through Phased Arrays	131
5.2.4 High Performance Phased Array.....	132
5.3 Increasing Baseline Length.....	134
5.4 Multipath Filtering	135
5.5 Delay Lock Architecture	135
Chapter 6: Flight Experiments	136
6.1 Ground Validation Tests	136
6.1.1 Static Accuracy Tests with a Theodolite.....	136
6.1.2 Preflight Trials with Pivoting Steel Roof Array.....	139
6.1.3 Testing Cycle Ambiguity Resolution for Space.....	139
6.2 Aircraft Experiments.....	141
6.2.1 Piper Dakota Flight Tests.....	141
Concurrent Kinematic Positioning.....	145
Aircraft Natural Modes	145

6.2.2 Dynamic Testing with Inertial Measurement Units	150
6.3 Spacecraft Experiments.....	151
APEX	151
RADCAL	154
Chapter 7: Applications of GPS Attitude.....	156
7.1 Aircraft Applications.....	156
Heading and Attitude Sensing.....	157
Estimation of Winds Aloft/Wind Shear Detection.....	157
Precision Landing.....	158
System Identification and Autopilot Synthesis	158
Formation Flying.....	159
7.2 Spacecraft Applications.....	159
Maneuvering Vehicles.....	159
Adjunct Attitude Sensing	160
Small Satellites.....	160
Space Vehicle Flexibility	160
Launch Vehicle Guidance	160
Precision Orbit Determination	161
Rendezvous and Docking.....	161
The Distributed Satellite.....	162
Chapter 8: Conclusion.....	164
Suggestions for Future Work	166
Epilogue	167
Appendix A: Measurement Time Tags	168

Appendix B: User Spacecraft Almanac	170
Appendix C: Kinematic Carrier Phase Positioning.....	173
References	175

List of Tables

6.2.1.1	Piper Dakota Self-Survey Results	143
6.2.1.2	Longitudinal Modes	148

List of Figures

1.1.1	Measuring Differential Range.....	2
1.3.2.1	Differential Ranging Error	7
1.3.2.2	The Performance Envelope of GPS-Based Attitude Determination	8
1.4.1	Attitude Determination Hardware.....	10
1.4.2	Receiver and Antenna Housing.....	11
1.4.3	GPS Receiver/Differential Phase Sensor Board.....	12
1.4.4	Top Level Processing Architecture.....	13
2.1	The Common Local Oscillator (Non-Multiplexing Design).....	17
2.1.1	Trimble GPS Receiver Architecture	19
2.1.1.1	The PRN Code of the GPS Signal.....	20
2.1.2.1	I and Q correlator sampling.....	22
2.1.3.1	Coherent Detector Architecture	23
2.1.3.2	Coherent phase detector	24
2.2.1.1	Splitting and Buffering the LO.....	25
2.2.1.2	Crosstalk Measurement.....	25
2.2.1.3	Unacceptable (5mm range error) Crosstalk	26
2.2.3.1	Modeling the Correlator as a Differencing Block.....	29

2.2.3.2	Arctangent Look-up Table Construction.....	30
2.2.3.3	Differential Phase Signal Processing Loop Architecture.....	32
2.2.3.4	Differential Phase Tracking	34
2.3.1.1	Antenna Multiplexing Top-level Design.....	35
2.3.2.1	Stability Ranges of Various Frequency Standards.....	37
2.3.3.1	Frequency Plan of GPS Signal and MUX Receiver.....	38
2.3.3.2	Detecting Bit Transitions During Multiplexing	40
2.3.3.3	Antenna switch timing	42
2.3.3.4	Data bit timing with respect to antenna switching	43
2.3.3.5	Phasing antenna switching to catch bit transitions (4 Antennas).....	44
2.3.7.1	Three Types of Aliasing Encountered in Multiplexing (4 Antennas).....	50
2.3.7.2	Aliasing Detection.....	52
2.3.7.3	Example of Disturbance Aliasing Mode	54
2.3.8.1	Signal Processing State Transition Diagram.....	57
2.3.8.2	Carrier Loop Lock Acquisition	59
2.4.1	Ephemeris Feed-Forward	62
2.5.1	Geometry of a Spinning Platform	64
2.5.2	Harmonic Correlator Signal Processing for Rotor-Mounted Antennas	65
3.1	Cycle Ambiguity Resolution.....	67
3.2	Baseline Length Constraint.....	68
3.3	The Integer Search	69
3.1.1.1	Differential Phase Geometry.....	72
3.1.2.1	Large Angle Motion Initial Guess.....	74
3.1.2.2	Non-Coplanar Baseline Rotation	77
3.1.3.1	Small Angle Motion.....	79
3.2.1.1	Flex Matrix Definitions.....	84

4.2.1	Examples of “balanced” antenna configurations	108
4.2.2	Solution benchmark results	109
5.1	Geometry of the Near-Infinite Ground Plane.....	114
5.2	SV Setting over Still Water	114
5.1.1	Multipath Error and Multipath Repeatability	115
5.1.2.1	Multipath Calibration Experimental Setup	119
5.1.2.2	Tracks of GPS Satellites Across the Celestial Hemisphere	120
5.1.2.3	Repeatability of Residuals among Different Space Vehicles.....	121
5.1.2.4	Best Fit Multipath Calibration (8th Order Model)	122
5.1.2.5	Calibration Results with Reflector	124
5.1.2.6	Sky Map of Phase Residuals	124
5.1.2.7	RMS Error Remaining after Calibration	125
5.2.1	Antenna Pattern Shaping.....	127
5.2.2	Multipath Phasor Interpretation	128
5.2.1.1	Directivity and Multipath	129
5.2.2.1	Antenna Pattern Shaping Experiment	129
5.2.2.2	Screen Wire Mesh Experiment	130
5.2.3.1	Pattern Shaping Antenna Array.....	131
5.2.3.2	Antenna Pattern Shaping Specifications	132
5.2.4.1	Composite-Embedded GPS Antennas: “Smart Skin”	133
5.2.4.2	High Performance GPS Attitude	134
6.1.1.1	Steel Rooftop Array	137
6.1.1.2	Steel Rooftop Array (Photograph)	137
6.1.1.3	Static Pointing Accuracy Test Results	138

6.1.2.1	Displacement Vectors in Horizontal Plane	138
6.1.3.1	Orthogonal Baseline Configuration	139
6.1.3.2	Raw, Unprocessed Phase Measurements	140
6.1.3.3	Integer Solution History	141
6.2.1.1	Piper Dakota with GPS	142
6.2.1.2	GPS Antenna on Right Wing	142
6.2.1.3	Aircraft Geometry	144
6.2.1.4	Landing at Half Moon Bay.....	144
6.2.1.5	Short Period Mode Response	147
6.2.1.6	Piper Dakota Phugoid Mode State Estimates.....	147
6.2.1.7	Phugoid Mode Relative Position.....	148
6.2.1.8	Phasor Diagram of Phugoid	149
6.2.1.9	Dutch Roll Mode Response.....	150
6.2.1.8	Phasor Diagram of Phugoid	149
6.3.1	APEX Spacecraft, Deployed Configuration.....	152
6.3.2	RADCAL Satellite Configuration.....	153
6.3.3	Top View of RADCAL Satellite.....	153
7.2.1	The Distributed Solar Power Satellite.....	163
A.1	Raw Measurement Time Tags.....	169

Chapter 1:

Introduction and Overview

Determining the attitude (orientation) of a flight vehicle is a fundamental aspect of aerospace Guidance and Control. The goal of this work is to utilize the precise positioning potential of the Global Positioning System (GPS) for attitude determination. The specific research contributions presented in this dissertation are as follows:

Definition of the Performance Envelope

The *Performance Envelope* of attitude determination using GPS is defined, such that a user can easily determine the expected pointing accuracy given the specific parameters of the flight vehicle [Chapter 1].

Antenna Multiplexing GPS Receiver Design

Reducing the size, weight, power, and cost of an attitude determination receiver by at least a factor of four, the innovation of antenna multiplexing is developed [Chapter 2].

Matrix Methods for Cycle Ambiguity Resolution

Driving the reliability of cycle ambiguity resolution as close to perfection as possible, new matrix methods are presented based on the geometry of motion [Chapter 3].

Generalization of Wahba's Problem

Reducing the computation required for an attitude solution by nearly an order of magnitude, efficient solutions to *Wahba's Problem* (attitude determination using vector observations) are generalized to attitude determination using GPS [Chapter 4].

Multipath Calibration

Improving the angular pointing accuracy by a factor of 10 or more for small satellites, a proposed method of *Multipath Calibration* is demonstrated experimentally [Chapter 5].

Estimation of Aircraft Dynamics and Aeroelasticity Using GPS

As a means of enhancing the overall safety and efficiency of aircraft operation, a series of experimental flight tests are carried out using GPS to estimate the aircraft dynamic model and aeroelastic effects [Chapter 6].

1.1 Overview

Relative range, Δr , to a GPS satellite between a pair of antennas (such as those shown in Figure 1.1.1) can be inferred by measuring the GPS carrier phase through each antenna. Using three or more antennas mounted on a flight vehicle, the GPS receiver developed herein translates these precise differential phase measurements to two or more GPS satellites into attitude estimates in real time. Relative range measurements using the GPS carrier at 1575 MHz have a precision of about 5mm RMS. Flight testing has demonstrated attitude accuracies better than 0.1 deg at a 10Hz output rate.




Figure 1.1.1: Measuring Differential Range

The scope of this dissertation has both theoretical and experimental dimensions. In the direction of theory, the ultimate performance limits of attitude determination using GPS are studied, with special emphasis on mathematical techniques in the areas of signal processing, cycle ambiguity resolution, and attitude solutions. In the experimental direction, a prototype attitude determination receiver is designed, constructed, and flight tested. The experimental implementation establishes the reliability and efficiency of the mathematical and geometrical foundations underlying the receiver design. Furthermore, the prototype serves as a test bed, examining the applications and ultimate capabilities of attitude determination using GPS.

1.2 History of Attitude Determination Using GPS

Attitude determination using GPS was originally proposed by Spinney¹ in 1976. During the early 1980s, Greenspan et al.² and others helped clarify the virtues and techniques of precise ranging using GPS carrier phase. Concurrently, Brown, Bowles, and Thorvaldsen proposed that carrier phase be applied to attitude determination.³ The first published attempt to build such a receiver was in 1983, when Joseph and Deem⁴ performed static tests with a single-baseline receiver.

In 1988, Kruczynski et al.⁵ performed the first three-axis, real-time test of heading and attitude sensing in an experiment on the U.S.S. Yorktown. The first test of attitude determination on an aircraft was performed by Purcell et al. from JPL,⁶ who used a pair of independent receivers whose antennas were mounted along the fuselage of a DC-8 to measure aircraft heading in post-processing. In 1991 van Graas and Braasch⁷ from Ohio University performed the first three-axis, real-time flight test on a DC-3. The receiver was a 24-channel, double-difference unit with a 1Hz output rate built by Ferguson et al.⁸ No space-based tests are known to have occurred as of this writing.

1.3 Performance of Attitude Determination using GPS

With a view toward ultimate performance, this section outlines the principal error sources of attitude determination using GPS and weights their relative importance. Next, the most significant error sources, including multipath, are combined together to form the *Performance Envelope* of attitude determination using GPS.

Could GPS-based attitude determination acquire the unique distinction of becoming the highest performance regime of GPS operation in existence? Not only does it have the capacity to achieve higher relative position accuracy than GPS surveying, it has a bandwidth potential approaching that of inertial instruments. As a complete, self-contained package providing differential positioning capability, attitude determination using GPS offers a unique test bed for exploring the error sources and ultimate limits of performance intrinsic to GPS itself.

1.3.1 PRINCIPAL ERROR SOURCES

The principal error sources in attitude determination are multipath, received carrier-to-noise ratio, bending of the received ray by atmospheric refraction (troposphere), and receiver systematic errors. Selective Availability (SA), GPS satellite clock errors, and correlated atmospheric errors cancel out in the differential operation of a receiver.

Multipath

By far the largest error source in attitude determination is multipath, the unwanted reflections of the GPS carrier from the environment surrounding the antenna (see Chapter 5). To quantify the magnitude of differential positioning error due to multipath, it is clear that one must consider the specific antenna installation environment.

A realistic and useful guideline is the case of a pair of microstrip patch antennas. Each antenna has an hemispherical pattern and an open view of the sky. Each is mounted without a ground plane in the vicinity of an abundance of reflective objects. For this example, the differential ranging multipath error for L1 is about 5mm RMS.

Since multipath is strongly dependent on geometry, the spectral color of multipath error also depends on the time scale of vehicle dynamics and the antenna separation.

Troposphere

Tropospheric refraction can cause the apparent direction of origin of a wave front to differ by more than a tenth of a degree for satellite elevations below 10 deg. This error source is easily modeled out for the relatively short baselines (<30m) under consideration for attitude determination. Based on the estimated index of refraction of air, Snell's law is used to apply a correction to the calculated line-of-sight vectors to the GPS satellites.

Carrier to Noise Ratio

The white noise on the reconstructed carrier dictates the differential range measurement error,⁹ as given in the white noise equation

$$\sigma_r = \sqrt{\frac{f_N}{C/N_0}} \frac{\lambda}{2\pi}$$

where f_N is the noise bandwidth of the carrier tracking loop, C/N_0 is the carrier to noise ratio (40 dB-Hz is assumed in the following section), and λ is the GPS L1 wavelength. If it were not for multipath, the received carrier-to-noise ratio would dominate the error budget for attitude determination.

Receiver-Specific Errors

Receiver-specific errors, including cross-talk, line bias, and inter-channel bias can be significant sources of error if they are not treated appropriately in a receiver design.

Cross talk between the R.F. paths for each antenna is an issue because there is often more than 100 dB of gain along each path.

Line bias is the nearly constant offset in phase from one antenna to another. A function of both cable length and temperature, line bias is all that remains of the relative clock offset in the common local oscillator design (see Chapter 2).

Finally, inter-channel bias is a by-product of the fact that different hardware channels are used to measure the carrier phase for each satellite. Note that the multiplexing receiver design presented in Chapter 2 effectively eliminates most of these systematic effects.

1.3.2 THE PERFORMANCE ENVELOPE

To characterize the accuracy of attitude determination using GPS, it is useful to define its *Performance Envelope*,¹⁰ or theoretical limit of performance. Such a performance envelope is most appropriately described in terms of accuracy and bandwidth. For a given tracking loop noise bandwidth (set to be a few octaves above the highest important characteristic frequency of platform attitude motion) and baseline length (separation between antennas), the envelope describes what pointing accuracy can be achieved.

The two limiting error sources, carrier-to-noise ratio and multipath, are compiled into a single set of curves which collectively define the performance envelope of GPS-based attitude determination. Each of these error sources is first treated in terms of carrier

tracking ranging error, σ_r . Multipath error is modeled conservatively as a constant differential range error of 5mm, which is independent of tracking loop bandwidth.

These two error sources can be combined (RSS) to provide an estimate of the total differential phase ranging error. This conservative estimate does not take into account the improvement in accuracy from multiple measurements taken with additional baselines and satellites. Figure 1.3.2.1 shows this nominal performance curve along with representative multipath and C/N_0 lines. Usually, either one or the other principal error source dominates. Note that a practical limit of bandwidth is imposed by the carrier loop lock threshold. At the bandwidth where the total ranging error crosses over that limit, the accuracy of the phase measurements is judged too poor to hold lock effectively.

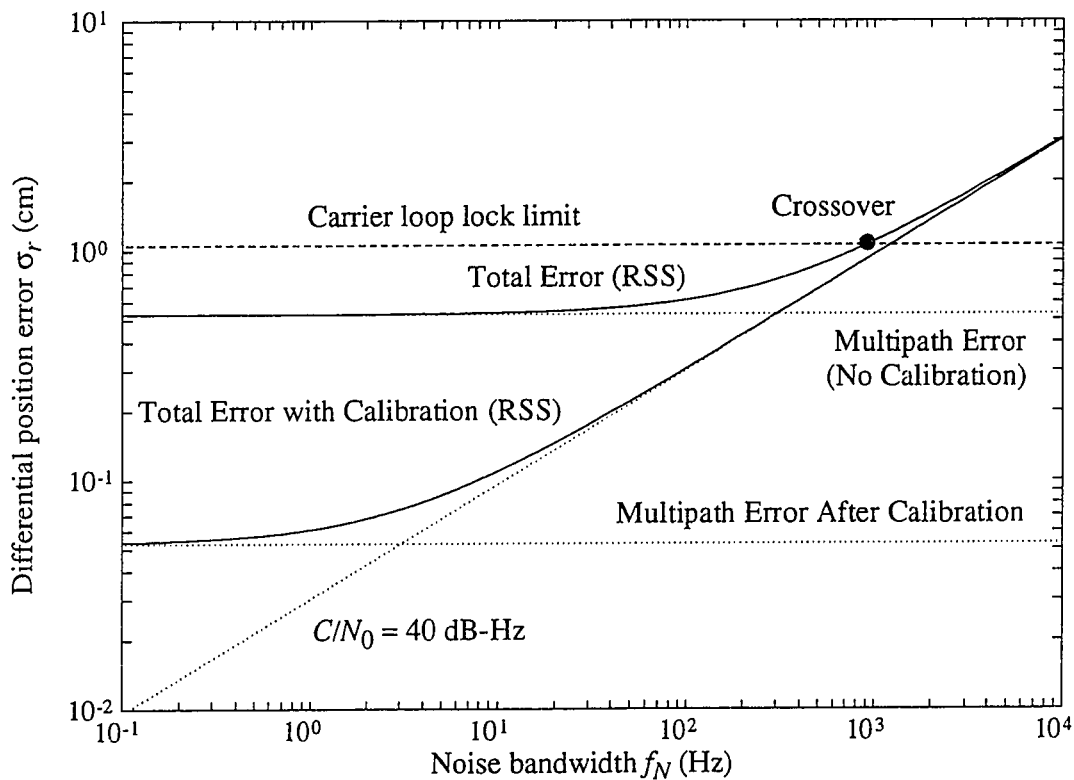


Figure 1.3.2.1: Differential Ranging Error




Figure 1.3.2.2: The Performance Envelope of GPS-Based Attitude Determination

A second curve indicates the effect of antenna pattern shaping or multipath calibration (see Chapter 5) on differential range measurement accuracy. If multipath error is reduced by an order of magnitude to 0.5 mm RMS, the carrier-to-noise ratio dominates the error budget starting at a lower noise bandwidth, resulting in superior accuracy.

The set of curves shown in Figure 1.3.2.2 comprise the performance envelope of GPS-based attitude determination. The resulting angular pointing error, σ_θ , is given in terms of the relative ranging error due to multipath, σ_r , and the baseline length, L , as

$$\sigma_\theta = \frac{\sigma_r}{L}$$

The curves in Figure 1.3.2.2 are identical to those shown in the previous figure, except they are divided by the baseline length, L , to yield the pointing error. Pointing accuracy improves considerably with increasing antenna separation.

1.4 Development of an Attitude Receiver for Aerospace

For the development of the high precision differential phase attitude receiver chronicled herein, extensive modifications were made to an off-the-shelf Trimble Advanced Navigation Sensor (TANS). The TANS is a small, ruggedized 6-channel GPS receiver designed for high vibration and extreme temperature environments.

Although starting from this foundation was certainly appropriate in view of the potential test flights envisioned for the attitude receiver, the original, unmodified TANS was never designed to make carrier phase measurements. Another possible approach would have been to start with a receiver that had originally been designed as a carrier phase sensor for static surveying applications, modify it extensively for real-time attitude determination,

and then ruggedize it for use in harsh flight environments. Unfortunately, the survey-class hardware is significantly more complex. There was little desire to pursue this approach because it was felt that for attitude determination using GPS to be truly viable, it had to be based on the simplest possible hardware to promote the highest possible reliability and, in the long run, lowest possible cost.

Only the single-frequency, C/A code signal component of GPS was employed to demonstrate the performance capability of attitude. This approach ensures that the results are completely general, given the current uncertainty as to the future status of the P code and dual-frequency availability.

System Architecture

The system architecture is presented in Figure 1.4.1. It consists of four antennas (with built-in pre-amplifiers), the GPS receiver/differential phase sensor (enclosed by the




Figure 1.4.1: Attitude Determination Hardware

dotted rectangle), and a 486-based PC for attitude calculations and display. The original TANS architecture is very similar to that shown here with the exception of the PIN diode switching circuit¹¹ (designed at Trimble based on specifications provided) and the additional antennas.

Commanded by the receiver computer, the switcher selects between any one of the four antennas connected to the receiver and routes its output into the R.F. section. In the R.F. section, the signal is amplified, down-converted, and digitized. A bank of custom digital signal processing channels is used to track up to six GPS satellites simultaneously. Each channel is directly readable and commandable by the 68000 receiver computer. This computer closes all receiver tracking loops, calculates position fixes, and routes raw






Figure 1.4.2: Receiver and Antenna Housing



Digital Side



Analog Side

Figure 1.4.3: GPS Receiver/Differential Phase Sensor Board

differential phase measurements to the 486 PC. The 486 PC then performs the cycle ambiguity resolution and attitude solutions.

Figure 1.4.2 shows a photograph of the receiver housing and hemispherical GPS microstrip patch antenna. The pre-amplifier is built into the antenna housing and biased by the receiver unit through the 50 Ω coaxial cable connecting the two. The receiver housing has four SMA connectors for the antennas along with power and data connectors.

The GPS receiver/differential phase sensor fits on the single board shown in Figure 1.4.3. On the analog side, the TCXO is visible. This oscillator serves as the frequency reference for the entire receiver. Also visible is the switching circuit, mounted directly to the R.F. section. On the digital side, the six receiver channels reside in two custom digital ASICs. The rest of the digital side of the board is filled with computer components, including the 68000 microprocessor, RAM, ROM, and DUART.

Top-Level Processing Design

The processing scheme in the attitude determination system is shown in Figure 1.4.4. A bank of signal processing loops (see Chapter 2) keeps track of differential phase with a separate process running for each combination of antennas and satellites. The integer estimator block uses a variety of processing schemes (see Chapter 3) to resolve the cycle ambiguities. This block is invoked at system power-up or if integer lock is ever subsequently lost. Once the integer ambiguities

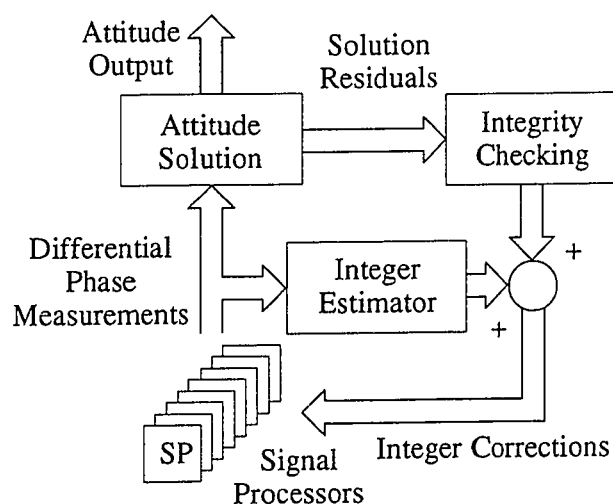


Figure 1.4.4: Top-Level Processing Architecture

are resolved, the attitude solution block (see Chapter 4) converts differential range measurements into attitude solutions. The integrity checking block (also in Chapter 4) combs the solutions to detect and isolate potential cycle slips. If a cycle slip is found, this block automatically initiates corrections by adjusting the integer value of the offending measurement(s).

Chapter 2:

Receiver Hardware and Signal Processing

The foundation of high performance attitude determination using GPS is the common oscillator (LO). Shared among all the antennas on the platform, this common oscillator serves as the reference for every phase measurement made by the receiver. Receiver oscillator noise is common-mode and cancels out when the phase measurements from each antenna are differenced. The result is exceptionally clean measurements.

This chapter discusses the hardware and signal processing approaches used to obtain the precise differential phase measurements on which attitude solutions are based. The hardware discussion focuses on several possible implementations of the common local oscillator, all of which were evaluated experimentally. The most economical approach of those considered is antenna multiplexing, which is described in greatest detail. The signal processing discussion covers the state space implementation of the carrier tracking loops developed specifically for the receiver.

Since signal processing is closely coupled with the receiver hardware itself, these two areas are addressed concurrently in this chapter.

In constructing the differential phase receiver, the level of integration of the separate R.F. paths is fundamentally important. A brief overview is presented of receiver design approaches in order of increasing levels of integration.

Using Multiple, Independent Receivers

Certainly the minimal level of integration is to run two or more survey-quality receivers together, each running on completely independent oscillators. Several experiments on attitude determination have been performed in this manner.^{12,13} The fundamental problem with this approach is that one additional satellite is always required to solve for the relative oscillator offset between a pair of independent receivers.

Use of a Common Reference Oscillator

The next higher level of integration is to run the two receivers from the same external reference oscillator, as has been done in several cases.^{8,14} This eliminates the additional measurement required to solve for the relative phase offset between the receivers which now stays roughly constant. However, since the reference oscillator usually drives additional circuitry (such as a phase-locked loop synthesizer) to form the local oscillator, the relative phase offset is apt to exhibit undesirable drift with temperature or supply voltage. Therefore, the perceived benefit from running from the same reference oscillator is not generally realized in practice.

The Common Local Oscillator

An even better design consists of splitting the local oscillator as it is distributed to the dedicated downconverter for each antenna. The resulting relative phase offset between R.F. paths is much less sensitive to environmental changes. This approach was pursued for the initial attitude determination experiments performed here. A conceptual layout of

the common local oscillator architecture (non-multiplexing design) is shown schematically in Figure 2.1.

From each antenna, the GPS signal (still modulated by the 1MHz C/A code) is mixed with the local oscillator and down-converted from 1575.42MHz to the intermediate frequency, 4MHz. The signal is band-pass filtered, quantized, and routed to a digital signal processor. A coherent detector then extracts the relative carrier phase for each satellite tracked among the different antennas.

Since the same local oscillator is distributed to each mixer in the R.F. path of each signal, the oscillator phase noise is common mode, canceling out perfectly in the differencing step between the multiple antennas. The result is an exceptionally clean differential phase measurement that is largely independent of the stability of the reference oscillator.

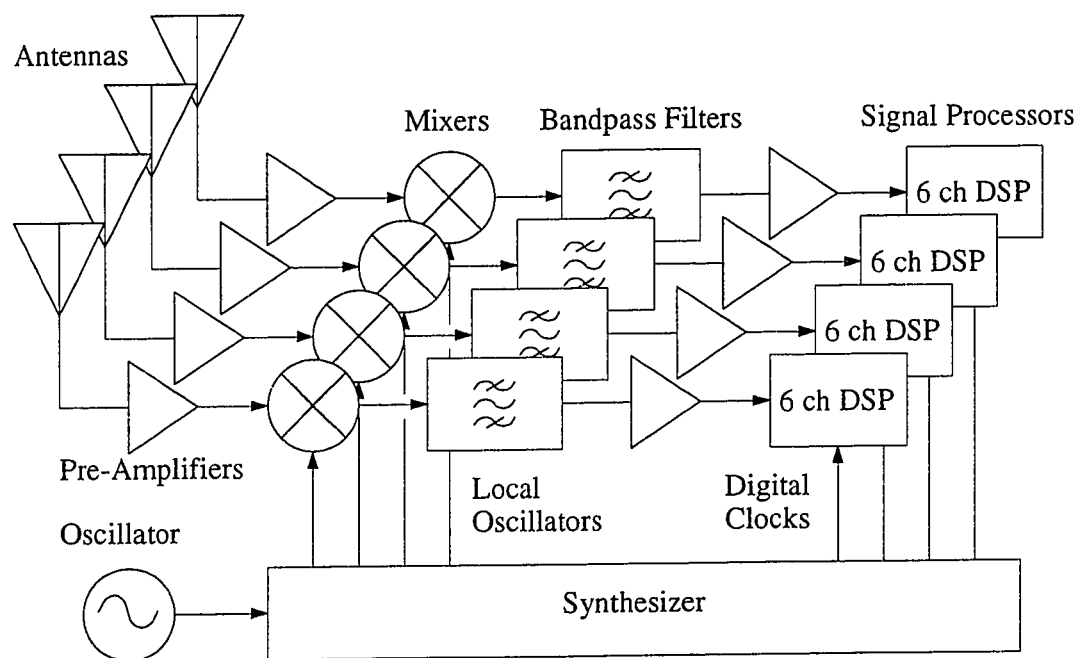


Figure 2.1: The Common Local Oscillator (Non-Multiplexing Design)

Antenna Multiplexing

The concept of multiplexing multiple antennas through a single GPS receiver¹⁵ (see Section 2.3) takes the level of integration to the ultimate limit. With this concept, a PIN diode or FET switch is used to feed the signals from separate antennas rapidly and sequentially through the same R.F. section. Two distinct advantages result from this approach. First, *the overall cost and complexity of the receiver is reduced to a minimum.* The entire receiver now draws only 3.5 watts of power and fits on the 5 by 7 inch circuit board shown in Figure 1.4.3. Second, line biases from one R.F. path to the next are—by definition—non-existent beyond the switching point in the processing chain. The differencing operation is based on signals which have traversed exactly the same signal path. *Line bias drift is effectively eliminated.*

2.1 Architectural Foundations of the Receiver

This section describes the existing state-of-the-art prior to the research described in this dissertation. The hardware that was subsequently developed is founded on the Trimble GPS receiver architecture.¹⁶ This architecture, based on an innovative and efficient one-bit signal quantization scheme, allows the bulk of signal processing to be performed in digital electronics.

As shown in Figure 2.1.1, the GPS signal received at the antenna is amplified and down-converted to an intermediate frequency of 4MHz. At this point it is amplified to the point that it saturates the output of the amplifier stages and hard-limits. The effect of this high gain is to quantize both the signal and noise at the input to the receiver into one of two states: 0 and 5 volts (or, alternatively, +1 and -1, or 0 and 1). Then, the limiter output may be sampled at a periodic interval. From here on, the receiver implementation is completely digital. All subsequent mixers, correlators, and oscillators can be rendered very efficiently in digital electronics.

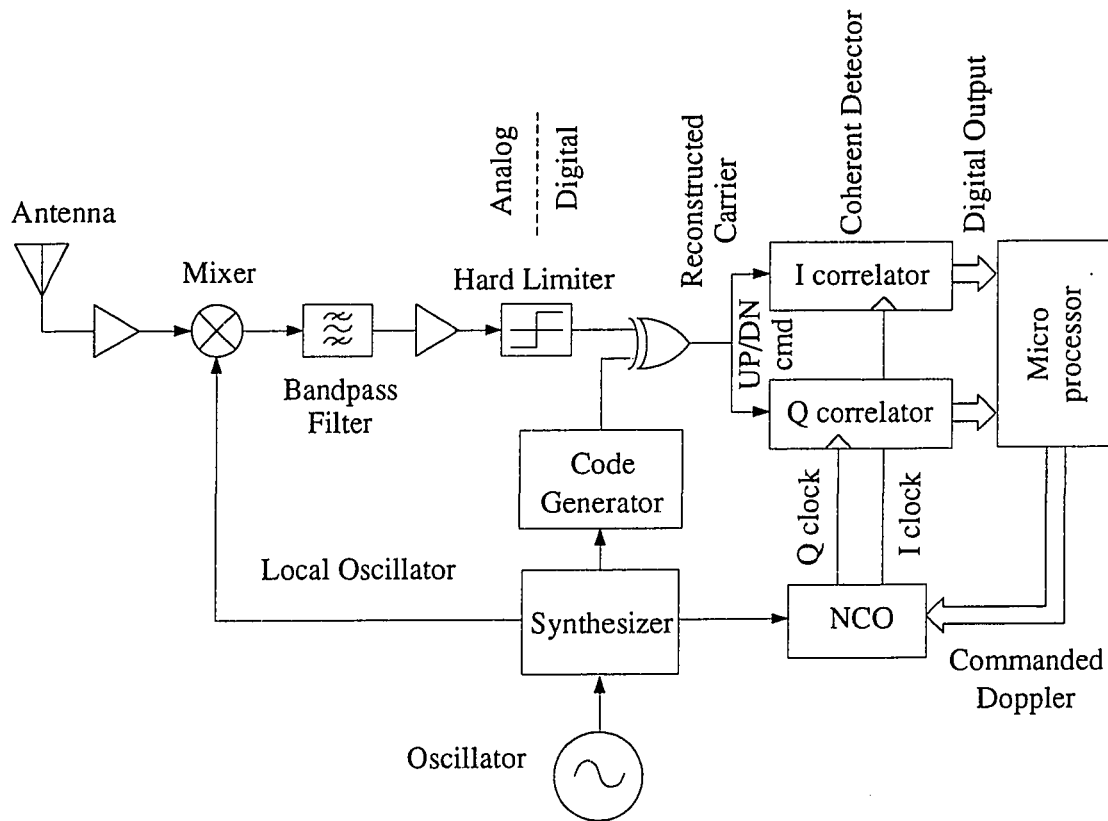


Figure 2.1.1: Trimble GPS Receiver Architecture

The digital correlator is implemented in the form of a digital up/down counter. The binary input to the correlator determines whether the counter increments or decrements when an in-phase or quadrature (I or Q) clock pulse is applied. After approximately 4,092 clock pulses, the value of the counter reads out the average value of the signal amplitude.

The carrier is effectively down-converted to baseband via sampling. This final sum or average value is latched into a holding register for readout, and the counter is reset for the next interval. A microprocessor reads the output of the correlators to measure the carrier phase tracking error. A new commanded frequency is computed by the control law and

latched into the Numerically Controlled Oscillator (NCO) to drive the carrier tracking error to zero for each channel.

The reason that single-bit quantization is viable is that the log signal-to-noise ratio of the GPS signal is negative. Spread-spectrum modulation is applied to the GPS signal in the form of a pseudo-random noise (PRN) digital code. The PRN modulation spreads the signal energy over a 1MHz bandwidth. Spilker¹⁷ shows that the power loss incurred by hard-limiting the PRN signal is a factor of $2/\pi$ or about 2 dB. This is not a bad trade for such an enormous reduction in receiver complexity.

2.1.1 GPS SIGNAL STRUCTURE

As described by Spilker¹⁸ the coarse/acquisition (C/A) code signal, $s(t)$, transmitted with amplitude A from each GPS satellite is a product of three components:

- (1) $c(t)$, the predefined PRN code at 1.023 MHz
- (2) $d(t)$, the 50 bits per second data message
- (3) $\sin(1540 f_0 2\pi t)$, the narrow-band GPS carrier

such that $s(t) = A c(t) d(t) \sin(1540 f_0 2\pi t)$. As shown in Figure 2.1.1.1, the values of $c(t)$ can be either ± 1 . The code changes value at a predefined frequency called the “chipping




Figure 2.1.1.1: The PRN Code of the GPS Signal

rate”, which, for C/A code, is set at 1.023 MHz. This sequence repeats continuously every millisecond (1023 chips). The parameter f_0 is a fundamental frequency defined to be 1.023 MHz in concert with the repetition period of 1023 chips.

When the PRN sequence is modulated onto the GPS carrier at the transmitter, it has the effect of spreading the signal power over a bandwidth roughly equal to the chipping rate. Since the product $c(t) \times c(t)$ is always unity, the original carrier can be recovered by multiplying the received signal by a replica of $c(t)$ generated by the receiver. In Figure 2.1.1, this mixing is performed by the exclusive-or gate. The code tracking loop has been omitted for clarity. The data modulation which remains—also set to be either ± 1 —can be stripped away easily because its characteristic frequency of 50 bps is slower than typical correlator integration times and faster than the vehicle kinematics.

Note that $c(t)$ is a unique, predetermined sequence for each individual satellite. The different codes are chosen such that they are nearly orthogonal. Therefore, the coded signals from any number of GPS satellites can share the same electrical path and frequency range. The carrier for any desired satellite can be reconstructed by mixing the received signal with the code for that specific satellite. The coded signals from the unwanted satellites still resemble broad band noise, which is filtered out in the correlation process.

2.1.2 CARRIER PHASE TRACKING

Once the carrier has been reconstructed, it is converted to baseband via sampling. The Numerically Controlled Oscillator (NCO) is adjusted to output both in-phase and quadrature clock edges with respect to the incoming carrier as shown in Figure 2.1.2.1. Accumulating and averaging roughly 4092 samples over a millisecond, the output of the I




Figure 2.1.2.1: I and Q correlator sampling

and Q correlators provides a nearly instantaneous measurement of the relative phase error between the NCO and the received carrier.

A standard representation of the phase error is to plot Q versus I as shown on the far right in Figure 2.1.2.1. The resulting locus of points forms a circle centered about the origin whose radius is proportional to the carrier-to-noise ratio (C/N_0) of the received signal. The angle of any point on the locus represents the phase tracking error with respect to the nominal phase, $I=C/N_0$, $Q=0$.

In nominal receiver operation, the NCO frequency is continuously regulated by the control law in the microprocessor to drive the phase tracking error to zero. The

commanded frequency is available as a measurement of the instantaneous Doppler shift of the satellite signal.

2.1.3 DIFFERENTIAL PHASE SIGNAL PROCESSING

As a foundation for differential phase measurements, Braisted and Eschenbach¹⁹ expanded this architecture to a coherent detection scheme. In this approach, a master and slave antenna share the same NCO, as shown in Figure 2.1.3.1. The NCO is controlled such that the phase error from the master antenna is driven to zero. In effect, this loop filters out the translational motion of the vehicle. Since the slave clocks are referenced with respect to the same NCO, the slave I and Q correlators yield the differential phase, as shown in Figure 2.1.3.2. In the initial design, the I and Q correlator output was averaged for half a second. Then, the

differential phase, $\Delta\phi$, was computed from the averaged samples:

$$\Delta\phi = \text{atan}_2(Q, I).$$

Kruczynski et al.⁵ applied this architecture to ship-board attitude determination. Their results were encouraging; however, they were still plagued with receiver crosstalk in the R.F. section, bandwidth

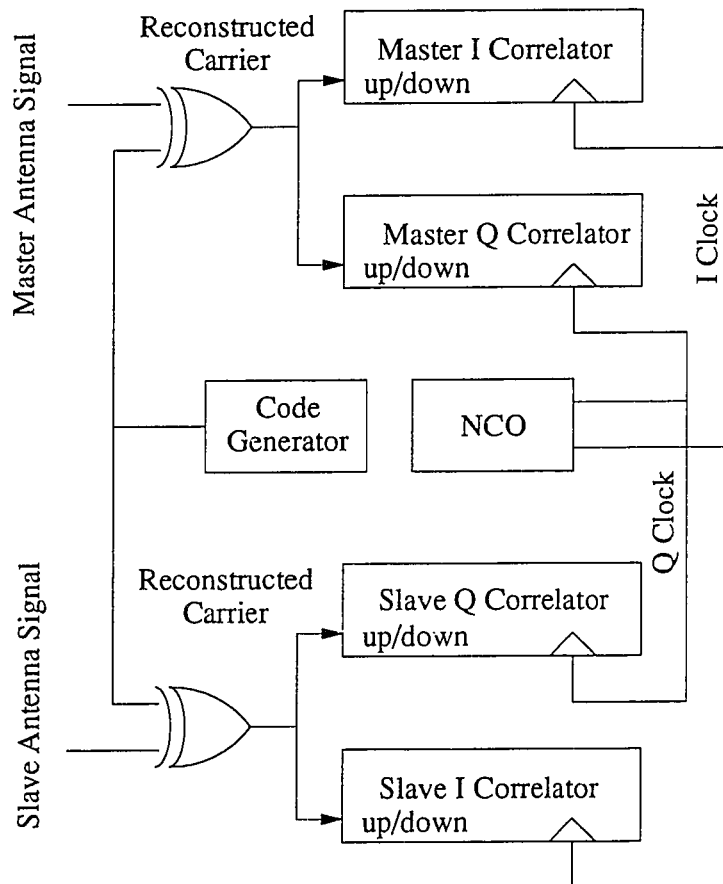


Figure 2.1.3.1: Coherent Detector Architecture




Figure 2.1.3.2: Coherent phase detector

limitations in tracking dynamics, and limitations in initially resolving and holding on to integer ambiguities.

2.2 The Continuous (Non-Multiplexed) Receiver Design

This section documents the efforts undertaken in this research to overcome the limitations in signal processing encountered in the previous work. The first step was to eliminate the R.F. crosstalk. Then, the signal processing techniques were redesigned to improve the performance of differential phase measurements in handling attitude rates and other kinematics.

2.2.1 CROSSTALK

In a modification of the shipboard receiver design, a single synthesizer was used for the entire receiver. The local oscillator was split before distribution to dedicated downconverters for each antenna. The resulting design is shown in Figure 2.2.1.1. The local oscillator is split 4 ways through a network of hybrids. Then it is amplified and buffered by broad band amplifiers.

Isolation between the R.F. sections is extremely important. The signal level of the GPS signal as received through a hemispherical antenna is specified to be -130 dBm. At the

receiver R.F. section, more than 100 dB of gain is required to bring the signal up to a usable level. Therefore, an equivalent level (>100 dB) of isolation is required between R.F. paths to keep crosstalk within an acceptable level. Unfortunately, there are many potential paths for crosstalk.

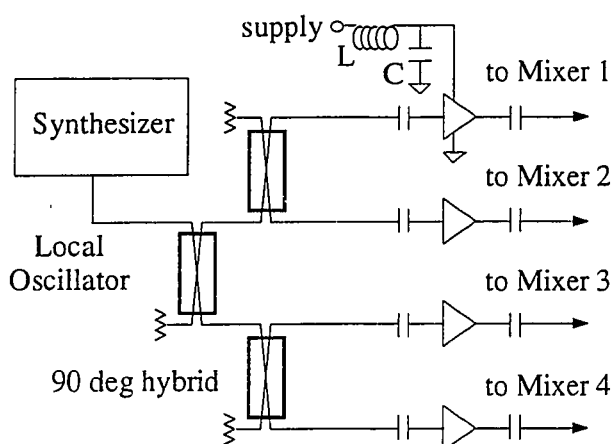
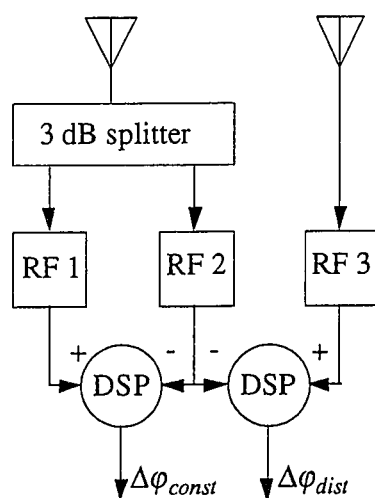


Figure 2.2.1.1: Splitting and Buffering the LO

Blatant crosstalk can be detected with a network analyzer. However, due to the large dynamic range of signal levels present in the receiver R.F. section, the most sensitive testing is done most practically with real satellite signals.

Crosstalk can be evaluated conveniently with a receiver that has at least three R.F. inputs. An appropriate configuration is shown in Figure 2.2.1.2. One antenna is connected to an R.F. splitter whose output is directed to two of the receiver inputs. Because the signal is



**Figure 2.2.1.2:
Crosstalk Measurement**

originating from the same source, the differential phase measured between this pair of inputs should always remain constant.

Next, the output of a second antenna is connected to the third R.F. input. Natural satellite motion will cause the differential phase between the two antennas to shift with time. The differential phase between the two antennas acts as a disturbance to the pair fed from the splitter. If there is any crosstalk from the third R.F. path to the

others, the measured differential phase between the first two R.F. sections will *not* stay constant.

Figure 2.2.1.3 shows an example of unacceptable crosstalk. In this figure, two distinct ovals are shown, with Q plotted against I for each differential phase measure-

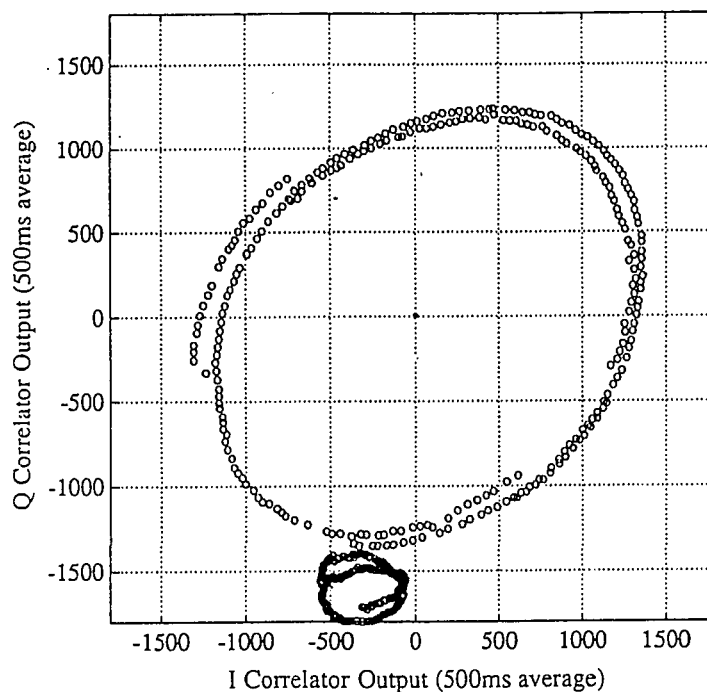


Figure 2.2.1.3: Unacceptable (5mm range error) Crosstalk

ment. The large oval is the disturbance phase changing at a constant rate. The second shows how the measured phase difference from the splitter is affected by the disturbance. Ideally, the phase difference between the first two antenna inputs should stay constant. Due to crosstalk, however, the magnitude and phase move in step with the disturbance. Since this plot is based on an amplitude scale, the effective isolation can be calculated by taking the area ratio of the two ovals. (In this case, the isolation is only about 14 dB, or 5mm of range error as given by half the angle subtended by the small oval from the origin.)

Careful attention to buffering and power supply filtering reduced the crosstalk level to a more reasonable level of 0.5 mm of relative ranging error, or 35 dB of isolation. The definition of “reasonable” comes from the total error budget. Some crosstalk is accepted as long as it does not dominate the error budget. An appropriate specification is to keep it very much smaller than the largest error source, multipath error.

2.2.2 LINE BIAS

Theoretically, the common LO offers a significant benefit in that the relative offset between measurements from two antennas (line bias) remains constant. This bias naturally varies with temperature due to thermal expansion in the antenna cables and electrical variations in the receiver. Based on an experiment in a temperature chamber, the receiver line bias temperature sensitivity of the common LO receiver was measured to be $190\mu\text{m/K}$. For a spacecraft passing out of the Earth's shadow, it is conceivable that centimeter-level changes in the line bias could result.

As a last resort, line bias can always be estimated continuously as a state parameter during normal operation at the expense of redundancy and overall solution accuracy. However, the most appealing solution is to minimize it in hardware to the point where it can be neglected altogether. Fortunately, changes in line bias are minimized by the differential nature of the receiver. Any variation is mostly common mode so that the difference cancels out. In fact, the multiplexing receiver described in Section 2.3 virtually eliminates the receiver component of line bias error. That which is attributable to thermal expansion of the antenna cables is minimized by keeping them as short as possible (or, at the very least, roughly the same length) and locating them where differential temperature changes are small.

2.2.3 STATE SPACE TRACKING LOOP IMPLEMENTATION

As a crucial enhancement to coherent phase detection, a phase-locked loop was closed around the differential phase measurements. Previously, I and Q samples were simply averaged over half a second to obtain a differential phase measurement. While this approach worked fine in the static case, it was inadequate in the presence of any relative velocity between the antennas.

Because of the sheer volume of digital data which must be processed to close a tracking loop and the limited resources of computing power in the microprocessor, the actual estimation and control software is best written in assembly code. The approach taken here was to develop an optimal estimator that executes with maximum speed and efficiency. A state space implementation with steady-state Kalman gains fits this problem beautifully for a number of reasons:

- *Speed.* For implementation in assembly code, the gain values can be selected to be a binary number which comes reasonably close to the optimal gain. The resulting estimator may then be implemented very efficiently in fixed point arithmetic, performing all math operations in binary via shifts and adds.
- *Convenience.* Each state variable is continuously accessible for readout.
- *Control.* Pole locations are less sensitive to binary quantization of state space gain values than are other methods of implementing digital control such as linear recurrence difference equations.
- *Adaptability.* Occasionally it is necessary to skip a measurement or two due to momentary antenna shading or noise. (*When both I and Q both lie near zero, the phase error becomes extremely sensitive!*) In this case, simple logic allows the estimator to skip a measurement update and “flywheel” through the transient.

In order to execute such a tracking loop with high throughput, it is necessary to implement the look-up table shown in Figure 2.2.3.1 as an efficient means of transforming raw I and Q measurements to phase angles. In the figure, the correlators




Figure 2.2.3.1: Modeling the Correlator as a Differencing Block

and look-up table can be considered collectively as a summing junction which outputs the numerical difference in phase between the incoming signal and the NCO.

The arctangent look-up table, shown in Figure 2.2.3.2, is constructed in an 8K block of ROM. Each cell in the table contains the phase angle, $\Delta\phi$, corresponding to the I and Q values along each coordinate axis. The 8-bit data byte is used to encode the value of the phase error such that the range 0-256 corresponds to an angular range of 0-360 deg or 0-1 cycle. The most natural way to represent carrier phase angles in fixed point binary arithmetic is in units of cycles. Therefore, the decimal place is positioned just to the left of the fractional phase from the look-up table.




Figure 2.2.3.2: Arctangent Look-up Table Construction

The locus of phase angles is not circular since the I and Q clocks are not exactly 90 deg out of phase. Therefore, the arctangent function also includes a correction for this built-in phase perturbation. The raw output words of the 13-bit I and Q correlators are truncated, leaving their 7 most significant bits. Since the mapping function is symmetrical, the most significant bit of the I coordinate can also be eliminated by reflecting I and Q about the origin. During normal operation, these two binary values are concatenated together to form the address in the look-up table.

The contents of the look-up table were calculated using a short MATLAB program based on the equation

$$\Delta\phi = \arctan_2\left(Q \sin \frac{5\pi}{12}, I - Q \cos \frac{5\pi}{12}\right)$$

The program output was scaled into integer values ranging from 0 to 255 and downloaded into the receiver ROM.

A second order loop implemented in state space was set up to estimate the differential phase, $\Delta\phi$, and differential phase rates, $\Delta\dot{\phi}$, for each satellite and baseline in track at each discrete time, k . This loop is shown schematically in Figure 2.2.3.3. The correlator is represented by the summing junction. The compensator $C(z)$ maintains lock on the translational component of signal dynamics, while the compensator $H(z)$ maintains lock on the differential range measurements.

Based on the correlator integration time, T , of 1ms, the state transition equation is

$$\begin{bmatrix} \Delta\phi \\ \Delta\dot{\phi} \end{bmatrix}_{k+1} = \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta\phi \\ \Delta\dot{\phi} \end{bmatrix}_k + \begin{bmatrix} \varepsilon \\ 1 \end{bmatrix} w_k$$

where a discrete zero mean disturbance, w_k , with variance W is applied to the velocity state. Since the correlator (the sensor) directly outputs the instantaneous (but noisy) differential phase angle, the observation equation is

$$z_k = \begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} \Delta\phi \\ \Delta\dot{\phi} \end{bmatrix}_k + v_k$$

subject to measurement noise v_k with variance V . The measurement noise variance V is known very well from the carrier-to-noise ratio (C/N_0) by the white noise equation:

$$1/V = C/N_0 T.$$




Figure 2.2.3.3: Differential Phase Signal Processing Loop Architecture

However, the character of the disturbance, w_k , depends strongly on the receiver environment, whether it be static or dynamic (mounted, for example, on a high performance jet). Given that the disturbance may even be correlated from one time to the next, the interpretation of W is open to debate. Given values of W and V , the optimal steady state gain, L , obtained through the algebraic Ricatti equation,²⁰ yields a transfer function from z to $\Delta\phi$ with a noise bandwidth, B . There is a direct mapping between V/W and B . Therefore, the tracking loops have been characterized not in terms of process noise, but in terms of bandwidth.

For the example case where a noise bandwidth of about 6 Hz is desired, the optimal estimator gain matrix turns out to be $L = \begin{bmatrix} 0.0237 \\ 0.2852 \end{bmatrix}$. An excellent approximation to this gain matrix is $L = \begin{bmatrix} 2^{-5} \\ 2^{-2} + 2^{-5} \end{bmatrix}$. Furthermore, an excellent approximation to the sampling rate, T , of 1ms happens to be given by 2^{-10} . The implication is that the state space implementation of the loop can be implemented in fixed point arithmetic very simply by shifts and adds.

The following simplified 68000 assembly code is an example of the implementation of tracking loop time and measurement updates. It is assumed that the differential phase estimate is in data register D0, the differential phase rate estimate is in register D1, and the tracking error (difference of measured and estimated phase) is in D2.

```
; Measurement update
ASR    #2,D2      ; divide error by 4 (2^2)
ADD    D2,D1      ; apply adjustment to phase velocity estimate
ASR    #3,D2      ; divide error again by 8 (2^3, total of 2^5)
ADD    D2,D0      ; apply adjustment to phase estimate
ADD    D2,D1      ; apply adjustment to phase velocity estimate

; Time update
MOVE   D1,D2      ; make a copy of velocity estimate
ASR    #10,D2     ; multiply by the sampling period (1ms)
ADD    D2,D0      ; add it to the phase estimate
```


Of course, the actual code requires additional logic to compute the phase error from the correlator output and the phase estimate, taking into account cycle wrap-around, as well as the arc tangent look-up table to convert the raw I and Q measurements to the phase error.

The end result is an efficient processing structure which executes in minimum time. An example of the processing output is shown in Figure 2.2.3.4 which shows both the raw correlator samples and the estimator output. There is a relative range rate between the two antennas of about 3 m/s.

For comparison, the GPS data bits modulated on the carrier are also shown. The GPS data bits in the figure hint at a key performance advantage of direct mea-




Figure 2.2.3.4: Raw Differential Phase

surement of differential phase. In other forms of GPS operation, the carrier tracking loop bandwidth cannot easily be set faster than the 50 bps data stream because the loop would begin to follow the 180 degree phase shifts of the data itself. This limitation does not apply to *differential* phase tracking because sign changes affect the carrier identically at each antenna.

2.3 The Antenna Multiplexing Receiver

Antenna multiplexing represents the very highest level of integration for the differential phase receiver. By multiplexing several antennas into a single GPS receiver, the ultimate performance and cost for attitude determination improves substantially. First, the overall cost, size, weight, and power of the receiver are effectively reduced to that of a standard

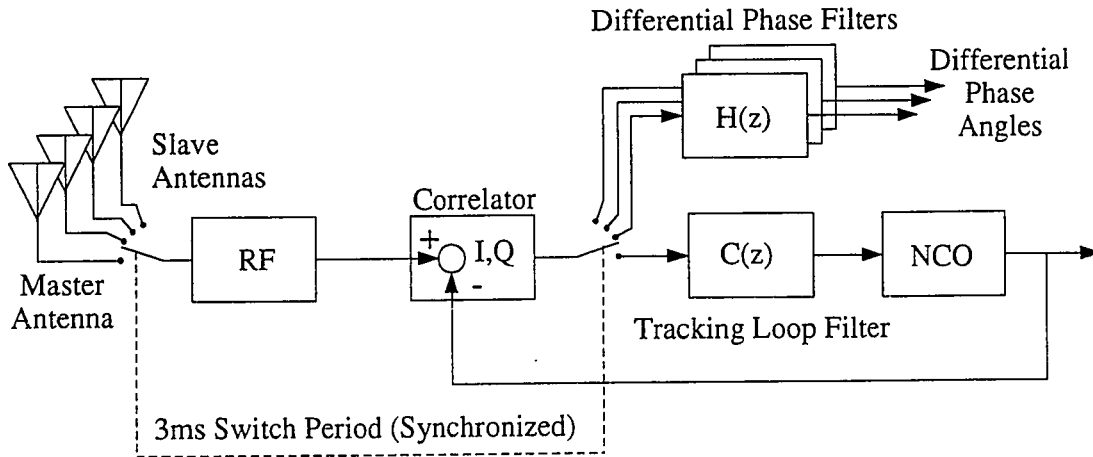


Figure 2.3.1.1: Antenna Multiplexing Top-level Design

positioning receiver. Second, variations in line bias with environmental changes are effectively eliminated.

The disadvantage to implementing antenna multiplexing is that there is not one function going on at the signal processing level of the receiver that is not completely incapacitated by antenna switches. To make this concept work, the signal processing core assembly code had to be almost completely rewritten. A number of the problems which arose are presented here.

The performance penalty incurred for adopting this architecture is that the single R.F. path is time-shared among the different antennas, resulting in a lower overall signal-to-noise ratio. During flight testing, this effective reduction has not been observed to cause any problems with performance.

2.3.1 FUNDAMENTAL IMPLEMENTATION

The fundamental implementation is shown in Figure 2.3.1.1. The salient feature of this design is that the signals from each antenna are time-multiplexed *through a single*

hardware path (except for the short antenna cables), using the receiver oscillator as the reference to tie together the separate carrier phase measurements from each antenna.

A PIN diode R.F. switch sequentially feeds the output of each antenna into the single hardware R.F. path. One antenna (designated the master) serves as the carrier phase reference. The master may be either assigned manually to be a specific antenna or assigned automatically to follow the antenna for which the signal strength is a maximum.

When the master antenna signal passes through the R.F. path, the traditional Costas loop is used to track the carrier. The output of the correlators at this time is the carrier phase tracking error. When signals of any of the slave antennas pass through the hardware channel, the NCO runs *open-loop*. The output of the correlators at these times corresponds to the differential carrier phase angle between a slave antenna and the master antenna. The transfer function $H(z)$ filters the raw differential phase measurements. The switching frequency is chosen to be as rapid as possible to track differential phase dynamics (typically $\sim 300\text{Hz}$).

2.3.2 IMPACT OF OSCILLATOR STABILITY

The theoretical feasibility of the multiplexed receiver is based on the stability of the receiver oscillator. Since the receiver oscillator is effectively the only foundation tying together phase measurements across different antennas, its stability over the sampling period is of paramount importance. Figure 2.3.2.1 shows a plot of the stability ranges of several types of high quality oscillators.

For reasons described in the following section, a switch period of 3 ms was selected. Assuming that the receiver returns to the reference antenna every four switch cycles, a stability interval of 12 ms results. From the chart, the stability ($\Delta f/f$) of a quartz oscillator

Reprinted from J.R. Vig, "Quartz Crystal Resonators and Oscillators", U.S. Army Electronics Technology and Devices Laboratory (LABCOM), Fort Monmouth, NJ, SLCET-TR-88-1 (Rev 4.5), October, 1991.




Figure 2.3.2.1: Stability Ranges of Various Frequency Standards




Figure 2.3.3.1: Frequency Plan of GPS Signal and MUX Receiver

over this period is much better than 10^{-8} . The resulting *conservative* range error is on the order of a centimeter over this period.

In practice, the performance is much better. This is also what one would expect from experience, since many atomic clocks employ quartz oscillators to flywheel over intervals shorter than one second. Since the size of the oscillator error in running open-loop is much less than the expected carrier-to-noise ratio over the time interval of one switch dwell period, the multiplexing scheme based on conventional quartz oscillators does not degrade accuracy, and more exotic clocks are unnecessary.

2.3.3 SELECTION OF SWITCH RATE AS A DESIGN PARAMETER

Choosing the switch rate is an important consideration in the receiver design. For reference, Figure 2.3.3.1 shows a full-page chart of the frequency scales of every fundamental process going on, both in the GPS signal structure and in the receiver. The figure also shows the lines of demarcation between hardware and software, microwave and digital electronics, and assembly and high-level software.

Two factors determine the actual range of viable rates. A lower limit on switch frequency is imposed by the GPS data rate of 50 bits per second. A switch rate lower than 50 Hz may mean loss of the ability to read data from the satellite data stream, especially if one or more antennas is shaded.

The integration time of the correlator hardware places an upper limit on switch frequency. The integrate-and-dump correlator in the DSP chip design currently cycles on the C/A code epoch, which occurs every millisecond. Therefore, 1 kHz is the fastest viable switch rate.

Aside from the two constraints setting the maximum and minimum switch rates, there are two competing factors which must be considered: detecting data bit transitions and dealing with host platform dynamics.

Receiver positioning depends critically on detecting the edges of the data bit transitions in the GPS signal data stream. The complete C/A code sequence repeats every millisecond. Therefore, the exact times of the data bit transitions are used to resolve the resulting integer millisecond ambiguity in C/A code range for conversion into pseudorange. Improper edge detection is potentially catastrophic, leading to GPS position errors well in excess of 300 km (the distance light travels in a millisecond).

As shown in Figure 2.3.3.2, to detect a clean bit transition, it is highly desirable to dwell on a single antenna for at least two consecutive correlator integration periods (2 ms) which straddle the potential bit transition. At least one correlator period coincides with one data bit, while a following correlator period coincides with the following bit of possibly opposite sign. Ideally, the receiver should spend a maximum amount of time on each antenna so as to maximize the probability of detecting each bit transition. In other words, lower switch rates are desirable from the standpoint of detecting bit transitions.

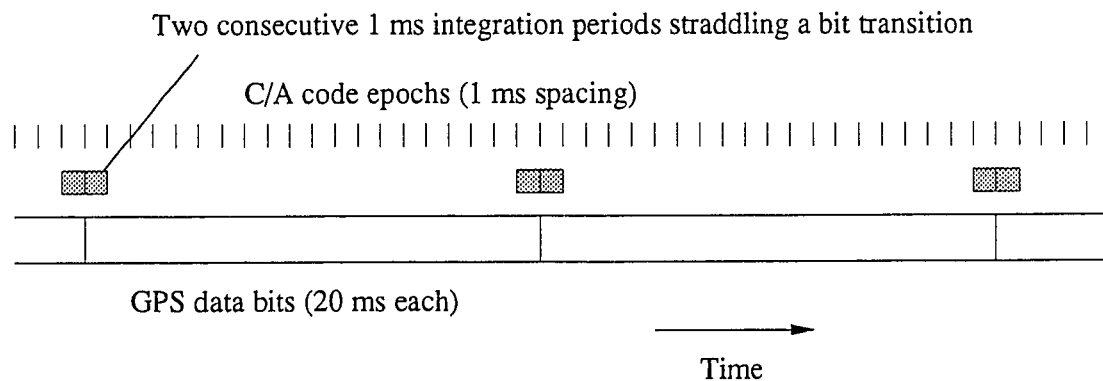


Figure 2.3.3.2: Detecting Bit Transitions During Multiplexing

On the other hand, consideration of the differential phase tracking loop response to dynamics drives the switch rate up. Generally, the switching time should be as fast as possible, not only to minimize the potential for aliasing (see sub-section 2.3.7), but to accommodate quickly changing carrier phase due to dynamics. In terms of phase-locked loop implementation, it is simplest to set the pre-detection interval (total correlator integration time per processed sample) equal to the dwell time on each antenna. To maximize tracking loop response and minimize smear, it is desirable to spend only the minimum amount of time to form a sample at the integrate-and-dump stage of processing.

As a compromise, an antenna dwell period of 2 ms was chosen for the initial design. This value represents middle ground between each of the competing factors. It represents the fastest possible switch rate which still allows for robust detection of bit transitions. Furthermore, it preserves consecutive millisecond integration periods which may also be used for detection of aliasing.

There is one other factor which, at least for now, needs to be included in consideration of switch rate. The asynchronous relationship between receiver time and received signal time (considered in the following section) results in a wasted millisecond sample whenever there is an antenna switch. The issue is that the antenna switch must, in general, occur right in the middle of a C/A code epoch interval, effectively destroying that sample. Therefore, a secondary consideration is that slower switch rates also lead to fewer lost samples.

Reconciliation Of Receiver Time And Received Signal Time

Unfortunately, the C/A code phase epochs (and, hence, time boundaries of integration in the receiver correlators) must necessarily be out of phase with one another at almost all times. (In fact, it is this very phasing that provides the fundamental range measurement

of GPS!) It is almost always impossible to choose an ideal time to switch antennas. Figure 2.3.3.3 shows that while it might be possible to switch at precisely the C/A code epoch of one particular satellite and not lose data, the code phase of the other satellites in track will not, in general, line up with the same phase.

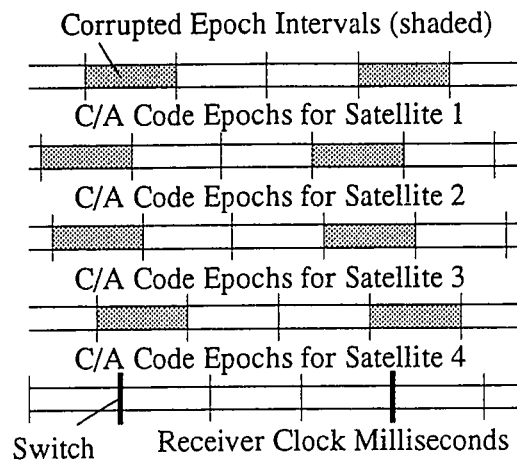


Figure 2.3.3.3: Antenna switch timing

The result is that no matter when the antenna switch occurs, it must corrupt at least some correlator samples. Currently, there is not a provision for latching the correlator output at a time other than the C/A code epoch; therefore, the sample straddling the antenna switch is simply discarded.

The antennas are switched open loop based on the receiver millisecond clock. The switch timing is appropriately structured for handling the entire switch event, from initiation of the switch, through integration of the correlator output, to interpretation of the correlator output. However, its greatest virtue comes in handling the pathological case when the C/A code epoch crosses over the receiver millisecond clock.

Under this fleeting yet frequently recurring condition (persisting for tens of milliseconds every few minutes in normal receiver operation) the receiver millisecond clock is effectively coincident with the C/A code epoch. From one millisecond to the next, the C/A code epoch may occur either slightly before or slightly after the millisecond edge. In these cases, it is possible for either 0, 1, or 2 epoch interrupts to occur in the interval of a single receiver clock millisecond. Even in this pathological situation, the timing structure allows antenna switching and correlator interpretation to occur smoothly.

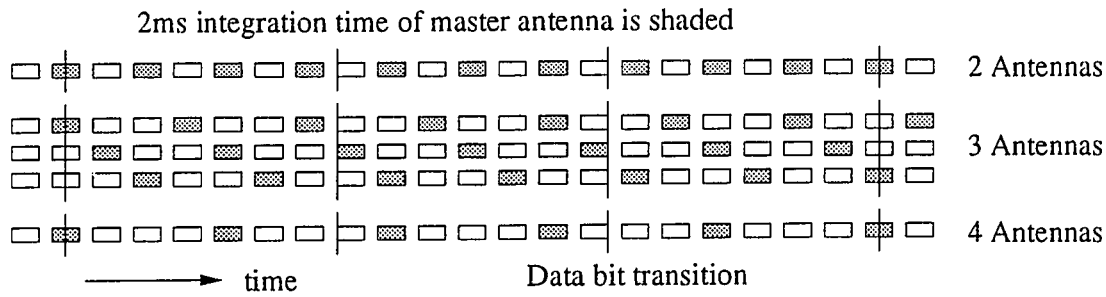


Figure 2.3.3.4: Data bit timing with respect to antenna switching

Detecting Bit Transitions

Given that the antenna switch rate does not preclude the detection of bit transitions, the issue is how does one guarantee that the receiver will always be able to detect bit transitions in the midst of an antenna switch events.

The issue is portrayed graphically in Figure 2.3.3.4. The pre-detection interval is shown as a box occupying 2 ms. The corrupted millisecond during the switch separates each consecutive block. The receiver sequentially cycles through all the enabled antennas. Depending on how many antennas are enabled, the master antenna (indicated by a shaded box) is switched on once per complete cycle through all the antennas.

The figure shows the case for either 2, 3, or 4 antennas enabled. Superimposed on the figure is the time of data bit transitions. A bit transition is detected when two consecutive millisecond integration periods exactly straddle the data bit edge. In the case of 3 antennas enabled, it can be seen that any initial phasing of the master antenna switching with respect to the timing of the GPS data bits eventually evolves into every other possible phasing. Therefore, it is only a question of a few tenths of a second before a bit transition is detected.




Figure 2.3.3.5: Phasing antenna switching to catch bit transitions (4 Antennas)

For the cases of 2 or 4 antennas, however, bit transition detection is not guaranteed, depending on the relative initial phasing. The case for 4 antennas is shown in Figure 2.3.3.5. The three possible times that the master antenna can be enabled for a given initial phasing are shown in the data bit of the first column. The union of these times is given immediately below the first column in the form of adjacent 2ms blocks. To ensure that all the possible relative phasings occur between the switching and the data bit transitions, the switching logic periodically adds an extra millisecond of dwell on an antenna. The result is that the antenna switching “window” slowly drifts across the bit transitions with time. To increase reliability in the presence of noise, logic requires that a bit transition be seen in the same place several consecutive times before it is finally accepted. This logic is implemented in assembly code for maximum execution speed.

This bit transition detection strategy is independent of the number of antennas shaded by obstructions. Only one antenna need be detecting a signal. The antenna with the strongest signal is designated as the master antenna. During normal multiplexing operation, several bit transitions on average are detected each second—more than a sufficient number for reliable operation and acquisition.

Reading Data

The ability to read data is another implicit requirement placed on the multiplexing receiver. From Figure 2.3.3.5 it is apparent that between 2 and 4 ms of integration time

are available for data collection when only 1 out of 4 antennas are receiving a signal. While this presents a significant power reduction when compared with the full 20 ms which are available without multiplexing, there is no problem with data collection in the non-jamming environments encountered in routine practice. In general, data collection is not impeded for a satellite whose signal is strong enough to be used for differential phase measurements anyway.

The data collection logic is set up around the scenario that only 1 of the enabled antennas is able to receive the signal from a given satellite in view. To keep the integration in synchronization with the bit transitions, a counter is maintained in assembly code to carry the receiver between observed bit transitions.

2.3.4 GAIN SELECTION

A gain table is used to select all tracking loop gains depending on the signal processing state selected by the user. Key factors governing gain selection are

- (1) Number of antennas enabled for multiplexing
- (2) High or low loop bandwidth settings (static or dynamic environments)
- (3) 1st or 2nd order tracking loops for differential phase
- (4) 2nd or 3rd order tracking loops for absolute position (see Section 2.3.5)
- (5) initial signal acquisition (see Section 2.3.6) or steady state operation

For picking nominal, steady state gains as a function of the number of antennas enabled, the following logic was applied. From the discussion of optimal estimation in Section 2.2.3, the only effective degree of freedom one has in selecting tracking loop gains is the loop bandwidth. For a given carrier-to-noise ratio, loop bandwidth is a trade-off between

ability to track rapidly changing dynamics (high bandwidth) and minimizing unwanted cycle slips (low bandwidth).

In other terms, the loop bandwidth may also be expressed in terms of the number of time samples required for the underdamped impulse response to cross zero. Based on practical experience over a variety of dynamic conditions with the nominal range of signal-to-noise ratio, excellent performance results from setting the tracking loop natural frequency, ω_n , to 10 samples per half cycle, or

$$\omega_n = \frac{\pi}{30nT}$$

where n is the number of enabled antennas, and T is the fundamental correlator integration period, 1ms. For four antennas enabled, the corresponding natural frequency of the loop is about 4 Hz. If fewer antennas are enabled, the natural frequency can be made considerably larger.

2.2.5 ABSOLUTE CARRIER PHASE TRACKING

In modifying the receiver core code to handle antenna multiplexing, the state space approach to signal processing was also applied to the absolute carrier tracking loops. The key advantages (in addition to those enumerated in Section 2.2.3) are

- *Flexibility*: State space is easier to reconfigure as a function of number of antennas enabled, loop bandwidth, and 2nd or 3rd order operation.
- *Consistency*: Since both the differential phase and frequency loops must operate over the entire I-Q plane instead of a small, linear region, the arc tangent discriminator (already implemented for differential phase) is ideal. To maintain

consistency between the differential and absolute carrier tracking processes, the $\text{signum}(\)$ detector originally used for carrier tracking was replaced by the I-Q look-up table, and a common state space structure was implemented throughout.

- *Economy*: Since the look-up table renders the phase discrimination insensitive to signal amplitude, additional signal processing logic normally required to handle varying signal amplitude can be eliminated.
- *Stability*: Third order tracking loops have three poles at the point $z=1$ on the unit circle. Therefore, at low values of loop gain, there must always be an unstable root. Use of the look-up table eliminates the destabilizing tendency of the signum detector to lower the loop gain when the signal amplitude diminishes.

The interesting issue of duality arises in the implementation of the state space carrier phase tracking loop. The question is whether the receiver channel is an estimator or a controller. It is really both. In one sense, the NCO is being *controlled* to track the GPS signal. On the other hand, the signal phase is being estimated, and the state of the NCO is simply one variable of the instantaneous state estimate. Treated as an estimator, the time and measurement update equations for a 3rd order tracking loop are

$$\begin{bmatrix} \overline{\varphi} \\ \overline{\dot{\varphi}} \\ \overline{\ddot{\varphi}} \end{bmatrix}_{k+1} = \begin{bmatrix} 1 & T & \frac{1}{2}T^2 \\ 0 & 1 & T \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \hat{\varphi} \\ \hat{\dot{\varphi}} \\ \hat{\ddot{\varphi}} \end{bmatrix}_k$$

$$\begin{bmatrix} \hat{\varphi} \\ \hat{\dot{\varphi}} \\ \hat{\ddot{\varphi}} \end{bmatrix}_k = \begin{bmatrix} \overline{\varphi} \\ \overline{\dot{\varphi}} \\ \overline{\ddot{\varphi}} \end{bmatrix}_k + \begin{bmatrix} L_1 \\ L_2 \\ L_3 \end{bmatrix} \Delta\varphi_k$$

where the state estimate consists of the instantaneous phase, phase rate, and phase acceleration. The NCO phase is driven by commanding a rate, u_k . The equations given above can be decomposed such that the NCO phase serves as the *a priori* phase state estimate *without explicitly appearing in any of the signal processing algorithms*.

The NCO can be modeled as the first order system, $\varphi_{k+1} = \varphi_k + Tu_k$. Defining the NCO state to be the *a priori* phase state estimate, $\bar{\varphi}_k$, and combining the NCO model with the first rows of both the time and measurement update equations, the resulting system is

$$\begin{bmatrix} \bar{\dot{\varphi}} \\ \bar{\ddot{\varphi}} \end{bmatrix}_{k+1} = \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \hat{\dot{\varphi}} \\ \hat{\ddot{\varphi}} \end{bmatrix}_k$$

$$\begin{bmatrix} \hat{\dot{\varphi}} \\ \hat{\ddot{\varphi}} \end{bmatrix}_k = \begin{bmatrix} \bar{\dot{\varphi}} \\ \bar{\ddot{\varphi}} \end{bmatrix}_k + \begin{bmatrix} L_2 \\ L_3 \end{bmatrix} \Delta\varphi_k$$

$$u_k = \hat{\dot{\varphi}}_k + \frac{1}{2}T\hat{\ddot{\varphi}}_k + \frac{L_1\delta\varphi_k}{T}$$

where now only the estimates of velocity and acceleration need to be maintained in memory, and the NCO rate is commanded by the given equation.

Analysis of Unmodeled Timing Errors

Note that the sample interval can deviate from its nominal value in normal operation, especially if extra milliseconds are added periodically to resolve bit transitions or if the C/A code epoch happens to be aligned with the receiver millisecond. To gain confidence that the tracking loop stability would not be adversely affected by changes in the actual sample rate, an augmented system model was developed using the method described in

Franklin & Powell.²¹ Given an unmodeled delay, mT , in the system, the state transition equation is

$$\begin{bmatrix} \hat{\phi} \\ \hat{\dot{\phi}} \\ \hat{\ddot{\phi}} \\ u_{prev} \end{bmatrix}_{k+1} = \begin{bmatrix} 1 & (1-m)T & \frac{1}{2}(1-m)T^2 & mT \\ 0 & 1 & T & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & \frac{1}{2}T & 0 \end{bmatrix} \begin{bmatrix} \hat{\phi} \\ \hat{\dot{\phi}} \\ \hat{\ddot{\phi}} \\ u_{prev} \end{bmatrix}_k + \begin{bmatrix} (1-m)(L_1 + L_2T + \frac{1}{2}L_3T^2) \\ L_2 + L_3T \\ L_3 \\ \frac{L_1}{T} + L_2 + \frac{1}{2}L_3T \end{bmatrix} \delta\phi_k$$

As expected, the system roots only move by a few percent for worst case values of m equal to $2/(3n)$, where n is the number of antennas enabled.


2.3.6 HALF CYCLE AMBIGUITY FROM DATA MODULATION

The data modulation on the GPS signal introduces a half cycle ambiguity between the carrier phase tracked on any two multiplexed antennas. Additional signal processing logic is used to check for sign inconsistencies within the same data bit between the slave antennas and the master. If this condition is detected, the phase of the inconsistent slave antenna is advanced by half a cycle.


2.3.7 ALIASING

Aliasing is perhaps the most troublesome aspect of the antenna multiplexing receiver. As shown in Figure 2.3.7.1a, two distinct sample periods are required for receiver operation: the fundamental correlator integration time ($T=1\text{ms}$) and the time it takes to cycle through all of the enabled antennas, $3nT$ (in this example, 12ms for $n=4$ antennas). Figure 2.3.7.1.b shows the corresponding aliasing modes in the frequency domain. Each sampling rate generates a set of aliasing modes. Furthermore, a third type of aliasing mode appears that is associated with the antenna cycle time.

(a) Time domain



(b) Frequency domain



(c) I-Q plane




Figure 2.3.7.1: Three Types of Aliasing Encountered in Multiplexing (4 Antennas)

Sample Period Aliasing

The aliasing mode associated with the sample period is relatively easy to detect and is typically encountered only upon initial acquisition. In a single sample period, the carrier phase traverses the semicircle shown on the in-phase and quadrature (I-Q) plane in Figure 2.3.7.1c. The value measured by the correlator is the arc centroid lying on the I axis. The presence of this aliasing mode is revealed by the sign of I changing each measurement sample.

Multiplexing Period Aliasing

The aliasing mode associated with the antenna multiplexing period is also encountered upon initial acquisition. Unfortunately, since it is relatively close to the nominal frequency, it can be induced occasionally by measurement noise or vehicle dynamics. This mode is active if, in the time it takes to cycle through all enabled antennas, the carrier phase error has moved by exactly half a wavelength. In allowance for the data demodulation, the I and Q values are reflected through the origin to yield *exactly zero phase tracking error*.

Multiplexing period aliasing occurs when the tracking frequency is offset from the true frequency by $\frac{1}{2} \frac{m}{3nT}$, where m is a non-zero integer. Note that as the number of enabled antennas, n , increases, the spacing between aliasing frequencies narrows (worsens).

This mode is combated by differencing the two adjacent, consecutive millisecond phase measurements. If no alias is present, the expectation value of the difference between these phases is zero. If the aliasing mode is present, this expectation value should be equal to $60m/n$ deg of phase. Therefore, logic is added to keep a running average of this difference.

Based on the signal-to-noise ratio of the received signal, it is possible to sketch out a statistical basis for choosing detection thresholds. Figure 2.3.7.2 shows a worst-case example of four antennas being multiplexed to track a relatively weak signal of 6 amplitude measurement units (AMU). [The amplitude measurement unit, defined for Trimble hardware, is given by the value of $I/64$, where 64 is approximately the square root of the number of averaging samples in a millisecond (~ 4092) such that the AMU is roughly the signal-to-noise ratio on an amplitude scale over a 1kHz bandwidth.] At 6 AMU, the noise level (1σ) for a single millisecond sample is $1/6$ radians, or about 10 deg of phase.




Figure 2.3.7.2: Aliasing Detection

Averaging 32 differences between the two consecutive millisecond samples, the resulting noise level is $\frac{\sqrt{2}}{\sqrt{32}} 10$ deg or about 2.5 deg, resulting in the pair of Gaussian probability distributions shown in the figure.

Setting the detection threshold is a trade-off between catching the aliasing as early as possible and minimizing false alarms. Tantamount to a detected cycle slip, an infrequent false alarm is not catastrophic. However, in carrier tracking applications, all occurrences of cycle slips should be minimized. It is also advantageous to minimize the possibility of the false velocity solutions (an error of about 10m/s is typical) attributable to aliasing. The threshold is set at the -1σ mark on the aliasing curve. Averaging 32 consecutive differences then provides an 83% chance of detecting the aliasing condition within a third of a second.

In static tracking, the aliasing condition is observed once every few hours, especially when a satellite is low on the horizon and deep fading due to multipath is occurring. As currently implemented, aliasing has not been observed for signal strengths above ~ 9 AMU. The nominal signal strength for a satellite in the middle of the antenna pattern is in the mid twenties.

In the long run, it appears that even the relatively infrequent occurrences of this form of aliasing are unacceptable in future applications of GPS attitude determination where system integrity is extremely important. The best prospect for eliminating the problem completely is to push the fundamental sampling frequency, $1/T$, up by an order of

magnitude or two. While there would be some minor impact on the DSP ASIC design, better overall system integrity would result.

Disturbance Aliasing

The last type of unwanted tracking loop behavior due to multiplexing is shown at right in Figure 2.3.7.1c. This mode simultaneously manifests characteristics of aliasing, limit cycling, and a steady-state disturbance.

Starting at point 1 with a phase rate (Doppler) estimate lagging behind by about half of the multiplexing period alias frequency, the NCO is commanded to *speed up* to catch up with the carrier phase. However, due to the erroneous phase rate, the tracking loop gain is insufficient to keep the carrier phase from crossing over to a negative value of I for the following sample at point 2a. Unfortunately, the GPS data demodulation logic requires that negative values of I be reflected through the origin, moving the sample to point 2b. On the basis of this negative phase tracking error, the control law commands the NCO to *slow down*, thus perpetuating the condition.

Figure 2.3.7.3 shows a sample of actual tracking data displaying this weakly stable mode. The plot is a one second trace of signal processing commanded NCO frequency, u , and phase velocity estimate, $\hat{\phi}$. After being stuck in the disturbance aliasing mode for several seconds, the state is eventually dislodged by measurement noise. What is interesting about this particular plot is that the state happens to fall the wrong way into the first multiplexing period aliasing mode, whereby it is promptly returned to the correct frequency. The plot clearly shows how the disturbance aliasing mode falls in between adjacent aliasing frequencies.




Figure 2.3.7.3: Example of Disturbance Aliasing Mode

The approach taken in solving this problem was based on the observation that the mode seemed to be weakly stable. Rather than combat it by adding special logic to the signal processing software, it was preferred to see if there was a way to choose the estimator gains so as to render the mode *weakly unstable*, and thereby *preclude its occurrence altogether*. It is now shown that precluding this mode *is* possible.

Mathematically, the mode can be analyzed by modeling two consecutive state transitions and assuming steady-state behavior. Given an initial *a priori* state estimate, \bar{x}_0 , the measurement error, $e_0 = -H\bar{x}_0$, leads to the following measurement update (assuming no noise):

$$\hat{x}_0 = \bar{x}_0 - LH\bar{x}_0$$

The L matrix is the estimator gain matrix (set to steady state Kalman gains), and H is the observation matrix. Again, it has been assumed that the NCO state is the same as the *a priori* state estimate. The time update is then given by the state transition matrix, Φ

$$\bar{x}_1 = \Phi\hat{x}_0 = \Phi(I - LH)\bar{x}_0$$

When the following measurement update is performed, there is a $\frac{1}{2}$ cycle wrap-around to bring the phase error back within the range $-\frac{1}{4}$ to $\frac{1}{4}$ cycles. This adjustment effectively acts as a disturbance to the controller/estimator loop, driving the nominal state estimate error away from zero. The resulting measurement error, $e_1 = \frac{1}{2} - H\bar{x}_1$, leads to the following measurement update:

$$\hat{x}_1 = \bar{x}_1 + L\left(\frac{1}{2} - H\bar{x}_1\right)$$

Applying the final time update and noting that the phase state at time 2 must be the same as that at time 0 but advanced by half a cycle, the state transition from time 0 to time 2 is

$$\bar{x}_2 = \bar{x}_0 + \frac{1}{2} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \Phi(I - LH)\Phi(I - LH)\bar{x}_0 + \frac{1}{2}\Phi L$$

From this expression, it is possible to solve for the state estimate in steady state

$$\bar{x}_0 = \frac{1}{2} [\Phi(I - LH)\Phi(I - LH) - I]^{-1} \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} - \Phi L \right)$$

The presence of this aliasing mode is assured if the steady state phase angle, $\delta\phi = H\bar{x}_0$, ranges between 0 and $\frac{1}{4}$ cycles. However, if L is set sufficiently large such that the steady state phase angle is greater than $\frac{1}{4}$ cycle, reflections through the origin will no longer be necessary, and the system will never fall into this mode.

The behavior of this aliasing mode is very similar to the steady state error resulting from a disturbance to a control system. In this case, the steady state error is given by the difference $e = \frac{1}{2} - H\bar{x}$, where the $\frac{1}{2}$ arising from the data demodulation represents the disturbance. For higher loop gains, the disturbance rejection improves to the point that the steady state error is not large enough to cross the $\frac{1}{4}$ cycle boundary and initiate aliasing.

2.3.8 ACQUISITION

The process of search and acquisition of the GPS signal is greatly complicated by antenna multiplexing. The first approach considered was to see if search and acquisition would be possible without multiplexing. The original idea was to initiate antenna switching only after each desired satellite was completely locked up. Unfortunately, this approach

was destined to pose too many operational restrictions. It was concluded that the most prudent method of structuring the acquisition is to have each channel be able to search and acquire *with antenna multiplexing in full operation*. The operation of each channel should, therefore, be completely independent. One channel in the process of acquiring a signal should not be able to influence or disrupt any other channel already successfully tracking and making measurements.

In order to achieve lock-up in the fastest possible time, the state transition structure shown in Figure 2.3.8.1 was adopted. Searching for a given satellite is performed throughout the space defined by PRN code phase, Doppler, and (optionally) by an-

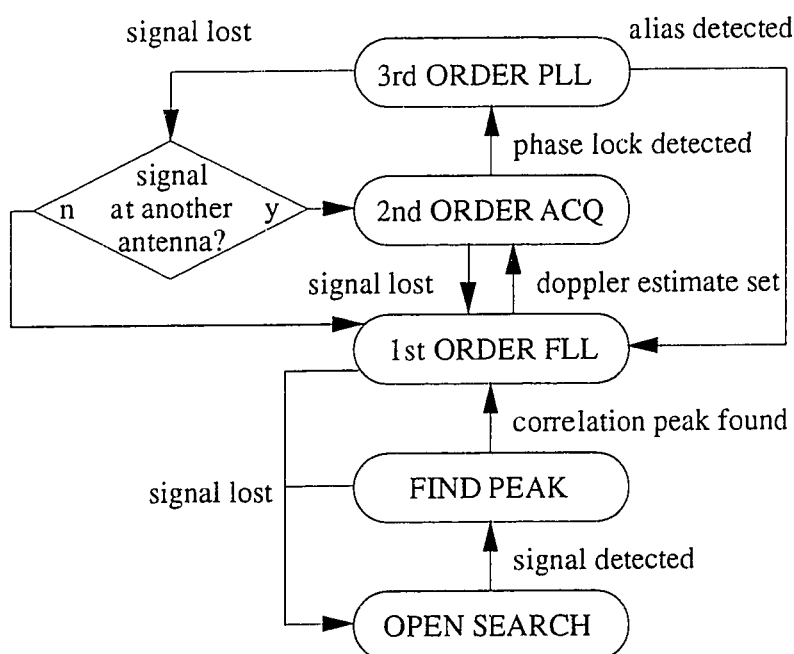


Figure 2.3.8.1: Signal Processing State Transition Diagram

antenna. If searching through each antenna independently is selected, the search time naturally becomes proportional to the number of antennas enabled.

Once a signal has been detected and the code phase positioned at the apparent correlation peak, the carrier loop is closed in a series of steps. First, a first order frequency lock loop is used to bring the estimated carrier frequency as close as possible to the actual value. This frequency loop employs the following discriminator function:

$$e_k = \frac{\Delta\varphi_k - \Delta\varphi_{k-1}}{T}$$

such that the error, e_k , is roughly proportional to phase velocity and each $\Delta\varphi$ measurement is obtained from the look-up table based on raw I and Q measurements. The frequency loop was employed for three reasons:

- *Acquisition dynamic range* — Search time can be minimized if the frequency search bins are as large as possible. The frequency loop funnels the wide search bin into the narrow acquisition window of the phase-locked loop.
- *Reducing multiplexing period aliasing during acquisition* — The frequency loop brings the Doppler estimate toward the correct tracking frequency before phase acquisition. This decreases the probability that the loop will fall into this aliasing mode upon initial acquisition.
- *Favorable performance with weak signals* — Since cycle ambiguities have no effect on frequency loop operation, lower gains can be used over a wide range of carrier frequencies. The net result is an improved probability of locking onto weaker signals.

After the frequency loop has run for several tens of milliseconds, the Doppler estimate is close enough to ensure successful phase acquisition. At this point, a second order phase locked-loop is closed with acquisition gains set as specified in the previous section to preclude disturbance aliasing. Once the signal processing logic senses that the loop is in lock, the system state graduates to the full tracking mode, closing the loop around both code and carrier using the nominal regulator gains and third order tracking (if third order




Figure 2.3.8.2: Carrier Loop Lock Acquisition

tracking has been enabled by the user). Note that during the signal acquisition process, either of the first two aliasing modes can occur. If aliasing is detected, the carrier frequency is immediately adjusted by the appropriate amount, and the state machine is reinitialized into the frequency mode.

Figure 2.3.8.2 shows an eventful example of the NCO commanded frequency during acquisition. During the code phase search at a constant frequency, a potential signal is discovered. After centering the code phase correlation peak, the frequency loop is closed. After two loop time constants, the tracking loop state transitions to the 2nd order phase-locked loop. After the transient settles out, phase lock is detected, and the tracking loop gains are set to 3rd order operation. Almost immediately, the alias detection logic uncovers the presence of the sample rate alias and adjusts the tracking loop state by the appropriate offset. This time, the acquisition proceeds normally, achieving steady-state, 3rd order operation.

2.4 Considerations for Launch and Low Earth Orbit

The extremely high Doppler shifts induced onto the signal by orbital velocities places additional requirements on a GPS receiver intended for the space environment. Since it is a goal to fly the attitude receiver in space, considerable attention was paid to ensure that the signal processing would be able to handle the additional challenges of this environment.

Oscillator Vibration and g Sensitivity

The crystal oscillator used in any receiver design is naturally sensitive to the vibrations and g forces imposed by the launch vehicle. Dealing with vibration is much less a signal processing issue than a hardware issue. The vibration frequencies are typically well

above the tracking loop bandwidth. Therefore, the oscillator hardware has to be designed to an acceptable level of vibration and g -sensitivity performance from the very outset.

Time Tagging

At orbital velocities, the time tag of each range measurement becomes much more critical than for terrestrial operation. If a time tag is off by even a millisecond, the spacecraft carrying the receiver will have moved by 8m along track—a significant position error, especially if centimeter-level kinematic positioning is being carried out with respect to a ground-based reference network. Therefore, as described in Appendix A, the signal processing software was carefully scrutinized to eliminate all errors in measurement timing.

User Spacecraft Almanac for Search and Acquisition

On Earth, the range of GPS Doppler shifts is about $\pm 5\text{kHz}$. In space, it is about $\pm 40\text{kHz}$. The dynamic range of receiver operation must be expanded by almost an order of magnitude to accommodate operation in space. With the additional requirement of being able to search 4 antennas individually for a signal, one quickly arrives at the conclusion that a blind search in space will easily take longer than the time the desired GPS satellite is visible in the sky.

By incorporating a user spacecraft almanac inside the receiver, the frequency search window can be reduced such that, if the receiver oscillator offset and GPS constellation almanac are known, the orbiting receiver can lock up within seconds. If any of these initialization parameters are inaccurate, the user spacecraft almanac helps speed the subsequent search considerably. Appendix B describes how these calculations are performed.




Figure 2.4.1: Ephemeris Feed-Forward

Figure 2.4.1 shows an example of the search windows for a case where the oscillator bias rate has been set off nominal by an extreme of 5kHz and four antennas are enabled. In this plot, the state of the signal processor search has been interrogated periodically to trace the dimensions of the search window and the actual search frequency. Note how the center of the search window is updated to follow the nominal orbital motion. Also, note that the acceleration has been fed forward and is being propagated by the signal processor at the assembly code level.

Steady-State Acceleration

Handling the steady state acceleration of a launch vehicle is a primary concern in the realm of signal processing, not only for nominal tracking under acceleration, but also for initial acquisition as covered in Appendix B. Third order tracking loops (as described in Section 2.2.5) were designed especially to handle the steady-state acceleration encountered during boost and the nominal $\sim 1g$ acceleration in low Earth orbit.

Code Phase Bias Rate

To further accelerate searches prolonged by the addition of 4 antennas, a bias rate applied to the code phase was extended to cover the entire search process. At high Doppler frequencies, the code phase (range) rate will be a significant fraction of the code phase search rate. This idea works because the code phase bias rate is already known by the current Doppler search bin. The resulting search duration is absolutely minimized.

2.5 Carrier Phase on a Spinning/Oscillating Platform

Measurements of either differential or absolute carrier phase made from an oscillating or spinning platform could benefit significantly by the signal processing design proposed in this section. An interesting variation of the attitude determination problem is the possibility of using a single antenna mounted away from the rotation axis of a spinning




Figure 2.5.1: Geometry of a Spinning Platform

platform. Cohen²² has demonstrated that the Doppler observable can be used to extract the attitude of a spinning spacecraft. Estimation of spacecraft nutation using GPS has also been proposed.²³

The rotating platform could also be a helicopter blade with the caveat that the plane of rotation may not always coincide with the body reference frame. At a typical rotation rate of 160rpm, the 1kHz sampling rate of the signal processor is certainly appropriate. An efficient signal processing concept is now proposed, which allows the rotation-induced phase modulation to be tracked in real time.

Whether a platform is spinning or oscillating, the measured carrier phase oscillates sinusoidally in proportion to the projection of the antenna position into the line-of-sight to the GPS satellite. Both attitude and position information can be extracted in real time from a single antenna.

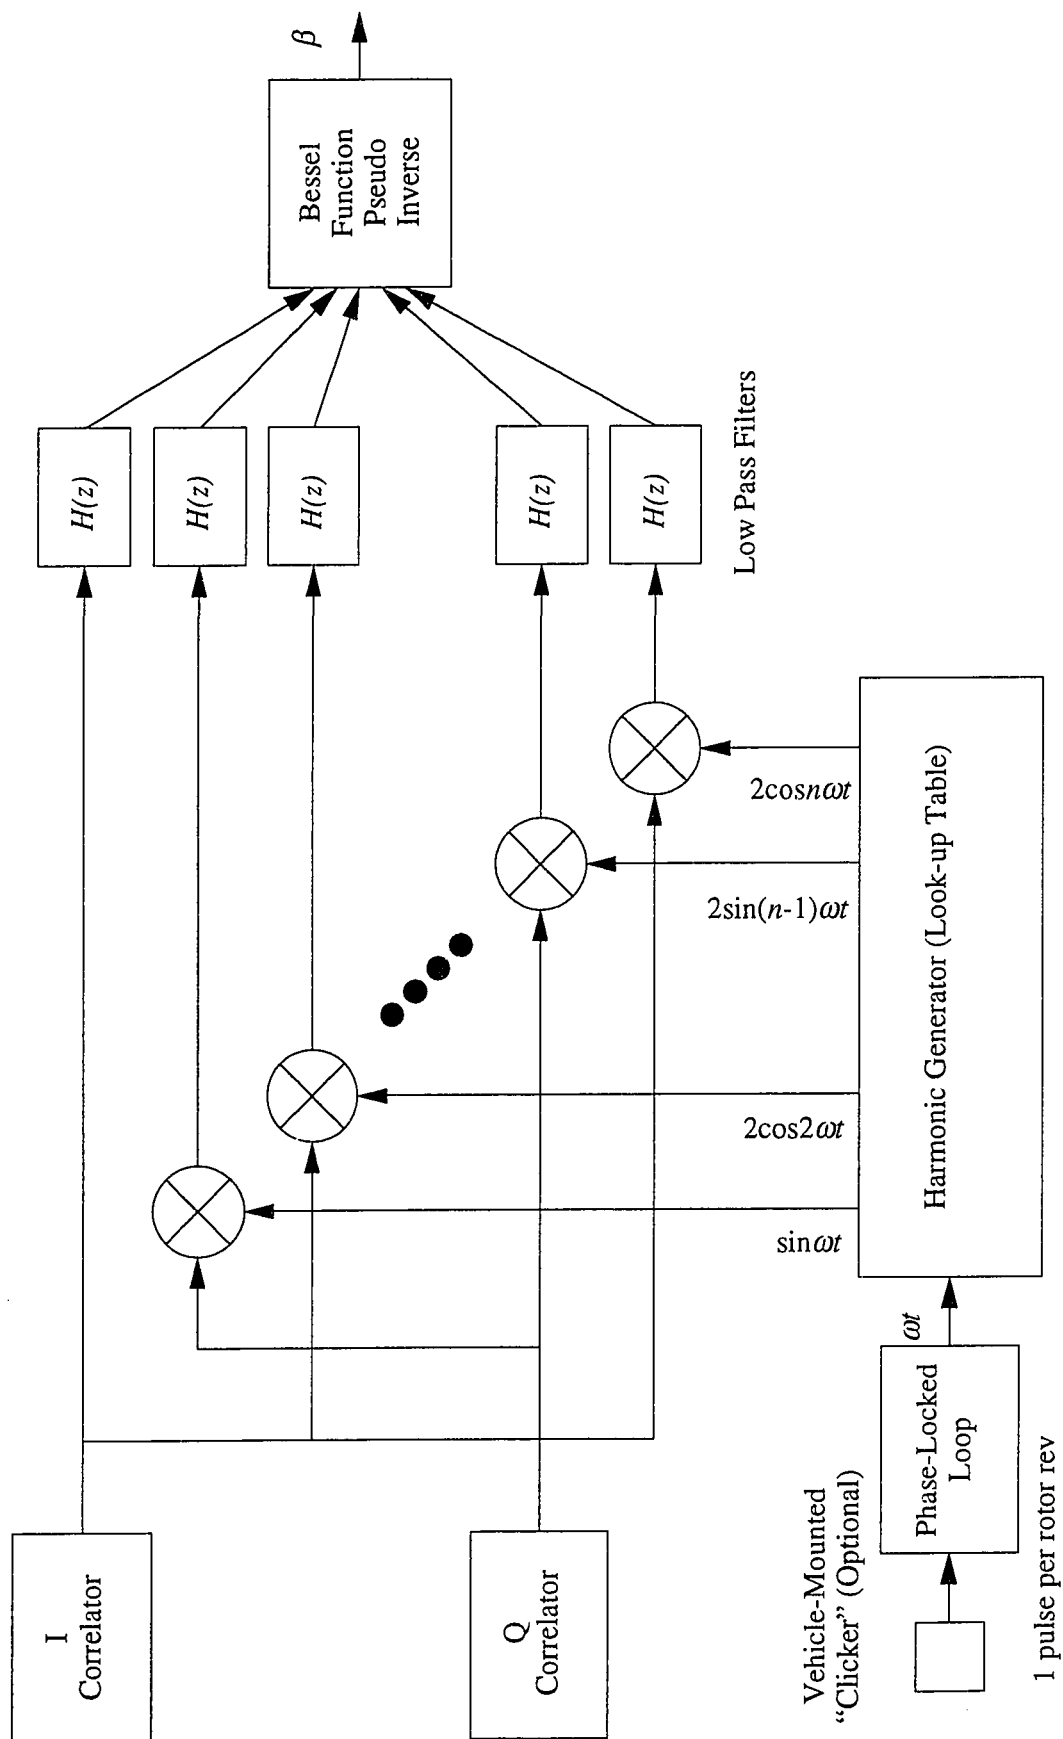


Figure 2.5.2: Harmonic Correlator Signal Processing for Rotor-Mounted Antenna

Figure 2.5.1 shows a diagram of the rotation geometry. Reading samples directly from the I-Q correlators, the phase modulation can be expanded in terms of Bessel functions

$$I = \cos(\beta \sin \omega t) = J_0(\beta) + \sum_{n \text{ even}}^{\infty} 2J_n(\beta) \cos n\omega t$$

$$Q = \sin(\beta \sin \omega t) = \sum_{n \text{ odd}}^{\infty} 2J_n(\beta) \sin n\omega t$$

where $\beta = a \sin \theta$. The correlator output can be processed either by the bank of phase-locked loops at harmonics of the spin frequency shown in Figure 2.5.2 or by applying an FFT. The Fourier coefficients of the correlator output can then be matched to Bessel functions to extract a real-time estimate of the angle between the spin axis and the line-of-sight vector to the GPS satellite. Note that this processing scheme does not require any knowledge whatsoever of the cycle ambiguities.

This same processing approach also has considerable potential when applied to flexible space structures or aeroelastic structures. Placing omnidirectional antennas at various positions along a flexible structure, the highest practical frequency that can be tracked by an ordinary phase-locked loop is on the order of 20 Hz. However, by observing the raw correlator I and Q measurements in the frequency domain, the measurable frequencies can reach as high as 1kHz.

If only one structural mode is excited, the sinusoidal oscillation will generate harmonics of the fundamental frequency. The same Bessel function estimation scheme can be applied to the Fourier spectrum to extract the amplitude and phase of the input oscillation. If more than one mode is present, harmonics of the fundamental input frequencies are accompanied at the correlator output by sums and differences of the input frequencies.

Chapter 3:

Cycle Ambiguity Resolution

Determining how many integer wavelengths lie between a given pair of antennas in the direction of a given GPS satellite is the crux of cycle ambiguity resolution. It is the most important initialization step that must be performed before attitude determination using GPS can achieve full accuracy. Since the GPS receiver only measures the fractional component of phase (see Figure 3.1), knowledge of the integers is required before differential *phase* measurements can be translated into differential *range* measurements. Many solutions to this problem have been proposed in the literature with varying degrees of applicability and effectiveness. Several new, efficient, and reliable methods are presented in this chapter in an effort to eliminate shortcomings of previous approaches.




Figure 3.1: Cycle Ambiguity Resolution

In GPS applications, the issue of cycle resolution first became clear in surveying, where the relative position of two different antennas is precisely determined down to the centimeter level. The static nature of the problem helps considerably. Unfortunately for attitude determination, static antenna placement may no longer be assumed; however, the known relative position constraints among the antennas *can* be applied.

Note that in some applications of attitude determination, the rigid body assumption does not apply. For example, the wing flexure in aircraft applications is a function of the flight condition and is completely unknown at initialization time. As shall be demonstrated in this chapter, the new methods of integer resolution handle wing flexure easily and even provide a quantitative measure of its magnitude.

Baseline Length Constraint

Consider a platform with a single baseline constructed from two antennas. The baseline vector originates at the master antenna and ends at the slave. Since differential positioning is employed, no generality is sacrificed by assuming that the tail of the vector stays fixed in space as




Figure 3.2: Baseline Length Constraint

shown in Figure 3.2. The possible positions of the slave antenna are constrained to lie on the surface of a virtual sphere of radius equal to the baseline length.

Integer Searches

In an integer search, different combinations of integer candidates (which can number in the hundreds of millions for antenna separations of even just a few meters) are

systematically checked against a cost function until, hopefully, the correct set is found. Several authors^{7,24,25} have resorted to this approach.

The search technique is depicted to scale in Figure 3.3 for a single 4λ baseline, three satellites, and 3σ multipath error. Satellite line-of-sight vectors are depicted with arrows.

Possible integer values are shown as concentric bands about the line-of-sight vectors. The correct integer set is indicated in white. Any place on the sphere where the concentric bands for all three satellites intersect (indicated in black) is a viable baseline orientation candidate. *It is clearly evident that there is not a unique solution.*




Figure 3.3: The Integer Search

Although search techniques work well for smaller baselines (on the order of a couple carrier wavelengths), they are vulnerable to erroneous solutions with longer baselines or when fewer satellites are visible.

While many innovative schemes have been synthesized for improving the mathematical efficiency of the search process and using code phase to narrow the search space, none of these can possibly address the fundamental liability of their foundation: the search method itself. Its greatest weakness is its propensity to arrive at wrong solutions.

Motion-Based Methods

The motion-based approach presented in the following sections has none of the weaknesses of the search technique and provides some rather distinct advantages. The key differences are outlined as follows:

- *Reliability:* When treated as a bias estimation problem, the integers are obtained analytically. A final sanity check can be applied to verify that the biases—estimated as continuous parameters—actually fall on integer values. Furthermore, it is straightforward to determine analytically exactly if and when there is enough information to resolve the integers rather than having to guess.
- *Robustness:* The motion-based approach degrades gracefully in situations where the vehicle parameters are way out of specification. As examples, consider wing flexure, or suppose that the host vehicle is an aircraft in an emergency situation (somehow damaged in flight). Search methods fail altogether in these situations, while the motion-based approach continues to provide accurate (although partially degraded) results.

Motion-based integer resolution algorithms make use of additional information provided by platform or satellite motion. Brown and Ward²⁶ describe a transformation for identifying the integers for a single baseline based on large angle motion. The following sections describe other novel motion-based methods as well as a new extension of this transformation. The cases discussed are platform motion (where the time scale of platform motion is very much faster than the time scale of GPS satellite motion), SV motion (static platforms), and quasi-static motion (where the time scale of platform motion is comparable to that of the GPS satellites).

3.1 Platform Motion

An abundance of information useful for integer resolution results from the attitude motion of the platform. With an appropriate antenna configuration, little platform motion is required, even if the baseline length, L , is very much greater than the wavelength, λ . For aircraft applications, cycle ambiguities may be resolved through the natural attitude motion, consisting of banks, turns, or small attitude motion excited by turbulence. If the baseline vectors are non-coplanar, even a turn on the ground (about a single axis) is sufficient. Thus initial ambiguities may be resolved with the first taxi turn.

Experiments applying an extended Kalman filter to the integer ambiguity resolution problem yielded very poor results. Convergence was slow, marginal, and always sensitive to the initial conditions, which generally had to be fairly close to the actual attitude for the filter to converge. Other authors²⁷ have corroborated this behavior. In the interest of achieving maximum speed and efficiency, an approach culminating with a batch non-linear least-squares fit proved to work much better.

The typical sequence of processing involved is to map the raw differential phase measurements into a set of displacement (delta position) vectors. An analytical mapping is used to transform these displacement vectors into an initial guess of the platform attitude. The initial guess—parameterized in terms of quaternions—is then refined using a full non-linear least-squares fit to best match the raw phase measurements. The previously unknown integers follow immediately.

3.1.1 DISPLACEMENT VECTORS

Without knowledge of the integers, it is possible to determine the Cartesian position of the slave antenna relative to its unknown starting point. As shown in Figure 3.2, positioning of the slave antenna is always with respect to the master. Since both antennas share a common local oscillator, only three satellites (one for each Cartesian spatial dimension) are required to extract this information. As shown in Figure 3.1.1.1, the measured differential range (phase), $\Delta\phi$, measured in wavelengths, is proportional to the projection of the baseline vector, x (3×1), measured in wavelengths, into the line of sight unit vector to the satellite, \hat{s} , for baseline i and satellite j :




Figure 3.1.1.1: Differential Phase Geometry

$$\Delta\phi_{ij} = \hat{s}_j^T x_i - k_{ij} + \beta_i + v_{ij}$$

The measurement, $\Delta\phi$, is subject to an integer ambiguity, k , an additive line bias, β ($0 \leq \beta < 1$), and noise, v . This differential phase measurement equation may be expressed in compact vector and matrix notation for the n satellites in view:

$$\Delta\phi_i = S^T x_i - k_i + \beta_i + v_i$$

where

$$\Delta\phi_i(n \times 1) = \begin{bmatrix} \Delta\phi_{i1} \\ \Delta\phi_{i2} \\ \vdots \\ \Delta\phi_{in} \end{bmatrix}, S(3 \times n) = [\hat{s}_1 \quad \hat{s}_2 \quad \cdots \quad \hat{s}_n], k_i(n \times 1) = \begin{bmatrix} k_{i1} \\ k_{i2} \\ \vdots \\ k_{in} \end{bmatrix}, \text{ and } v_i(n \times 1) = \begin{bmatrix} v_{i1} \\ v_{i2} \\ \vdots \\ v_{in} \end{bmatrix}$$

The line bias, β_i , is still a scalar, since, from the antenna multiplexing hardware, there are no inter-channel biases.

Note that only 3 satellites are required to determine the relative location of the slave antenna undergoing motion. Note that this is one fewer than the 4 satellites required for kinematic survey, since the differencing operation between antennas is performed immediately against the common oscillator in the receiver hardware.

Suppose a baseline moves from Cartesian position x_0 to x . The measured change in differential range, Δr , (with noise neglected) is given by

$$\Delta r = \Delta \varphi - \Delta \varphi_0 = S^T (x - x_0) = S^T \Delta x$$

Since the line bias and integer ambiguity cancel out of the above expression, one can solve for the displacement vector, Δx (3×1), explicitly using a linear least-squares fit. Note that it has been assumed that the position displacement is occurring on a much faster time scale than that of the satellite line of sight vectors, \hat{s} .

Section 6.1 provides an example of experimentally-determined displacement vectors.

3.1.2 COMPUTING THE INITIAL GUESS FOR LARGE ANGLE MOTION

As baseline motion is occurring, it is possible to accumulate a set of displacement vectors for over a short interval of time. If the platform moves by a large angle, the set of displacement vectors can be used to calculate an initial guess to initialize a non-linear least-squares fit.

In the following development, the motion-based approach of Brown and Ward is extended to multiple baselines by including terms associated with cross-correlation of measurements across different baselines. In doing so, a significant limitation of the original method is overcome, namely that *integer resolution is possible with rotation about only a single axis of rotation* (when baselines are non-coplanar).

A two-dimensional representation of the rigid body antenna mounting constraints is shown in Figure 3.1.2.1. Suppose that the baseline vector, x rotates in space (about the master antenna without loss of generality). Here the baseline vec-




Figure 3.1.2.1: Large Angle Motion Initial Guess

tor is moved (displaced) by the vector Δx to two different locations at two different times, 1 and 2. If a plane is constructed perpendicular to each Δx vector passing through its midpoint, note that the center of the sphere must be included in the plane.

By simultaneously considering each Δx vector, the center of the sphere can be located along with the initial position, x , of the baseline. Mathematically, the solution may be developed by constructing the square of the norm of the rotated baseline vector, $x + \Delta x$, given as

$$(x + \Delta x)^T (x + \Delta x) = x^T x + 2\Delta x^T x + \Delta x^T \Delta x$$

Noting that the left side is equal to the square of the baseline length as is $x^T x$, the two terms may be canceled, leaving

$$2\Delta x^T x = -\Delta x^T \Delta x$$

Different Δx vectors taken at N different times (indicated by a superscript in parentheses) may be stacked into matrix form as follows

$$2 \begin{bmatrix} \Delta x^{(1)T} \\ \Delta x^{(2)T} \\ \vdots \\ \Delta x^{(N)T} \end{bmatrix} x = - \begin{bmatrix} \Delta x^{(1)T} \Delta x^{(1)} \\ \Delta x^{(2)T} \Delta x^{(2)} \\ \vdots \\ \Delta x^{(N)T} \Delta x^{(N)} \end{bmatrix}$$

Then the baseline solution x may be obtained through a linear least-squares fit. Each Δx vector defines a subspace in which the slave antenna must lie. The baseline is then the point which comes closest to this condition in a least-squares sense. A convenient shorthand notation for this least-squares fit is given by

$$2\Delta X^T x = -\text{diag}(\Delta X^T \Delta X)$$

$$\text{where } \Delta X (3 \times N) \equiv [\Delta x^{(1)} \quad \Delta x^{(2)} \quad \dots \quad \Delta x^{(N)}]$$

and $\text{diag}()$ corresponds to a vector ($N \times 1$) comprised of the diagonal elements of the argument ($N \times N$).

Unfortunately for the single baseline case, large angle rotation about the two axes perpendicular to the baseline is *always* required to completely resolve the integer ambiguities using motion. [Brown and Ward erroneously state that it is possible.] To see

why, note that Figure 3.1.2.1 depicts the special case where rotation occurs about only a single axis (the one coming out of the page). The intersection of the planes perpendicular to the Δx vectors is a line rather than a point. Imagine if the true baseline location were actually out of the plane of the page. Given any noise on the measurements, it is difficult to establish the correct component of position out of the plane of the page. The observation matrix is not full rank. Rotation about a single axis does not lend itself to robust cycle resolution for the single-baseline case.

To address the shortcoming of requiring two-axis motion perpendicular to each baseline, it is suggested here that information from multiple baselines be combined into a single simultaneous estimation equation. The idea is to use the constraint equation between different baselines i and j as well:

$$(x_i + \Delta x_i)^T (x_j + \Delta x_j) = x_i^T x_j + x_i^T \Delta x_j + \Delta x_i^T x_j + \Delta x_i^T \Delta x_j$$

Again the dot product of each baseline pair is constant, and the corresponding term may be canceled from both sides leaving

$$\Delta x_j^T x_i + \Delta x_i^T x_j = -\Delta x_i^T \Delta x_j$$

Combining Δx measurements from N different times (typically 20 to 60 epochs), this form may be expanded as follows:


$$\begin{bmatrix} \Delta x_j^{(1)T} \\ \Delta x_j^{(2)T} \\ \vdots \\ \Delta x_j^{(N)T} \end{bmatrix} x_i + \begin{bmatrix} \Delta x_i^{(1)T} \\ \Delta x_i^{(2)T} \\ \vdots \\ \Delta x_i^{(N)T} \end{bmatrix} x_j = - \begin{bmatrix} \Delta x_i^{(1)T} \Delta x_j^{(1)} \\ \Delta x_i^{(2)T} \Delta x_j^{(2)} \\ \vdots \\ \Delta x_i^{(N)T} \Delta x_j^{(N)} \end{bmatrix}$$

Invoking the same matrix notation as above, the entire initial guess for the case of three baselines shown in Figure 3.1.2.2 can be combined into a single, unified least-squares fit equation:

$$\begin{bmatrix} \Delta X_2^T & \Delta X_1^T & 0 \\ \Delta X_3^T & 0 & \Delta X_1^T \\ 0 & \Delta X_3^T & \Delta X_2^T \\ 2\Delta X_1^T & 0 & 0 \\ 0 & 2\Delta X_2^T & 0 \\ 0 & 0 & 2\Delta X_3^T \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = - \begin{bmatrix} \text{diag}(\Delta X_1^T \Delta X_2) \\ \text{diag}(\Delta X_1^T \Delta X_3) \\ \text{diag}(\Delta X_3^T \Delta X_2) \\ \text{diag}(\Delta X_1^T \Delta X_1) \\ \text{diag}(\Delta X_2^T \Delta X_2) \\ \text{diag}(\Delta X_3^T \Delta X_3) \end{bmatrix}$$

The left-hand matrix is now $6N \times 9$, the solution vector is 9×1 , and the right hand vector is $6N \times 1$. This same matrix structure applies to any number of baselines. For the case of three or more baselines, two important advantages fall out of this approach.

First, with motion about any two *arbitrary* axes, no *a priori* information about the antenna placement is required to unambiguously solve for all 3 of the three initial baseline vectors (3 components each, 9 total dimensions). Note that it is the measurements themselves that are providing all the geometrical information. Therefore, this approach could even be used to perform an *in-flight self-calibration of GPS baselines*. No prior knowledge whatsoever about their relative positions is required.



Second, by incorporating the known baseline constraints in the case of

Figure 3.1.2.2: Non-Coplanar Baseline Rotation

non-coplanar baseline configurations, *a single axis of rotation is entirely adequate for an unambiguous baseline vector solution*. This eliminates much of the dependency on the *type* of motion that the antennas must undergo before there is sufficient information to resolve the cycle ambiguities.

In practice, it is likely that the integers will be resolved through static means long before the aircraft engines are turned on. But as a second tier of redundancy, this approach allows the aircraft to identify the integers with only a small turn on the taxiway. The integers are practically guaranteed to be known within the first few seconds of taxiing. Once the integers are established, it is rare that they would ever need to be resolved again for the remainder of the flight.

The most remarkable aspect of this approach is that the inherently non-linear problem of attitude estimation is recast into a completely linear estimation problem. Of course it is advocated that the linear solution be treated as the initial guess for the full non-linear least-squares problem.

3.1.3 SMALL ANGLE MOTION: THE “EIGENAXIS” METHOD

As originally proposed by Cohen,¹⁰ the technique described in the previous section is now extended to small angle motion. Although the large angle method works very well, there are many circumstances in which there simply is not enough large angle motion to obtain a satisfactory initial guess. Examples include the rocking of a boat on choppy seas or the small angle motion experienced by an aircraft in turbulence or on final approach. In each of these cases, it turns out that there is usually still enough information available in the motion to successfully resolve the cycle ambiguities.

Presented in this section is a matrix approach which may be used to reduce displacement vectors into an initial guess based on *small angle motion*. This approach is called the “Eigenaxis” method because it consists of finding a unique axis in space that reveals the platform attitude and, hence, the integer ambiguities. In practice, it is useful both for initial ambiguity resolution and, during receiver operation, for integrity checking.

In this section, the displacement vector from the initial baseline position is written as δx (3×1) instead of Δx as a reminder that only small angle motion is being considered. The single baseline case is discussed first. A sequence of N measurements from different times is concatenated together to form a matrix, δX , of delta Cartesian positions,

$$\delta X \equiv [\delta x^{(1)} - \bar{\delta x} \quad \delta x^{(2)} - \bar{\delta x} \quad \dots \quad \delta x^{(N)} - \bar{\delta x}], \text{ where } \bar{\delta x} = \frac{1}{N} \sum_{t=1}^N \delta x^{(t)}$$

The mean has been subtracted such that δX represents deviations from a “nominal” baseline location. Figure 3.1.3.1 shows an example of a sequence of δx vectors for an interval of baseline motion. For small angle motion, the vectors $\delta x^{(N)}$ collectively define a plane which is normal to x , the baseline vector. Mathematically, the derivative of the length constraint $x^T x = \text{const}$ is $2\delta x^T x = 0$. Combining all N of the measurements together in a single equation in a similar manner to that in the previous section, the matrix result is $2\delta X^T x = 0$.

Ideally, the null space of δX^T is exactly the unknown baseline vector. Once this null space is established, the baseline vector (subject only to a sign ambiguity) may be adjusted to its known length.




Figure 3.1.3.1: Small Angle Motion

Note that the presence of noise on the differential phase measurements effectively guarantees that δX^T will be full rank. However, it is relatively straightforward to make an appropriate determination of the null space of a matrix by using the singular value decomposition or SVD [see, for example Golub & Van Loan²⁸]. The matrix δX^T is decomposed as $\delta X^T = U\Sigma V^T$, where U and V are orthonormal matrices and Σ is a diagonal matrix of the (positive) singular values of δX^T arranged in decreasing order. Based on the singular values of the resulting matrix, one can establish criteria to define its rank in the presence of measurement noise. The null space of the matrix is simply the last column of V (associated with the smallest singular value). Note that the singular values can also be used to test for small angle motion.

An alternative method of achieving the same result is to form a “covariance matrix” $P = \delta X \delta X^T$ from the Cartesian measurements. The ellipsoids described by the equation $x^T P^{-1} x = \text{const}$ are an approximation for the shape of the collection of Cartesian data points. The best estimate of x is then given by the eigenvector of P associated with the smallest eigenvalue (the “flattest” dimension of the ellipsoid shown in Figure 3.1.3.1). This is exactly equivalent to the null space solution of the SVD since the eigenvalues of P are the squares of the singular values of δX^T .

Note that for a single baseline, two-axis motion (perpendicular to the baseline) is required to establish the baseline normal vector. The next problem is the issue of how to generalize this result to multiple baselines. Note that one could simply neglect the interconnection between the separate baselines and find the null space for each baseline individually. However, this approach discards useful information. The idea behind the methods presented is to take advantage of the *correlation* among motion of separate baselines rigidly constrained to one another.

The result for large angle motion is extended to small angle motion by simply noting that squared terms may be considered second order small. Therefore, in the case of the three baseline case, the $6N \times 9$ system of equations for the large angle case reduces to

$$\begin{bmatrix} \delta X_2^T & \delta X_1^T & 0 \\ \delta X_3^T & 0 & \delta X_1^T \\ 0 & \delta X_3^T & \delta X_2^T \\ 2\delta X_1^T & 0 & 0 \\ 0 & 2\delta X_2^T & 0 \\ 0 & 0 & 2\delta X_3^T \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = 0$$

With 3 non-coplanar baselines and small angle motion about *any* 2 or 3 axes, it is possible to determine attitude unambiguously in a single step. Note that this requirement is less restrictive than that for a single baseline where the 2 axis motion must take place *specifically* about axes perpendicular to the baseline.

Given the appropriate motion, the one dimensional matrix null space (the eigenaxis or single direction in 9 space) exactly yields the 3 Cartesian vectors defining the 3 axis attitude. Applying the SVD, this null space is most closely approximated by the 9th column of the V matrix (that direction associated with the smallest singular value).

The next step is to scale the eigenaxis such that it best matches the known baseline configuration. Consider the 2-norm of the aggregate baseline vector in 9-space, $\begin{bmatrix} x_1^T & x_2^T & x_3^T \end{bmatrix}^T$, given by

$$\left\| \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \right\|_2^2 = \begin{bmatrix} x_1^T & x_2^T & x_3^T \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = x_1^T x_1 + x_2^T x_2 + x_3^T x_3$$

Note that if the baseline vectors are transformed by any proper (det=1) orthonormal 3×3 coordinate transform matrix, the 2-norm is invariant.

$$\begin{aligned}
\left\| \begin{bmatrix} Ax_1 \\ Ax_2 \\ Ax_3 \end{bmatrix} \right\|_2^2 &= \begin{bmatrix} x_1^T A^T & x_2^T A^T & x_3^T A^T \end{bmatrix} \begin{bmatrix} Ax_1 \\ Ax_2 \\ Ax_3 \end{bmatrix} \\
&= x_1^T A^T A x_1 + x_2^T A^T A x_2 + x_3^T A^T A x_3 \\
&= x_1^T x_1 + x_2^T x_2 + x_3^T x_3 \\
&= \left\| \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \right\|_2^2
\end{aligned}$$

Therefore, once the null space has been obtained, it can be normalized and then scaled by the *a priori* 2-norm to bring the baseline dimensions to the correct vicinity.

Then, the only issue is the ambiguity in choosing the sign of the eigenaxis. For three or more non-coplanar baselines, this ambiguity is easily settled. Only one choice will result in the correct “handedness” of the baselines with respect to one another. A simple test is to check the sign of the quantity formed from the cross product of the first two baselines dotted into the third: $\text{sgn}((x_1 \times x_2) \bullet x_3)$. If this result matches that for the *a priori* baseline, the sign of the eigenaxis was chosen correctly.

If the baselines happen to form an orthonormal set, the matrix $[x_1 \ x_2 \ x_3]$ would be exactly the coordinate transformation from the body frame (defined by the baseline set) to the inertial frame in the absence of measurement noise. With measurement error, the elements will still be very close to forming an orthonormal set.

Coplanar Baseline Case

The less desirable coplanar baseline configuration is considered briefly. The eigenaxis formulation for two baselines is

$$\begin{bmatrix} \delta X_2^T & \delta X_1^T \\ 2\delta X_1^T & 0 \\ 0 & 2\delta X_2^T \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = 0$$

Unfortunately, for the two-baseline case and all other coplanar baseline configurations, the eigenaxis sign ambiguity yields an image solution. In most circumstances, it may be possible to reconcile the ambiguity by considering the relation between the hemispherical antenna pattern and the local vertical. Other clues may come from the GPS velocity vector or other external sources.

3.2 Refining the Initial Guess

The initial guess is seldom accurate enough to be used as is for integer resolution. Therefore, iterative refinement of the attitude estimate was implemented as a batch non-linear least-squares fit. The strategy is to estimate the integers as if they were continuous parameters. Following convergence, the continuous parameters are tested to see how near they lie to true integer values. If this difference exceeds a threshold, the solution integrity is suspect, and the process repeats after waiting for newer measurements.

Several approximations to the complete non-linear fit were considered in order to simplify the process. However, the most effective approach is to estimate the attitude at each time tag in the measurement data set while simultaneously estimating the integer offset which applies to each continuous time segment of a carrier phase measurement.

Note that the attitude matrices, A (3×3), corresponding to each time step in the initial guess can be obtained easily by using the method of attitude determination using vector measurements (Wahba's Problem), described in Chapter 4. For each time step, the optimal A is found which best matches the set of body-referenced baselines, b_i , to the set of inertially referenced baseline vectors, x_i .




Figure 3.2.1.1: Flex Matrix Definitions

3.2.1 ESTIMATING PLATFORM STRUCTURAL FLEXIBILITY

In many cases—especially in aircraft applications—the issue of platform flexibility must be considered during the integer resolution process. The difference in relative position due to wing flexure for an airplane on the ground and in flight can amount to a significant fraction of a wavelength. Furthermore, if the aircraft is in a banked turn, the wing flexure can be even more pronounced. Wing flexure is almost double its nominal cruise value when the aircraft is at a 60 deg bank angle. Therefore, it is most prudent to make no assumptions about the state of wing flexure prior to cycle ambiguity resolution.

An additional wing flexure state variable, δf , is introduced into the least-squares fit. The baseline that is used for the fit consists of the sum of the nominal (undistorted) baseline matrix, B_0 , and a flexure matrix, B_{flex} :

$$B = B_0 + B_{flex} \delta f$$

The matrix B ($3 \times m$) = $[b_1 \ b_2 \ \dots \ b_m]$ is a concatenation of the m Cartesian baseline vectors referenced with respect to the body frame. Figure 3.2.1.1 shows an example of the Piper

Dakota baseline vectors. The baselines in this case can be defined to run from the tail antenna to each of the other antennas, numbered from left to right by the same convention that aircraft engines are numbered. For this example, the flexure matrix is

$$B_{flex} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ -1 & 0 & -1 \end{bmatrix}$$

Wing flexure is then measured in wavelengths as a linear perturbation from the unstressed wing locations. Unfortunately, wing flexure does not happen to be observable with single-axis rotations on the ground for aircraft. This is of little consequence, however, because on the ground the wing flexure is always at zero (as defined by the surveyed locations of the GPS antennas at the time of installation). Additional logic must be present to disable estimation of flexure for integer resolution on the ground.

3.2.2 LINEARIZED OBSERVATION MATRIX

To refine the initial guess iteratively, the observation equation relating state estimate variables and measurements is linearized. The attitude component of the state variables consist of perturbations in the vehicle attitude about all three axes for every time sample, ℓ , under consideration. Wing flexure and the integer bias states complete the list of estimated parameters. The resulting linearized equations are

$$\begin{bmatrix} H_1 & 0 & 0 & 0 & f_1 & I \\ 0 & H_2 & & 0 & f_2 & I \\ \vdots & & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & H_N & f_N & I \end{bmatrix} \begin{bmatrix} \delta\theta^{(1)} \\ \delta\theta^{(2)} \\ \vdots \\ \delta\theta^{(N)} \\ \hline \delta f \\ \hline k \end{bmatrix} = \begin{bmatrix} \delta\varphi^{(1)} \\ \delta\varphi^{(2)} \\ \vdots \\ \delta\varphi^{(N)} \end{bmatrix}$$

where each $H_\ell(mn \times 3)$ is the sensitivity matrix of changes in measured differential phase with respect to rotations about each of the three axes of attitude (each row of H is given by $A_\ell \hat{s}_j \times b_i$), each f_ℓ is the sensitivity vector (whose elements are given by $b_{flexi}^T A_\ell \hat{s}_j$) of changes in phase with respect to the wing flexure state, δf , and $I(mn \times mn)$ is the identity matrix. The other state variables are

$$\delta\theta^{(\ell)} = \begin{bmatrix} \delta\theta_x^{(\ell)} \\ \delta\theta_y^{(\ell)} \\ \delta\theta_z^{(\ell)} \end{bmatrix} \text{ and } k(mn \times 1) = \begin{bmatrix} k_{11} \\ k_{12} \\ \vdots \\ k_{mn} \end{bmatrix}$$

which are vectors of small-angle, 3-axis attitude rotation corrections for each time sample, ℓ , and a vector of integer biases, respectively, for all m baselines and n GPS satellite combinations. The right hand side of the matrix equation is a vector of differential phase residuals corrected for line bias, β , such that

$$\delta\varphi^{(\ell)}(mn \times 1) = \begin{bmatrix} \delta\varphi_{11}^{(\ell)} \\ \delta\varphi_{12}^{(\ell)} \\ \vdots \\ \delta\varphi_{mn}^{(\ell)} \end{bmatrix} \text{ where } \delta\varphi_{ij}^{(\ell)} = \Delta\varphi_{ij}^{(\ell)} - b_i^T A_\ell \hat{s}_j - \beta_i$$

where A_ℓ is the current best estimate of the platform attitude for time sample ℓ .

The only problem with this formulation is that for a typical set of platform motion data, the size of this matrix system is overwhelming, especially for the modest computing power available inside the receiver. Considering a scenario with 6 satellites in view on 3 baselines, each H matrix is 18×3 . Combining together perhaps $N=60$ sets of measurements taken over a 30 second time span, the total size of the system is 1080×199 .

The resulting total storage for the matrix is 214,920 single precision floating point numbers or almost 1 MB of RAM.

Fortunately, as shown in the next sub-section, it is possible to take advantage of the sparsity of the matrix to make a drastic reduction in the required storage.

3.2.3 SPARSE CHOLESKY FACTORIZATION

The sparsity of the sensitivity matrix allows the required storage to be minimized. Considering the storage requirements of each block matrix individually, the storage comes out to be $60 \times (3 \times 18 + 18) = 4320$ single precision floating point numbers, or about 17 kB of RAM. Solving the least-squares problem using the normal equations, the intermediate symmetrical system becomes

$$\left[\begin{array}{cccc|cc} H_1^T H_1 & 0 & 0 & 0 & H_1^T f_1 & H_1^T \\ 0 & H_2^T H_2 & & 0 & H_2^T f_2 & H_2^T \\ \vdots & & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & H_N^T H_N & H_N^T f_N & H_N^T \\ \hline f_1^T H_1 & f_2^T H_2 & \cdots & f_N^T H_N & \sum_{\ell=1}^N f_\ell^T f_\ell & \sum_{\ell=1}^N f_\ell^T \\ \hline H_1 & H_2 & \cdots & H_N & \sum_{\ell=1}^N f_\ell & NI \end{array} \right] \begin{bmatrix} \delta\theta^{(1)} \\ \delta\theta^{(2)} \\ \vdots \\ \delta\theta^{(N)} \\ \delta f \\ k \end{bmatrix} = \begin{bmatrix} H_1^T \delta\varphi^{(1)} \\ H_2^T \delta\varphi^{(2)} \\ \vdots \\ H_N^T \delta\varphi^{(N)} \\ \sum_{\ell=1}^N f_\ell^T \delta\varphi^{(\ell)} \\ \sum_{\ell=1}^N \delta\varphi^{(\ell)} \end{bmatrix}$$

This symmetric system is then solved by using a block Cholesky factorization²⁸ of the symmetric matrix,

$$L = \left[\begin{array}{cccc|c} L_1 & 0 & \cdots & 0 & 0 \\ 0 & L_2 & & 0 & 0 \\ \vdots & & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & L_N & 0 \\ \hline G_1 & G_2 & \cdots & G_N & L_K \end{array} \right].$$

The block Cholesky decomposition conserves the sparsity of the original symmetric matrix. Each submatrix, L_ℓ , is a 3×3 lower triangular matrix. Each matrix G_ℓ is $(mn+1) \times 3$. The matrix L_K is a $(mn+1) \times (mn+1)$ lower triangular matrix which takes the positions formerly occupied by the flexure and integer bias sensitivity elements. The vector of state estimate corrections is then obtained through block back substitution of the factored symmetric system, $LL^T x = b$:

$$\begin{bmatrix} L_1 & 0 & \cdots & 0 & 0 \\ 0 & L_2 & & 0 & 0 \\ \vdots & & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & L_N & 0 \\ \hline G_1 & G_2 & \cdots & G_N & L_K \end{bmatrix} \begin{bmatrix} L_1^T & 0 & \cdots & 0 & 0 \\ 0 & L_2^T & & 0 & 0 \\ \vdots & & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & L_N^T & 0 \\ \hline 0 & 0 & \cdots & 0 & 0 \end{bmatrix} \begin{bmatrix} G_1^T \\ G_2^T \\ \vdots \\ G_N^T \\ \hline L_K^T \end{bmatrix} \begin{bmatrix} \delta\theta^{(1)} \\ \delta\theta^{(2)} \\ \vdots \\ \delta\theta^{(N)} \\ \hline \delta f \\ k \end{bmatrix} = \begin{bmatrix} H_1^T \delta\phi^{(1)} \\ H_2^T \delta\phi^{(2)} \\ \vdots \\ H_N^T \delta\phi^{(N)} \\ \hline \sum_{\ell=1}^N f_\ell^T \delta\phi^{(\ell)} \\ \hline \sum_{\ell=1}^N \delta\phi^{(\ell)} \end{bmatrix}$$

The state estimate is then adjusted by the new corrections, and the process is repeated through convergence.

The entire integer solution from start to finish takes less than a second to execute on a 25MHz 486DX computer. Furthermore, total storage requirements are only in the tens of kB of RAM.

3.3 Static Platform Resolution via GPS Satellite Motion

Integer resolution for the static platform is based on the orbital motion of the GPS satellites. As the unit vectors to the GPS satellites rotate with respect to the local horizontal frame, the changes in differential phase measured across the baselines eventually reveals the platform attitude and, hence, the integers.

Initial Guess

Again, the approach is to use a batch least-squares fit, starting out with an initial guess and refining the estimate until convergence. The initial guess may be computed easily by treating each baseline, i , individually and solving for the three Cartesian components of baseline coordinates in the local horizontal frame, x_i (3×1). The resulting matrix equation (which is completely linear) is given as

$$\begin{bmatrix} S_1^T & I \\ S_2^T & I \\ \vdots & \vdots \\ S_N^T & I \end{bmatrix} \begin{bmatrix} x_i \\ k_i \end{bmatrix} = \begin{bmatrix} \Delta\phi_i^{(1)} \\ \Delta\phi_i^{(2)} \\ \vdots \\ \Delta\phi_i^{(N)} \end{bmatrix}$$

where k is a vector ($n \times 1$) of biases with integers for all n satellites in view and the line bias for that baseline all lumped together, and each $\Delta\phi$ is a vector of raw differential phase measurements ($n \times 1$) for all satellites in view.

Once x_i is known for each baseline, the method of attitude determination using vector observations (see Wahba's Problem in Chapter 4) can be used to find the initial guess for the attitude matrix A (3×3) which best rotates the x_i baseline vectors referenced to the local horizontal frame into the b_i vectors (matrix B) referenced to the body frame.

Iterative Refinement

Next, the solution is refined iteratively by linearizing the observation equation in a manner similar to that in Section 3.2.2. However, since the static platform attitude is the same for all time samples, ℓ , the form of the linearized observation equation is

$$\begin{bmatrix} H_1 & I \\ H_2 & I \\ \vdots & \vdots \\ H_N & I \end{bmatrix} \begin{bmatrix} \delta\theta \\ \overline{k} \end{bmatrix} = \begin{bmatrix} \Delta\phi^{(1)} \\ \Delta\phi^{(2)} \\ \vdots \\ \Delta\phi^{(N)} \end{bmatrix}$$

Again, the right hand side is a vector of differential phase residuals corrected for line bias, β , such that the final solution of the vector k yields the integer ambiguities. For non-coplanar baselines, the estimation process usually has enough information to resolve the ambiguities (conservatively) after about 12 minutes of satellite motion—even with only two satellites in view.

3.4 Quasi-Static Integer Resolution

For the *platform motion* and *GPS satellite motion* algorithms, it has been assumed that the motion of either the platform or the GPS satellites is very much faster than the other. The case considered in this section is when the motion of the two is comparable. For non-spinning spacecraft or a boat on a constant heading, this scenario is almost always true. For aircraft, however, it is less important because it arises only during straight and level flight in the absence of turbulence. Achieving steady state implicitly requires platform motion to get there; therefore, the integers must already be known.

The quasi-static case can be treated by applying the same non-linear least-squares fit equations as in sub-section 3.2.2. The only difference is that the time span of consideration has to be increased over that used in the static case of Section 3.3. Because it is implicitly assumed that the platform is not static, additional states must be estimated. Furthermore, it is assumed that the *platform motion* method of cycle resolution has not been successful because insufficient platform motion has occurred.

3.4.1 CYCLE AMBIGUITY RESOLUTION FOR SPACECRAFT

On orbit, geometrical motion occurs from both rotation and translation. Even if the spacecraft attitude is inertially fixed, the line-of-sight vectors to the GPS satellites will rotate due to orbital motion. The integer resolution technique described here (due to Cohen²⁹) is non-linear and, therefore, requires an initial guess at the outset. Integer refinement may then be carried out in a manner similar to the *quasi-static method* described previously.

Initial Guess

One method of arriving at an initial guess is to use the received signal strength as a coarse indicator of attitude. The presence or absence of a GPS signal on hemispherical antennas whose normal vectors are non-collinear may be used to construct a coarse direction vector to the GPS satellite in the body frame.

The coarse direction vector in the body frame, s_B , is constructed from a linear combination of the normal vectors, a (3×1), of the antennas from which the GPS signal strength exceeds some threshold.

$$s_B = \sum_{\ell=1}^{\#ant} \delta_{\ell} a_{\ell}$$

where δ_{ℓ} is 1 if the signal strength exceeds the threshold and 0 otherwise. The relative weighting could also be a continuous function of signal strength based on the known gain pattern of the antenna and *a priori* knowledge of the absolute signal strength. The body frame vectors are normalized and collected so as to minimize Wahba's cost function (see Chapter 4) using the unit vectors to each GPS satellite. The resulting coordinate transformation matrix, A , serves as the initial guess to the solution refinement step.

Solution Refinement

For spacecraft integer resolution, the number of degrees of freedom required for the estimate can be significantly reduced if it is assumed that the spacecraft is rotating at a constant rate over the time span used to resolve the integers. Raw phase measurements are collected and stored for batch processing over an interval of at least a few minutes in duration, depending on antenna geometry and attitude dynamics. The measurements are then used to solve for an initial attitude and a constant attitude rate. The attitude, $A(t)$, is decomposed into the product of a constant attitude matrix and a constant rate rotation matrix

$$A(t) = A(q)A(\omega t)$$

where ω (3×1) is a Cartesian attitude rate in inertial coordinates and q (4×1) is the quaternion representation of the attitude at the reference epoch (time zero).

Based on the initial guess, the solution is refined in batch form to find the optimal solution to the following modified cost function:

$$J(q, \omega, k) = \sum_{i=1}^m \sum_{j=1}^n w_{ij} \int_{t_0}^{t_1} \left[\Delta \phi_{ij}(t) - k_{ij} - b_i^T A(q) A(\omega) \hat{s}_j(t) \right]^2 dt$$

where k_{ij} is the integer ambiguity associated with baseline i and GPS satellite j . Measurement weights, w_{ij} , are provided for generality. Again, the cycle ambiguities are treated as continuous parameters. Since the GPS receiver outputs measurements at discrete times, the time integral in the cost function may be replaced by a summation.

This cost function is complicated somewhat by the fact that over the interval that raw measurement data is being collected, GPS satellites come into and out of view of the

various baselines on the spacecraft. Continuous measurements over a specific interval all share the same integer offset, k .

The cost function is linearized as in sub-section 3.2.2, except that now the state estimate perturbation is given as the sum of a constant component and a component linearly proportional to time:

$$\delta\theta = \begin{bmatrix} \delta\theta_x \\ \delta\theta_y \\ \delta\theta_z \end{bmatrix} + \begin{bmatrix} \delta\omega_x \\ \delta\omega_y \\ \delta\omega_z \end{bmatrix} t$$

Solving for these corrections refines q and ω for the next iteration.

Integer acquisition time on orbit is strongly dependent on antenna placement geometry. *Note that the accelerated time scales for line-of-sight motion encountered in low Earth orbit enhance the rapidity of motion-based integer ambiguity resolution.* Section 6.1 provides a sample of experimental results of this method.

3.5 Fourier Cycle Resolution Method

In this section, a method for solving the integers is presented based on the Fourier transform. Its operational utility is likely to be realized in applications where large arrays of antennas are employed, although it is conceivable that it could be used to aid in analysis of array configurations with fewer elements. The concept is founded on taking a spatial Fourier transform of the magnitude and phase measured at each antenna.

The GPS carrier can be treated as a plane wave propagating through all of space. As an example, the E-field, $E(x,t)$ of such a wave is

$$\mathbf{E}(\mathbf{x}, t) = \text{Re}\{\mathbf{E}_0 e^{-j(\mathbf{k} \cdot \mathbf{x} - \omega t)}\}$$

where \mathbf{E}_0 is a complex vector describing the wave amplitude, polarization, and phase. The wave number vector, $\mathbf{k} = -\frac{\omega}{c} \hat{s}$, points in the direction of propagation.

Consider an abstract antenna treated as a homogeneous continuum which occupies all of space. The complex phase, $\varphi(\mathbf{x})$, due to the plane wave measured at each point in the antenna with respect to a reference oscillator is given by

$$\varphi(\mathbf{x}) = \varphi_0 e^{-j\mathbf{k} \cdot \mathbf{x}}$$

where the antenna polarization, initial phase, and receiver clock bias are collectively lumped into the complex scalar φ_0 . The wave number and, hence, the direction of origin of the GPS satellite, \hat{s} , can be extracted by taking the three-dimensional Fourier transform, $\Phi(\mathbf{k}')$ of the measured complex phase

$$\Phi(\mathbf{k}') = \iiint e^{j\mathbf{k}' \cdot \mathbf{x}} \varphi_0 e^{-j\mathbf{k} \cdot \mathbf{x}} dx dy dz = \varphi_0 \delta(\mathbf{k}' - \mathbf{k})$$

where $\delta(\mathbf{k})$ is the Dirac delta function with the wave number as an argument. For the single GPS frequency, the magnitude of \mathbf{k}' is already given. Therefore the wave number is constrained to lie on the surface of a sphere whose radius is ω/c . The direction of propagation is, therefore, simply that point on the sphere which maximizes the Fourier transform, $\mathbf{k}' = \mathbf{k}$.

The idea is now generalized to a more practical and realistic antenna. Consider an array of three or more point antennas spaced at discrete locations, \mathbf{x}_i . The measured phase can be treated simply as the summation

$$\varphi(\mathbf{x}) = \sum_i \varphi_0 e^{-j\mathbf{k} \cdot \mathbf{x}} \delta(\mathbf{x} - \mathbf{x}_i)$$

Now, applying the Fourier transform,

$$\Phi(\mathbf{k}') = \iiint e^{j\mathbf{k}' \cdot \mathbf{x}} \sum_i \varphi_0 e^{-j\mathbf{k} \cdot \mathbf{x}} \delta(\mathbf{x} - \mathbf{x}_i) d\mathbf{x} = \varphi_0 \sum_i e^{j(\mathbf{k}' - \mathbf{k}) \cdot \mathbf{x}_i}$$

Again, constraining \mathbf{k}' to lie on the surface of a sphere, the magnitude of the Fourier transform is maximized when $\mathbf{k}' = \mathbf{k}$. *The measured line-of-sight vector is that wave number vector which maximizes the magnitude of the spatial Fourier transform of the measured phase at each antenna in an array.* Based on the original measurements, this desired \mathbf{k}' can be written as

$$\max_{\mathbf{k}' \in |\mathbf{k}'| = \omega/c} |\Phi(\mathbf{k}')|, \text{ where } \Phi(\mathbf{k}') = \sum_i e^{j\mathbf{k}' \cdot \mathbf{x}_i} \varphi(\mathbf{x}_i)$$

This conclusion draws heavily from the Fourier relationship between the antenna aperture distribution and the far field pattern³⁰ discussed in Section 5.2. It is exactly analogous to analyzing the far field pattern which would result from rebroadcasting the carrier from each antenna element exactly in the phase that it was received.

The Fourier transform reveals the presence of cycle ambiguities if it turns out that there are several viable maxima of the Fourier magnitude. One would expect to see multiple global maxima for certain antenna configurations when the antenna spacing begins to exceed half a wavelength. If the antenna array is designed correctly (with a sufficient

number of elements spaced correctly for the desired pointing accuracy) there will be no cycle ambiguity at all.

3.6 Self-Survey

A very interesting and useful capability of the attitude determination receiver is that it can perform its own self-survey at installation time. For example, the antennas can be mounted on an aircraft with little attention to their exact locations. The aircraft can be left outside, preferably overnight, and the receiver itself can use GPS survey techniques to measure the precise relative locations of all the antennas as well as the line biases on each antenna cable.

Sequential Estimation

Estimating the Cartesian relative position vector of each baseline is a completely linear problem. Therefore, sequential estimation methods produce quite satisfactory results and require little onboard memory. Because of its superior numerical stability, a sequential square-root information filter was utilized.³¹

The estimator state is stored separately for each baseline, i , as $[U \mid y]_i$, where U is the square root of the state estimate covariance matrix factored as a square upper triangular matrix, and x is the state estimate such that $Ux=y$.

The Cartesian baseline vector, b_i , and the line biases and integer ambiguities lumped into a single bias parameter, k , for the n satellites in view are concatenated into the state variable, $x_i = [b_x \ b_y \ b_z \mid k_1 \ k_2 \ \dots \ k_n]^T$, $(3+n) \times 1$. Each time a new measurement becomes available for baseline i and satellite j , the system of equations is updated to form a revised cost function to be minimized:

$$J_i(x_i) = \left\| \begin{bmatrix} \hat{s}_j^T & e_i^T \\ -\frac{\phantom{\hat{s}_j^T}}{U_i} \end{bmatrix} x_i - \begin{bmatrix} \Delta\varphi_{ij} \\ -\frac{\phantom{\Delta\varphi_{ij}}}{y_i} \end{bmatrix} \right\|_2^2$$

where e_i is a column of zeros with a one in the i th row. This new form is almost upper triangular. Applying an orthogonal sequence of Givens rotations³², the sub diagonal elements shown below are eliminated.

$$J_i(x_i) = \left\| \begin{bmatrix} \text{rectangle} \\ \text{triangle with } \times \text{ marks} \end{bmatrix} x_i - \begin{bmatrix} \text{square} \\ \text{rectangle} \end{bmatrix} \right\|_2^2$$

The result is a new U and y state estimate (denoted by primes) which incorporates the new measurement. The scalar measurement residual, r , can be checked for solution integrity and then discarded.

$$J_i(x_i) = \left\| \begin{bmatrix} U'_i \\ 0 \end{bmatrix} x_i - \begin{bmatrix} y_i \\ r \end{bmatrix} \right\|_2^2 = r^2 + \|U'_i x_i - y_i\|_2^2$$

Additional logic and matrix operations are required to increment and decrement the number of bias states when GPS satellites rise and set or when cycle slips occur. At any given time, the state, x , is an optimal estimate of the baseline position.

Line Bias Estimation

After the bias states have converged sufficiently, they may be decomposed into the cycle ambiguity and line bias components. Once the integers, k_{ij} , are known (conservatively

after about an hour or so of filter operation), each baseline can be reduced to four state variables: three Cartesian position components, and the line bias, β_i . The new state vector, x_i (4×1), is $x_i = [b_x \ b_y \ b_z \ \beta_i]^T$. The simplified cost function is

$$J_i(x_i) = \left\| \begin{bmatrix} \hat{s}_j^T & 1 \\ -U_i \end{bmatrix} x_i - \begin{bmatrix} \delta\phi_{ij} \\ -y_i \end{bmatrix} \right\|_2^2, \text{ where } \delta\phi_{ij} = \Delta\phi_{ij} - k_{ij}.$$

Following completion of the self-survey, the baseline positions and line biases can be stored indefinitely in non-volatile memory on-board the receiver.

Platform Attitude During Self-Survey

One final piece of information is required to complete the self-survey: the attitude of the platform during the survey. The self-survey provides baseline position vectors relative to the local horizontal frame; however, the receiver has no information regarding how this ties to the platform body frame.

For aircraft applications, this final piece of information can be inferred very easily. If it is specified that the fuselage and tail antennas are aligned to the aircraft centerline, the self-survey heading can be determined by the azimuth of this baseline. If it is assumed that the aircraft was parked level during the survey, then pitch and roll can be assumed zero. Chapter 6 provides the self-survey results for the Piper Dakota test aircraft installation.

Chapter 4:

Operational Attitude Solutions

Once the cycle ambiguities are known, differential *phase* measurements can be treated explicitly as differential *range* measurements. The process of attitude determination consists of converting these differential range measurements into attitude solutions. This chapter describes several ways this may be accomplished, including one particularly fast algorithm based on Wahba's Problem, also known as attitude determination using vector observations.

4.1 Non-Linear Least-Squares Solution

An optimal attitude solution for a given set of range measurements, Δr_{ij} , taken at a single epoch for baseline i and satellite j is obtained by minimizing the quadratic cost function

$$J(A) = \sum_{i=1}^m \sum_{j=1}^n w_{ij} (\Delta r_{ij} - b_i^T A \hat{s}_j)^2$$

where b (3×1) is the baseline vector defined with respect to the body frame, \hat{s} (3×1) is the line-of-sight to the GPS satellite given with respect to the inertial frame, and A (3×3), the variable to be used in minimization, is the attitude transformation from the inertial frame

to the body frame. The parameters, w_{ij} , represent optional measurement weighting for generality.

Given a trial attitude matrix, a better estimate may be obtained by linearizing this cost function about the current solution, A_0 , and solving for a correction matrix, δA , of small angle rotations:

$$\delta A(\delta\theta) = I + \Theta^\times$$

where I (3×3) is the identity matrix and $\delta\theta$ (3×1) is a vector of small angle rotations about all three body frame axes

$$\delta\theta = \begin{bmatrix} \delta\theta_x \\ \delta\theta_y \\ \delta\theta_z \end{bmatrix}$$

and Θ^\times (3×3) is the skew symmetric matrix associated with the vector $\delta\theta$:

$$\Theta^\times = \begin{bmatrix} 0 & -\delta\theta_z & \delta\theta_y \\ \delta\theta_z & 0 & -\delta\theta_x \\ -\delta\theta_y & \delta\theta_x & 0 \end{bmatrix}$$

The attitude cost function becomes

$$J(\delta\theta)|_{A_0} = \sum_{i=1}^m \sum_{j=1}^n w_{ij} \left[\Delta r_{ij} - b_i^T (I + \Theta^\times) A_0 \hat{s}_j \right]^2 = \sum_{i=1}^m \sum_{j=1}^n w_{ij} \left(\delta r_{ij} - b_i^T \Theta^\times A_0 \hat{s}_j \right)^2$$

where $\delta r_{ij} = \Delta r_{ij} - b_i^T A_0 \hat{s}_j$. The geometrical model may be rewritten directly in terms of the three attitude correction angles about each axis

$$b_i^T \Theta^\times A_0 \hat{s}_j = -\hat{s}_j^T A_0^T \Theta^\times b_i = \hat{s}_j^T A_0^T B_i^\times \delta\theta$$

The linearized cost function may now be written as

$$J(\delta\theta)|_{A_0} = \left\| W^{1/2} (H\delta\theta - \delta r) \right\|_2^2$$

where δr is the vector formed by stacking all measurements, $W^{1/2}$ is a diagonal matrix formed by the square root of each weighting factor, and H is the observation matrix formed by stacking the measurement geometry for each separate measurement:

$$H = \begin{bmatrix} \vdots \\ \hat{s}_j^T A_0^T B_i^\times \\ \vdots \end{bmatrix}$$

The estimate for A is then refined iteratively until the process converges to the numerical precision of the computer.

Estimation of Aircraft Wing Flexure

A variation of the aircraft attitude problem results for the case where wing flexure is estimated. Depending on the specifics of the aircraft installation, flexure can be significant enough to traverse more than a wavelength of GPS carrier phase. Furthermore, as discussed in Chapter 6, the study of wing flexure itself provides new insight into aircraft dynamics as measured through GPS.

The linearized observation equations are augmented with a new flexure state variable, δf . As described in Section 3.2.1, the platform flexure is constrained to distort the baseline

configuration as specified by the matrix B_{flex} . Recalling that $B = [b_1 \ b_2 \ \dots \ b_m]$, the i th column of B_{flex} gives the direction in which the i th baseline may move, given the undistorted baseline matrix, B_0 , such that

$$B = B_0 + B_{flex} \delta f$$

The revised optimal solution is that point in state space that minimizes the cost function

$$J(A, \delta f) = \sum_{i=1}^m \sum_{j=1}^n w_{ij} \left[\Delta r_{ij} - (b_0 + b_{flex} \delta f)_i^T A \hat{s}_j \right]^2$$

where the notation $()_i$ indicates which baseline is being considered in the body frame.

This cost function is linearized as before, except that now $\delta r_{ij} = \Delta r_{ij} - (b + b_{flex} \delta f_0)_i^T A_0 \hat{s}_j$ and the observation matrix becomes one element wider

$$H_{flex} = \begin{bmatrix} \vdots & \vdots \\ \hat{s}_j^T A_0^T B_i^\times & \hat{s}_j^T A_0^T b_{i \ flex} \\ \vdots & \vdots \end{bmatrix}$$

Angular Velocity Estimation

Carrier phase sensing is inherently a *position* measurement. By taking some sort of derivative, rate information can also be derived through various means. From the discussion in Chapter 2, the differential phase *rate* is already estimated at the signal processing level. Based on the best estimate of the attitude matrix, an entire set of phase rate measurements can then be assembled into a rate estimate about all three body axes, ω (3×1). The final H matrix can be used again for the angular velocity solution to minimize the following cost function for angular velocity

$$J(\omega)|_{A_0} = \|H\omega - \Delta\dot{r}\|_2^2$$

Another option considered in this work was to employ a dynamic model for the vehicle relating the angular rate and the attitude. Then, in return for some additional processing computation burden, an extended Kalman filter could be used in real time. In many cases, a simple double integrator model would suffice. In the presence of more demanding dynamics, a more accurate dynamic model—perhaps including vehicle flexibility—might be more appropriate.

4.2 Wahba's Problem

In the interest of minimizing the processing time required to form a solution from raw differential range measurements, a new approach was developed based on Wahba's Problem³³ (attitude determination using vector observations). There are many efficient algorithms for solving Wahba's problem in the literature, including Markley's Fast Optimal Attitude Matrix algorithm³⁴ and Shuster's QUEST algorithm.³⁵ It is shown here how these solutions can be generalized for attitude determination using GPS. The end result is that attitude solutions can be carried out with unprecedented speed.

The problem of attitude determination using vector observations originates from early work on spacecraft attitude determination with celestial sensors. The problem is to correlate measurements made with respect to the body reference frame of the spacecraft of direction vectors to stars or other celestial bodies with the known, inertially-referenced vectors.

The problem is posed as finding the optimal attitude matrix, A (3×3), that minimizes the cost function

$$J(A) = \sum_{j=1}^n w_j \left| \hat{s}_B^{(j)} - A \hat{s}_I^{(j)} \right|^2$$

where $\hat{s}_B^{(j)}$ is the j th measured direction vector in the vehicle body frame, $\hat{s}_I^{(j)}$ is the known direction vector in the inertial frame, and w_j is an optional measurement weighting factor.

In many respects, attitude determination using GPS is similar to this problem. The idea is to correlate the known direction vectors of the GPS satellites with the incoming planar wave fronts measured with respect to the body frame. The traditional cost function is

$$J(A) = \sum_{i=1}^m \sum_{j=1}^n w_{ij} \left(\Delta r_{ij} - b_i^T A \hat{s}_j \right)^2$$

where Δr_{ij} is the differential range measurement for baseline i and satellite j , b_i is the i th baseline vector, \hat{s}_j is the j th satellite unit vector, and w_{ij} is a weighting factor included for generality.

The baseline and satellite line-of-sight vectors can be concatenated into matrices B ($3 \times m$) and S ($3 \times n$) where

$$B = \begin{bmatrix} b_1 & b_2 & \cdots & b_m \end{bmatrix} \text{ and } S = \begin{bmatrix} \hat{s}_1 & \hat{s}_2 & \cdots & \hat{s}_n \end{bmatrix}$$

Furthermore, differential range measurements can be assembled into a matrix ΔR ($m \times n$)

$$\Delta R = \begin{bmatrix} \Delta r_{11} & \Delta r_{12} & \cdots & \Delta r_{1n} \\ \Delta r_{21} & & & \\ & \ddots & & \vdots \\ \Delta r_{m1} & & \cdots & \Delta r_{mn} \end{bmatrix}$$

In the cases where all the baselines see a given satellite, the cost function can be converted to the matrix form

$$J(A) = \left\| W_B^{1/2} (\Delta R - B^T A S) W_S^{1/2} \right\|_F^2$$

where the Frobenius norm for a matrix, M , is defined as

$$\|M\|_F^2 \equiv \sum_{i,j} |m_{ij}|^2 = \text{tr}(M^T M)$$

The W_B and W_S are applicable to the rows (baselines) and columns (satellites), respectively, of the ΔR matrix. These weighting matrices take the place of the generalized weighting factors, w_{ij} , with the restriction placed on W_B described below. Although it has been assumed here that all antennas share the same field of view of the GPS satellites, this form may be adapted somewhat to include more general cases.

This matrix form can be expanded easily because the square of the Frobenius norm is equivalent to the trace of the matrix times its transpose. The result is

$$J(A) = \text{tr}\left(W_S^{1/2} \Delta R^T W_B \Delta R W_S^{1/2}\right) - 2 \text{tr}\left(W_S^{1/2} \Delta R^T W_B B^T A S W_S^{1/2}\right) + \text{tr}\left(W_S^{1/2} S^T A^T B W_B B^T A S W_S^{1/2}\right)$$

This cost function becomes much simpler if W_B is chosen properly. Such a W_B matrix may be obtained by performing the Singular Value Decomposition (SVD) of B such that

$B = U_B \Sigma_B V_B^T$, where U_B and V_B are orthonormal and Σ_B is a diagonal matrix of *positive* singular values.

For reasons to be discussed shortly, B must be full rank (non-coplanar baselines) to realize the full potential of this fast solution. The appropriate W_B is then

$$W_B = V_B \Sigma_B^{-2} V_B^T$$

where the notation -2 means taking the reciprocal squared of each diagonal element. The resulting first and last terms of the cost function then become independent of attitude and thus drop out. The result is expressed in a form identical to that used to solve Wahba's problem to *maximize* the new cost function

$$J'(A) = \text{tr}(A S W_s \Delta R^T W_B B^T) \equiv \text{tr}(A G^T)$$

where $G \equiv B W_B \Delta R W_s S^T$ and the invariance of the trace to cyclic permutations of its matrix factors has been invoked.

SVD Solution

Markley³⁶ describes how the optimal A can be obtained by performing the SVD of the matrix G (3×3) where

$$G = U_G \Sigma_G V_G^T$$

To ensure that the optimal attitude matrix, A_{opt} , comes out to be proper (right-handed, $\det=1$), the signs of the third column (associated with the smallest singular value) of both

orthonormal matrices U_G and V_G are adjusted to yield the proper matrices U_+ and V_+ . The optimal A (the attitude solution) is then given by

$$A_{opt} = U_+ V_+^T$$

Note that in order to maximize speed, Markley's fast algorithm for computing A_{opt} does not explicitly compute the singular value decomposition.

Quaternion Solution

Davenport³⁷ shows how the cost function may be recast as the eigenvalue problem when the attitude, A , is expressed in terms of the quaternion, q (4×1), through the substitution $A(q) = (q_4^2 - \mathbf{q}^T \mathbf{q})I + 2\mathbf{q}\mathbf{q}^T - 2q_4 Q^\times$, where Q^\times is the skew symmetric matrix

$$Q^\times = \begin{bmatrix} 0 & -q_3 & q_2 \\ q_3 & 0 & -q_1 \\ -q_2 & q_1 & 0 \end{bmatrix} \text{ and } q \equiv \begin{bmatrix} q_1 \\ q_2 \\ q_3 \\ q_4 \end{bmatrix} = \begin{bmatrix} \mathbf{q} \\ q_4 \end{bmatrix}$$

The modified cost function is then expressed as

$$J'(q) = q^T K q$$

where

$$K = \begin{bmatrix} G + G^T - \text{tr}(G)I & z \\ z^T & \text{tr}(G) \end{bmatrix} \text{ and } z \equiv \begin{bmatrix} G_{23} - G_{32} \\ G_{31} - G_{13} \\ G_{12} - G_{21} \end{bmatrix}$$

The optimal q is then the normalized eigenvector ($q^T q = 1$) associated with the maximum eigenvalue of the K matrix. Shuster's QUEST algorithm is a computationally efficient solution using this quaternion approach.

The Rank of the B Matrix

As mentioned previously, the antenna layout has a profound impact on system performance. One interpretation of the measurement process is that the direction vector to the GPS satellite with respect to the body frame is measured directly by the baseline array. Each baseline effectively samples one Cartesian component of the GPS satellite direction vector. Depending on the antenna layout, the baseline vectors may or may not form an orthogonal set.

The weighting matrix, W_B , is a measure of how much the raw measurements have to be scaled to transform the baseline configuration to an equivalent orthonormal basis. For this reason, coplanar baseline configurations are not compatible with the vector observations framework as presented here.

The ideal “balanced” baseline configuration is one for which $BB^T = aI$, where a is a scalar and I is the identity matrix. Figure 4.2.1 shows several examples of such balanced configurations. Another way to express this is that all singular values of B are equal. In this




Figure 4.2.1: Examples of “balanced” antenna configurations.

case (which also corresponds to the case in which reasonable definitions of the attitude dilution of precision are independent of the platform attitude) the Wahba cost function and the traditional cost function are exactly equivalent. The resulting optimal attitude solutions are, therefore, identical.

In the case where this condition is not true, increased solution sensitivity results as the unbalanced baseline configuration is mapped to its orthonormal equivalent. For “flat” baseline configurations (whose height is, for example, one fifth of the width and length dimensions), it is possible for the optimal attitude solutions for the two cost functions presented here to differ by as much as 0.5 deg. So it is with this caveat that the speed benefits should be weighed in cases where the baseline configuration is not balanced.

Fast Solution Results

For comparison, two different solutions to the GPS attitude determination problem were benchmarked against each other. The standard for comparison was the solution to the conventional cost function with a full non-linear least squares fit. The other solution was based on Wahba’s problem.




Figure 4.2.2: Solution Benchmark Results

For relative comparison, the benchmark was run on a 25 MHz 486DX processor. Figure 4.2.2 shows the resulting execution times for a single attitude solution based on four antennas and the given number of satellites in view.

Both methods are iterative. In each case, the logic is set up to iterate until the

machine precision is reached. For the benchmark testing, the iteration count was hard-wired to its average value in normal operation. Therefore, the benchmark solution corresponds to the average time required to achieve the full numerical precision of the machine.

Based on the results shown in Figure 4.2.2, *the solution based on Wahba's Problem is almost an order of magnitude faster than the conventional non-linear least-squares fit.*

4.3 Integrity Checking

Occasionally, the carrier phase tracking will suffer a cycle slip. While these events have proven to be relatively rare during normal operation, they must be readily manageable when they do occur. Cycle slips usually originate from satellites which lie in the periphery of the antenna pattern. Such satellites are either low on the horizon or partially obscured due to steep angle attitude maneuvers. Given the inevitability of occasional cycle slips, several layers of logic were built into the design to safeguard against incorrect solutions.

Each separate phase measurement being tracked has associated with it a logical flag named `INTEGER_VALID`. This flag is set whenever it is believed that the integer component of the measured phase is correct (valid) and reset whenever there is any reason to suspect that the phase might be in error. When attitude solutions are performed, only those measurements for which `INTEGER_VALID` is set are utilized.

Setting `INTEGER_VALID`

The `INTEGER_VALID` flag is globally accessible to all processing layers. While only the top-level software can set `INTEGER_VALID` (such as at start-up when the integers are initially solved), just about any process—from the signal processing level up—can

reset it. Typically, a measurement can be brought off line without interrupting the flow of attitude solutions. If attitude solutions are still being performed, the receiver will make every effort to bring the measurement back on line.

To bring a measurement back on line, an integer adjustment, ΔI , is calculated based on the current attitude solution, A , and the measured phase angle, $\Delta\phi_{ij}$

$$\Delta I_{ij} = \text{round}(\Delta\phi_{ij} - b_i^T A \hat{s}_j)$$

A potentially bad solution is flagged if the argument of the round-off function does not lie near an integer value. Once the adjustment has been executed, the `INTEGER_VALID` flag is automatically set as long as the signal strength from the given antenna and satellite exceeds the preset threshold. This adjustment process is illustrated in the top-level processing diagram given in Figure 1.4.4.

Clearing `INTEGER_VALID`

The first line of defense against cycle slips is placed at the signal processing level. Any number of events at this level can result in `INTEGER_VALID` being cleared. Examples of these include loss of lock detected on the master antenna, master antenna reassignment, signal strength below threshold, aliasing mode detected, differential range exceeding baseline length, or half-cycle slip detected. Detection of any of these events immediately brings the measurement off line.

A final RMS residual check is applied at the conclusion of the measurement solution. If this RMS residual exceeds a preset threshold (usually around 1.5 cm), a cycle slip is suspected. Since the least-squares solution is substantially overdetermined, the integrity algorithm systematically steps through the individual measurement residuals and clears

the INTEGER_VALID flag associated with the outlier—the residual with the largest absolute value. Once a good solution is obtained, the integers of the bad measurements are adjusted and brought back on line.

In almost all cases, the outlier is the measurement with the cycle slip. In the rare cases where there are two simultaneous cycle slips, the outliers are also usually the measurements with the cycle slip. Occasionally, the wrong measurement will be eliminated, but the algorithm provides the isolation process another chance when the subsequent set of measurements come through.

Special attention has been paid to eliminate pathological cases where computation and communication delays interfere with correct operation. For example, the logic has been designed to prevent an integer adjustment from being executed if a potential cycle slip has been detected *during the adjustment calculation*. Furthermore, it suppresses sending integer adjustments more than once. An example of such a scenario is given as follows:

After a successful attitude solution, the logic detects that INTEGER_VALID is clear for a particular measurement. It computes a correction, but transport lags prevent INTEGER_VALID from being set before the next set of measurements arrives. The uncorrected logic will repeat the same correction again; only this time, the adjustment drives the integer away from its correct value. Later, the post-solution residual check will bring the measurement off line again, and the entire process will continue indefinitely. The resulting oscillation is a perfect example of the destabilizing influence of a transport lag on feedback control systems.

Chapter 5:

Antennas and Multipath

Multipath, the unwanted reflections of the GPS signal from the environment surrounding an antenna, is by far the largest error source in attitude determination. While the accuracy of GPS differential range measurements is generally quite good (on the order of 5mm RMS), its ultimate potential is even better. In fact, the ultimate accuracy of GPS may be at least an order of magnitude better, since the null baseline measurement precision is better than 1mm.

The goal of this chapter is to study multipath and explore means by which the accuracy of GPS carrier phase measurements can be improved. Several approaches are considered to overcome multipath:

- Calibrating out repeatable component of multipath
- Rejecting multipath through antenna pattern shaping
- Diluting multipath with larger antenna separation
- Filtering out multipath in static environments
- Processing out multipath with alternate delay-lock loop architectures




Figure 5.1: Geometry of the Near-Infinite Ground Plane

Each tactic has its own range of applicability and effectiveness. The first three methods are emphasized in this work because they address dynamic applications most directly.

The Infinite Ground Plane Experiment

As an illustration of multipath, an interesting study in a near-ideal environment was carried out on a municipal levee in Sunnyvale adjacent to the San Francisco Bay. The water,




Figure 5.2: SV Setting over Still Water

lying perfectly still, takes on the characteristics of a nearly infinite ground plane. Figure 5.1 shows the experiment geometry. The differential path length, Δr , determines whether there will be constructive or destructive interference between the direct and reflected rays. Figure 5.2 shows the received GPS signal amplitude for a satellite setting on the horizon.

It provides a clear indication of the change in path length between the direct ray and its reflection. From the geometry, the small change in differential path length, δr , due to a small angle change in satellite elevation is given by $\delta r = 2h\delta\theta$.

Over the interval shown in Figure 5.2, the satellite elevation angle traversed about 6.5 deg, as it set on the horizon. From the plot, the measured differential path length during the interval is about 6.5λ , or 1.2m. Therefore, the calculated antenna height, h , would appear to be 5.3m. The independently-measured antenna height above the water was roughly 5m. Unfortunately, the multipath observed in practical situations is not so well-behaved. Therefore, many different approaches for dealing with multipath have been considered here.

5.1 Multipath Calibration for Small Satellites

In many applications, especially in aerospace, the multipath environment is highly repeatable. Figure 5.1.1a shows the multipath error (differential range residual) observed between a static pair of antennas on two separate days as a GPS satellite traverses virtually the same arc across the sky. For clarity, the residuals from the second day have been offset. In Figure 5.1.1b, the multipath signatures are differenced, revealing




Figure 5.1.1: Multipath Error and Multipath Repeatability

repeatability within a fraction of a millimeter RMS. In repeatable environments such as this one, it should be possible to calibrate out the effects of multipath and, in real time, recover the inherent sub-millimeter accuracy of the GPS signal.

5.1.1 Theory Of Calibration

For a small spacecraft in low Earth orbit, the only source of multipath is that originating from the spacecraft itself. A wave front originating from a given direction with respect to the body frame will always reflect and diffract off of the various surfaces on the vehicle in exactly the same way. Given such repeatability of the wave front scattering, it should be possible to calibrate out the multipath error as a function of direction of wave front incidence.³⁸

When an antenna is mounted on a vehicle, one may consider the vehicle as part of the antenna. Receivers should have a repeatable response independent of signal modulation since the carrier is reconstructed by mixing the received signal with the original pseudo-random sequence. Furthermore, it is assumed that the principle of superposition applies. Therefore, the correction to the repeatable component of multipath should be independent of both the signal strength and the specific source satellite and mainly dependent on the direction of the signal.

Note that the phase perturbations originating from a variety of sources may be conveniently lumped into the same repeatable category. Specifically, antenna phase center deviation (individual antenna response), inter-antenna interaction, and conventional multipath reflection all combine together to yield a repeatable phase response. In fact it is arguable that these separate phenomena are indeed *indistinguishable* if one simply expands the definition of the antenna to include the host vehicle as well.

It is postulated, therefore, that it should be possible to calibrate out the multipath error for a flight vehicle provided that the only sources of reflections in the immediate vicinity are due to the fixed geometry of the vehicle itself. Note that reflections arriving at the receiver delayed by a difference in path length exceeding one C/A code chip (300 m) will be suppressed by the autocorrelation characteristics of the pseudo-random noise code sequence.

A Model of Multipath

In the ideal point model of a pair of GPS antennas, the measured phase, $\Delta\varphi_{ideal}$, is given by the projection of the baseline separating the two antennas, x , into the line of sight to the satellite, \hat{s} . An additive line bias, β , is added to handle the cable length offset and other fixed delays.

$$\Delta\varphi_{ideal} = \hat{s}^T x + \beta$$

Added to the ideal point model, the measured differential phase, $\Delta\varphi_{measured}$, is perturbed by the repeatable multipath component of phase error, $\delta\varphi(\theta, \phi)$. The multipath perturbation is assumed to be a single valued function of the angular position of the satellite (θ, ϕ) in spherical coordinates with respect to the local body frame of the vehicle. Multipath error not corrected in the model and receiver random noise are lumped into the remaining error, $\Delta\varphi_n$. The measured differential phase is then

$$\Delta\varphi_{measured} = \Delta\varphi_{ideal} + \delta\varphi(\theta, \phi) + \Delta\varphi_n$$

This formulation departs from the traditional treatment of multipath error as a random process. Although its signature is typically modeled as zero-mean and band limited,³⁹ in

fact, multipath from the vehicle itself is *not* a random process. Much of the desired information in the signal is *still contained* in the measurements perturbed by multipath.

Spherical Harmonic Model

Although there are a number of ways to model the repeatable characteristic, spherical harmonics were chosen as the calibration function. This form of representation is desirable because

- Few coefficients are required to represent the smoothly varying function
- The set of body frame direction vectors covers the spherical domain
- For calibration, the representation must “interpolate” between satellite data arcs
- There is no projection distortion

The functional form of the spherical harmonic representation is

$$\delta\varphi(\theta, \phi) = \sum_{\ell=1}^n \left[J_{\ell} P_{\ell}(\cos \theta) + \sum_{m=1}^{\ell} P_{\ell m}(\cos \theta) (C_{\ell m} \cos m\phi + S_{\ell m} \sin m\phi) \right]$$

For an n th order model of the multipath error, the calibration coefficients, $J_{\ell m}$, $C_{\ell m}$, and $S_{\ell m}$, describe an approximation to $\delta\varphi(\theta, \phi)$. In the expression, $P_{\ell m}(x)$ functions are Legendre polynomials. Thus each baseline pair of antennas has a set of coefficients to describe the vehicle multipath specific to that pair. The capacity of this functional form to follow variations of increasing complexity is governed by the model order, n .

5.1.2 EXPERIMENTAL RESULTS

To demonstrate the feasibility of the multipath calibration technique, differential phase measurements were made between a single pair of antennas over a large range of satellite directions. To build up a multipath model, one accumulates measurements from signals

ideally originating from every possible direction in the vehicle body frame. In practice, measurements may consist of “slices” of the celestial sphere to create a mesh of sufficient density to solve for the coefficients of the spherical harmonic model.

Experimental Antenna Configuration

In lieu of actual calibrations of vehicles in flight, a ground-based analog was developed. By using a large (3λ radius) horizontal ground plane (shown in Figure 5.1.2.1) it was attempted to attenuate reflections coming up from below the antenna. To the extent that this was successful, the antenna was rendered immune to its multipath environment for azimuthal rotations. For the experiments performed here, the antenna was placed on the roof of a building such that it was the highest object for at least 500m in every direction. In this respect, the analogy corresponds to a flight vehicle operating well away from other reflective objects. With the ground-plane and antenna pair alone, the residual “noise floor” due to multipath was 2.6mm RMS relative position error.

To create intentional multipath errors, a rectangular sheet metal reflector was mounted on the ground plane as shown in the figure. The analog to attitude motion of a flight vehicle




Figure 5.1.2.1: Multipath Calibration Experimental Setup




Figure 5.1.2.2: Tracks of GPS Satellites Across the Celestial Hemisphere

is *rotation of the entire ensemble* (consisting of the antennas, ground plane, and reflector) about the zenith line by various azimuth angles, α , with respect to true north. The observed result is that the perturbations in carrier phase relative to the ensemble body frame *moved with the body* in a highly repeatable way about its azimuthal rotations.

GPS Satellite Sky Coverage

To build up a model of the phase perturbations, the antenna ensemble was left static for 24 hours in various azimuthal positions. The orbital motion of the GPS satellites in the sky naturally varied the directions of incidence of the GPS plane wave fronts. Over a 24 hour period, the repeating ground track of the GPS constellation traces out the arcs in the sky shown in Figure 5.1.2.2a.

The spherical harmonic model also acts as an angular (spatial) filter. The higher the model order, n , the greater the angular detail (angular or spatial frequency) the model can approximate. A rough rule of thumb is that the smallest variation with angle is roughly $180 \text{ deg}/n$. Therefore, gaps in the regions of the sky shown in Fig. 5.1.2.2a (excluding the region in the north where no satellites fly) preclude a higher order model than about 5.

For higher order, n , the best fit would not be accurate in the uncovered regions since there would be no measurements to shape the estimate.

To circumvent this limitation, measurements were collected with the antenna rotated among a total of five azimuthal positions, indicated in Figure 5.1.2.2a. Figure 5.1.2.2b shows the resulting collection of data points consisting of positions 1-4, each separated by 90 deg. Note that the areas which remain uncovered are much smaller, suitable to increasing the model order to around $n=18$. To reduce the size of the data set, the measurements were spatially prefiltered before the spherical harmonic fit.

Experimental Data

For a given 24 hour data run, the phase residuals were obtained by performing a least-squares fit of the antenna point model to all the measurement data, solving for x and β . The resulting best-fit baseline solution, x , is defined as the *differential phase center* of the antenna pair. Phase residuals—constructed by subtracting out the resulting $\Delta\phi_{ideal}$ from the actual measurement—are exceptionally repeatable, as shown in Figure 5.1.1 and as




Figure 5.1.2.3: Repeatability of Residuals among Different Space Vehicles

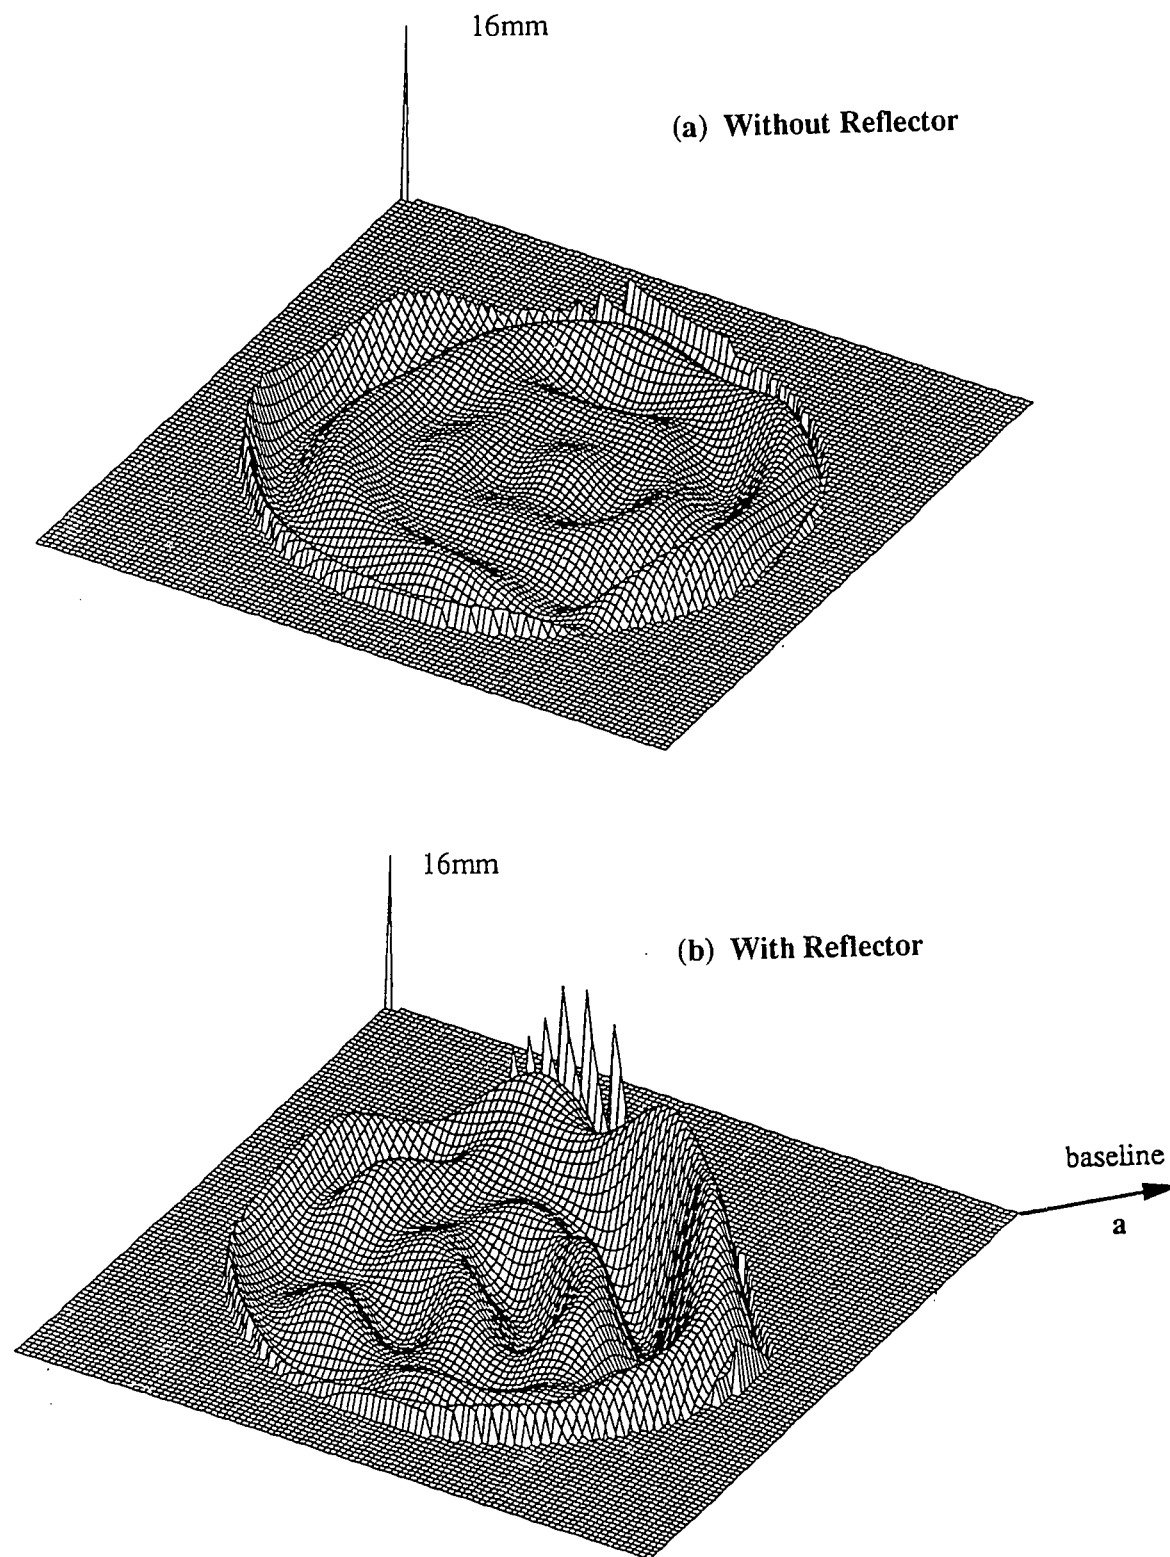


Figure 5.1.2.4: Best Fit Multipath Calibration (8th Order Model)

documented previously.^{39,40} As each satellite exactly retraces its sky track each day, the signature in the phase residuals is almost identical. Furthermore, the signatures of *different satellites* following closely spaced arcs at different times of day are also exceptionally repeatable as demonstrated in Figure 5.1.2.3. Note that SV20 is setting over the interval while SV13 is rising. Data points on the left are closer to the horizon. This observation supports the conclusion that multipath perturbations are largely independent of carrier modulation code and Doppler shift.

For comparison, 8th order fits of the spherical harmonic model were performed on the phase measurement residuals. The residual fits are shown in Figure 5.1.2.4 both *without* (a) and *with* (b) the rectangular reflector affixed to the antenna.

These mesh plots depict the residual error (represented as height) for the antenna pair as a function of direction in the sky. This representation could be overlaid on Figure 5.1.2.1, since the horizontal dimensions correspond to direction in the antenna body frame. The phase correction ranges between $\pm 4.8\text{mm}$ without the reflector and from -20 to $+14\text{mm}$ with the reflector. The spikes at the rear corner serve as a reference scale and represent a phase error of 16mm (30° of phase). The portion of the sky blocked by the reflector is masked off. Two 24 hour runs at azimuthal positions 1 and 2 were used for the fit without the reflector. Four 24 hour runs at azimuthal positions 1-4 were used for the fit with the reflector.

To illustrate the effectiveness of the calibration, Figure 5.1.2.5 shows the measured phase residuals from a completely *independent* set of data taken at position 5 against the calibration accumulated over positions 1-4. For reference, Figure 5.1.2.6 shows which region in the antenna body frame produced each residual plot. The circles correspond to the abscissa grid lines in the referenced figures. Also shown is a simple ray-tracing




Figure 5.1.2.5: Calibration Results with Reflector

model of the regions in the sky over which the reflector would influence the phase measurements. However, these regions are only for reference since diffraction effects severely limit the validity of the ray-tracing model. The reflector is less than three wavelengths in dimension.

Intentionally introducing multipath with the reflector doubled the uncorrected error residuals. This disturbance raised the residuals above the 2.6mm RMS minimal multipath “noise floor” to more than 5.2mm RMS. Applying the calibration reduced the error to below 3.2mm RMS, almost fully recovering the original accuracy achieved without the disturbance.




Figure 5.1.2.6: Sky Map of Phase Residuals




Figure 5.1.2.7: RMS Error Remaining after Calibration

The residuals resulting from the independent data set are plotted against model order in Figure 5.1.2.7. In general, the model order n required to handle a vehicle dimension, L , can be loosely approximated as $2L/\lambda$ (around $n=12$ for the reflector dimension used here). Reflections from objects farther away than the antenna contribute to the higher frequency components of the “noise floor” of the multipath signature.

Operating in space away from all other reflecting objects, it is expected that the noise floor will be substantially smaller. The RMS error is expected to decrease linearly until it is much less than 1mm at about model order 18.

Extrapolation of Results to Flight Vehicle Performance

The inherent measurement accuracy of the receiver used for these experiments is at least an order of magnitude better than the error due to multipath. If calibration is applied, multipath need not limit the angular accuracy that can be achieved. Instead, the GPS carrier-to-noise ratio ultimately drives the accuracy, as described in the discussion on the *Performance Envelope* of attitude determination in Chapter 1.

5.1.3 PRACTICAL IMPLICATIONS

The calibration technique outlined here is essentially a correction to the ideal phase based on the point model of an antenna. One should note that the correction applied to each signal is a function of satellite direction in the body frame—a quantity which is not known until the solution is performed. Fortunately, since the coarse attitude without the correction is already known, little iteration is generally required. Furthermore, if the vehicle attitude is changing slowly, the corrections need not be updated frequently. Note that multiple baselines may be accommodated by creating a separate calibration map for each antenna pair.

The initial calibration to determine the coefficients before flight could be performed in a number of ways, one of which is with an anechoic chamber. Where practical, the calibration could also be performed through simulation. In some circumstances, in-flight calibration is possible if some external truth reference, such as an IMU, is used for initial testing. Vehicle articulation could constitute an error source with this technique, although it could be accounted for by adding additional dimensions to the calibration function.


An important practical consideration associated with the technique is the model order, n . The number of coefficients for an n th order model is $(n+1)^2$. Storage and computation

requirements mount quickly as model complexity is increased. Recall, however, that the model order is roughly proportional to the vehicle dimension. For small vehicle dimensions (on the order of a few wavelengths of the GPS signal) the calibration model should be extremely effective and efficient in improving accuracy. For larger vehicles (above ~2 meters), one may choose to exploit the larger vehicle dimension and employ a longer baseline length. In either case, the multipath error is attenuated.

5.2 Antenna Pattern Shaping

Antenna pattern shaping is studied as a way to reject multipath altogether. In this section, conceptual designs of two new antennas are proposed. A set of experiments supports the theoretical framework of each of these techniques.

Antenna pattern shaping is another possible method of curbing multipath. As shown in Figure 5.2.1, the idea is to design the antenna pattern such that direct rays from the GPS satellite are accepted and multipath reflections by surrounding objects are rejected.



Note that antenna pattern shaping can also include polarization parameters depending on the composition of the reflective materials. The GPS carrier is right-hand circularly polarized. Upon reflection from a metallic surface, the polarization is reversed. Therefore, in certain specialized circumstances, the use of polarization could be used to block the reflected signal.

In practice it is quite difficult to construct the sharp cutoff in the antenna pattern shown in Figure 5.2.1. The shape of the antenna pattern is governed by the Fourier transform relationship between the far field pattern, $P(n,m)$, given as functions of the direction cosines, and the aperture distribution, $E(\xi,\eta)$, of the antenna⁴¹ given as a function of the local antenna position coordinates:

$$P(n,m) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} E(\xi,\eta) e^{j2\pi(m\xi+n\eta)} d\xi d\eta$$

In principle, any desired far field pattern may be engineered by ensuring that the corresponding aperture distribution exists at the antenna. Most omnidirectional antennas (since they are spatially small with respect the carrier wavelength) have a characteristically slow cutoff with angle near the horizon. This slow cutoff serves to worsen multipath for incident rays originating from low elevation angles.

As shown in Figure 5.2.2, the measured phase may be visualized as the sum of the direct phase and the resultant of all the reflected rays entering the antenna.

The relative amplitude attenuation for the multipath error is roughly the square root of the relative antenna gain between the direct and reflected rays.




Figure 5.2.2: Multipath Phasor Interpretation

5.2.1 HELICAL ANTENNA EXPERIMENT

As a qualitative example of the influence of antenna pattern shaping on multipath rejection, plots of L1 carrier amplitude versus time for two antennas of different directivity are shown in Figure 5.2.1.1. The satellite is passing within 5 degrees of zenith. As suggested in the figure, the signal strength from an antenna may be just as effective an indicator of multipath sensitivity as differential phase. The first antenna is an omnidirectional patch antenna. The second is a 10 turn helical antenna (of greater directivity). As a basis for comparison, the signal amplitude is normalized by its mean.




Figure 5.2.1.1: Directivity and Multipath

5.2.2 SCREEN WIRE FENCE PATTERN SHAPING EXPERIMENT

In another experiment, a large screen wire mesh was built around a pair of antennas, as shown in Figure 5.2.2.1. The differential phase measured inside this mesh is shown in Figure 5.2.2.2a. As a basis for comparison, measurements made between another pair of identical antennas equally spaced is shown in Figure 5.2.2.2b. However, in this case, no screen wire mesh is used; only 1.5λ radius ground planes are used. What is interesting




Figure 5.2.2.1: Antenna Pattern Shaping Experiment

about the measurements in (a) is the complete absence of high frequency components with respect to those in (b). The higher spatial frequencies—normally attributable to reflective





Figure 5.2.2.2: Screen Wire Mesh Experiment

objects situated much farther away from the antenna—are greatly attenuated. This crude antenna pattern shaping experiment demonstrates the feasibility of the pattern shaping approach.

5.2.3 ANTENNA PATTERN SHAPING THROUGH PHASED ARRAYS

A conceptual design of a low-multipath antenna is presented here. The Fourier relationship between the antenna pattern and the aperture distribution implies that any desired antenna pattern may be achieved by setting up the appropriate aperture distribution. A planar phased array of microstrip antennas several wavelengths in diameter is shown in Figure 5.2.3.1. The amplitude and phase of each element could be designed to suppress backlobes and sharpen the antenna pattern cutoff, greatly improving the phase measurement accuracy.




Figure 5.2.3.1: Pattern Shaping Antenna Array

Arrays of microstrip patch antennas have already demonstrated sidelobe suppression exceeding 35 dB.⁴² This level of antenna performance brings potential relative ranging errors down to the level of about 0.5mm.

Antenna pattern shaping could be conceived as being somewhat analogous to filter design. The multipath-free antenna pattern could be specified in terms of the beam width, maximum sidelobe level, and cutoff range, as shown in Figure 5.2.3.2. This specification

indirectly dictates the size of the array and the number of radiating elements required for synthesis. A satellite is not used if it falls above the minimum antenna boresight offset of the cutoff range. Similarly, the in-




Figure 5.2.3.2: Antenna Pattern Shaping Specifications

stallation is designed such that there is no reflective object lying within the maximum boresight offset of the cutoff range. The entire array could be manufactured as one monolithic unit.

5.2.4 HIGH PERFORMANCE PHASED ARRAY

The ultimate application of antenna pattern shaping comes in the conceptual design of “smart skin” for a launch vehicle or high performance aircraft. As shown in Figure 5.2.4.1, the composite materials from the vehicle structure could serve as the dielectric medium for the microstrip antennas. The phased array might effectively be woven into the composite structure itself.

As shown in Figure 5.2.4.2, this phased array can place an individual spot beam on each GPS satellite separately, simultaneously achieving higher accuracy (through rejection of multipath via antenna pattern shaping) and higher bandwidth (through the increased carrier-to-noise ratio provided by the spot beam).




Figure 5.2.4.1: Composite-Embedded GPS Antennas: "Smart Skin"

In real time, Fourier transforms of the instantaneous array aperture distribution are used to solve for the direction cosines of the GPS satellites in the body frame of the launch vehicle (see Fourier Cycle Resolution in Section 3.5). Attitude determination using vector observations (Wahba's Problem in Chapter 4) is used to convert the direction




Figure 5.2.4.2: High Performance GPS Attitude

cosine output into attitude solutions at the highest possible output rate. This phased array antenna sets the stage for the ultimate high performance application of GPS attitude determination: *closed-loop attitude and position control of a launch vehicle based exclusively on GPS.*

5.3 Increasing Baseline Length

Diluting multipath is based on the idea that differential multipath error is largely independent of antenna spacing. Where the dimensions of the platform permit it, using the antenna separation to increase accuracy is quite attractive. The angular pointing error, σ_θ , is given in terms of the relative ranging error, σ_r , and the baseline length, L , as

$$\sigma_\theta = \frac{\sigma_r}{L}$$

Even with standard omnidirectional patch antennas in difficult multipath environments (relative range errors up to 5 mm RMS), antenna separation on the order of meters is sufficient to achieve pointing accuracy at the milliradian level.

One important consideration, however, is that longer baselines tend to compound the problem of cycle ambiguity resolution. As baseline length increases, there are more possible integer combinations to consider. However, the motion-based cycle resolution algorithms presented in Chapter 3 are advocated precisely because they work independently of the baseline length. Therefore, baselines may be made as long as reasonably practical.

5.4 Multipath Filtering

The filtering of phase measurements in static environments is the last option considered for dealing with multipath. Certainly, straight averaging is the simplest means of improving accuracy. More elaborate filtering schemes are certainly possible as well. Filtering was not pursued directly because the future aerospace applications of attitude determination are almost exclusively dynamic.

An extension of filtering for the dynamic case has been proposed;⁴³ however, it requires an inertial measurement unit. The idea is to use the gyro-based attitude measurements to average out the GPS attitude measurements. In the presence of dynamics, the hope is that the changes in geometry will serve to “whiten” and randomize the measured multipath signature, thereby rendering it more amenable to standard filtering techniques.

5.5 Delay Lock Architecture

Multipath can also be countered at the signal processing level by structuring the delay lock loop more progressively. Hagerman⁴⁴ presents a thorough analysis of the effects of multipath on GPS delay lock loop operation. Along these lines, experiments and designs employing a narrower spacing between the early and late correlators have yielded superior multipath rejection.⁴⁵

Chapter 6:

Flight Experiments

Flight tests of the attitude determination receiver on a Piper Dakota have demonstrated the full operational effectiveness of GPS attitude for aircraft. Flight tests on two spacecraft, APEX and RADCAL, are scheduled for launch into low Earth orbit in May, 1993. This chapter begins with a description of the initial ground tests used for pre-flight validation of the receiver, followed by a description of the current and future flight experiments.

6.1 Ground Validation Tests

For ground validation, a 20 ft cruciform steel array was built to match the geometry of a small aircraft. This array was used for both static and dynamic tests and evaluation in preparation for the test flights. For experiments to validate cycle ambiguity resolution for space missions, a smaller 1m orthogonal array was constructed.

6.1.1 STATIC ACCURACY TESTS WITH A THEODOLITE

In a likeness of the Piper Dakota, the 20 ft steel array is shown schematically in Figure 6.1.1.1. The positioning of each antenna corresponds to antennas mounted on the tail, fuselage, and wing tips of the Piper Dakota. Figure 6.1.1.2 shows a photograph of the

steel array. The static tests were intended to address a realistic installation of the attitude determination receiver; no ground planes were used for the antennas, and the tests were conducted on a rooftop surrounded by reflective objects.




Figure 6.1.1.1: Steel Rooftop Array

The attitude of the array was established independently by using a theodolite to sight along the tail-fuselage baseline.⁴⁶ This bearing was then compared to that of Polaris. The results of these trials are shown in Figure 6.1.1.3. The plot shows the GPS array azimuth solution output at a rate of one independent point per second. The resulting standard deviation is 0.0653 deg. Two optical sightings on Polaris (estimated to be good to 0.0007 deg) are included in the plot. These differ from the mean GPS azimuth solution by 0.026 deg.



Figure 6.1.1.2: Steel Rooftop Array




Figure 6.1.1.3: Static Pointing Accuracy Test Results

The baseline length from tail to fuselage was about 3m, and the wingspan of the steel array was about 6m. Given these dimensions, the absolute accuracy of static operation is




Figure 6.1.2.1: Displacement Vectors in Horizontal Plane

exactly as expected. In aircraft applications, absolute pointing accuracy better than 0.1 deg during standard operation is expected (without calibration).

6.1.2 PREFLIGHT TRIALS WITH PIVOTING STEEL ROOF ARRAY

Before the flight tests on board real aircraft, the steel roof array was mounted on the pivot stand shown in the background of Figure 6.1.1.2. This setup provided a test bed for testing all aspects of receiver operation, including cycle ambiguity resolution. Qualitative dynamic testing was also performed on this facility to demonstrate extreme attitude rates of several radians per second on a 6m baseline without loss of lock.

Figure 6.1.2.1 shows displacement vectors (introduced in Section 3.1) derived from actual phase measurements. In this case, the steel roof array was rotated about the vertical axis by about 90 deg. Displacement vectors from all three baselines are shown. Further processing converted these displacement vectors into an initial attitude guess for cycle ambiguity resolution.

6.1.3 TESTING CYCLE AMBIGUITY RESOLUTION FOR SPACE

Founded on the cycle ambiguity resolution method described in Section 3.4, a ground-based experiment was performed with the antennas arranged in the smaller orthogonal baseline arrangement shown in Figure 6.1.3.1. The baseline length was $\sim 4\lambda$. For this *static* test, the Earth's rotation provided the attitude rate for the antenna configuration.

A set of completely *raw and unprocessed* measurements




Figure 6.1.3.1: Orthogonal Baseline Configuration




Figure 6.1.3.2: Raw, Unprocessed Phase Measurements

taken over a 20 hour interval for one of the baselines is shown in Figure 6.1.3.2. Line segments start and end when GPS satellites rise and set. Each line segment is offset by its heretofore unknown integer ambiguity. Some multipath is visible, especially at the end points of each line segment, where the satellite is low on the horizon and in the periphery of the antenna pattern. Occasional cycle slips are visible but only for weak, low-elevation satellites. Cycle slips are easily detected by the receiver signal processor or solution integrity logic.

Figure 6.1.3.3 shows the estimated integer states for all baselines as a function of the fit interval duration. The independent variable provides the time span covered by the phase data used for the batch least-squares fit. On Earth, the solution exhibits a comfortable margin against false integer estimates after 30 minutes. On orbit, resolution of integers will be much quicker, occurring within just a few minutes. As an unusual adjunct not




Figure 6.1.3.3: Integer Solution History

ordinarily associated with conventional static GPS survey, the Earth's rotation rate was measured to within 7%.

6.2 Aircraft Experiments

The receiver was flown on a Piper Dakota for operational tests and experiments in aircraft dynamic model system identification. In December, 1992, the receiver is scheduled to fly on the NASA Ames King Air 200 twin turbo-prop transport aircraft and compared with the output of an inertial measurement unit.

6.2.1 PIPER DAKOTA FLIGHT TESTS

The first flight tests of the new receiver were performed on a Piper Dakota single-engine aircraft.⁴⁷ The results from this series of flight tests were summarized in *Aviation Week*



Figure 6.2.1.1: Piper Dakota with GPS

and Space Technology.⁴⁸ Figure 6.2.1.1 shows the aircraft in flight. Four microstrip patch antennas were mounted on the airframe, located on the fuselage, tail, and wing tips. Figure 6.2.1.2 shows a close-up of a wing-mounted antenna.

The installation static self-survey was conducted at Palo Alto airport during an overnight data collection run. The data was processed into reference baseline vectors and line bi-ases using the filter described in Section 3.6. Table 6.2.1.1 provides the result of this self-survey. As standard practice for aircraft installation, the nominal azimuth was defined to be that of the baseline from the tail to the fuselage antenna (which was




Figure 6.2.1.2: GPS Antenna on Right Wing

mounted on the centerline of the aircraft). The nominal pitch and roll were defined to be zero.

For the flight tests, the carrier phase cycle ambiguities were resolved by using the platform motion approach from Chapter 3. Throughout the banks and turns of normal flight, the GPS satellites in view are frequently obscured. Each time a satellite becomes visible, its integers are assigned based on the current attitude solution that has been computed from the other satellites already in view. For steep angle maneuvers, this “integer hand-off” from satellite to satellite results in seamless operation under dynamics. *Throughout normal operation during the test flights, integer lock was never lost.* Cycle slips that occurred during the maneuvers were instantly detected and corrected.

Because the design of the receiver is based on a common local oscillator, it is possible to perform a complete attitude solution with only two GPS satellites. The payoff is that the receiver is able to hold on to attitude and integer lock continuously, even during steep angle maneuvers when some satellites become obscured. Although position solutions may be interrupted momentarily, attitude output is unaffected. During the flight tests, the receiver could routinely handle full 360 deg turns with a 45 deg bank angle (the highest angle tested) without losing attitude lock.

Table 6.2.1.1: Piper Dakota Self-Survey Results

Self Surveyed Azimuth: 85.83 deg; All state estimates are in units of L1 cycles.

Baseline from Tail:	to Left Wing	to Fuselage	to Right Wing
East Component	19.373	22.171	23.551
North Component	29.345	1.617	-26.128
Up Component	-5.152	-2.577	-5.547
Line Bias	0.386	0.398	0.639

As described in Chapter 4, each attitude solution is comprised of a four parameter state that includes heading, pitch, roll, and instantaneous wing flex. For reference, a diagram of the aircraft body axes, GPS antenna placement, and wing flex states is shown in Figure 6.2.1.3.




Figure 6.2.1.3: Aircraft Geometry

The capacity to estimate wing flexure is a remarkable facet of the attitude determination problem. Not only does including this state improve the accuracy of the attitude solution (especially in turbulence), some interesting applications result. Most importantly, it provides a real-time indication of wing loading. By scaling the units of measurement, wing flexure can be converted to g 's of vertical acceleration. The result is that *the structural deformation of the aircraft under stress as measured by GPS can be used as an accelerometer!*

Wing flex during landing is shown in Figure 6.2.1.4. At the left of the plot, the aircraft is still airborne and undergoing maneuvers that visibly strain the wings. At touchdown, the airspeed and angle of attack drop, and the wings return to their nominal unsupported displacement.




Figure 6.2.1.4: Landing at Half Moon Bay

The aircraft finally slows to taxi speed at the far right.

During parabolic low- g maneuvers in the Piper test flights, the feeling of weightlessness experienced by the aircraft occupants exactly followed the GPS wing flexure readout.

Concurrent Kinematic Positioning

As part of this research, the attitude receiver was modified to measure carrier beat phase for precision kinematic positioning (running concurrently with attitude solutions). These precise position solutions provide important insight into the aircraft dynamic model. In the future, it is anticipated that this capability will also play an important role in the precision landing problem. Note that the TANS was never designed to perform precise kinematic positioning using carrier phase. However, static and dynamic trials have successfully validated this mode of operation.

During many of the test flights, a base station receiver was set up at Stanford to take reference carrier phase ranging data. The reference measurements were used to obtain precise position solutions in post-processing as part of the dynamic response testing. Appendix C outlines the processing steps that were taken for kinematic positioning. An example of the results are shown for the phugoid mode in Figure 6.2.1.6 of the next section.

Aircraft Natural Modes

In pursuit of safer air travel and more efficient autopilots, an investigation of the aircraft natural modes was undertaken. The availability of high fidelity measurements of attitude and position from GPS offers the unique possibility of easily performing aircraft dynamic model estimation. Given the relatively low cost of receiver equipment, it appears that even continuous, real-time estimates of dynamic parameters is feasible.

In the Piper flight tests, position and attitude data were collected in response to impulsive inputs to the elevator and rudder. After the impulse was applied, the controls were left free.




Figure 6.2.1.5: Short Period Mode Response

On the Piper, the short period transient is over after only a couple of seconds. Figure 6.2.1.5 shows an example of the short period mode. The 10 Hz output rate of attitude solutions is critical. If the output rate were much slower, the receiver might miss much of the information in the transient.

The period of the phugoid mode for the Piper at nominal cruising speed is about half a minute. Figure 6.2.1.6 shows a sample of the phugoid mode data. The first plot shows the pitch angle of the aircraft plotted versus time. Kinematic positioning measurements (developed in the previous sub-section) reveal the relative altitude and along track position. The along track relative position was obtained by subtracting out the constant velocity component of position. Last, the instantaneous wing flex is shown with a precision of 1.4mm RMS. This figure shows the subtle aeroelastic strain on the aircraft structure induced by the natural dynamics of the phugoid mode.

Kinematic positioning enables the construction of the relative position plot shown in Figure 6.2.1.7. This spiral trajectory is the view that would be seen from another aircraft off to the side, flying parallel, straight, and level.

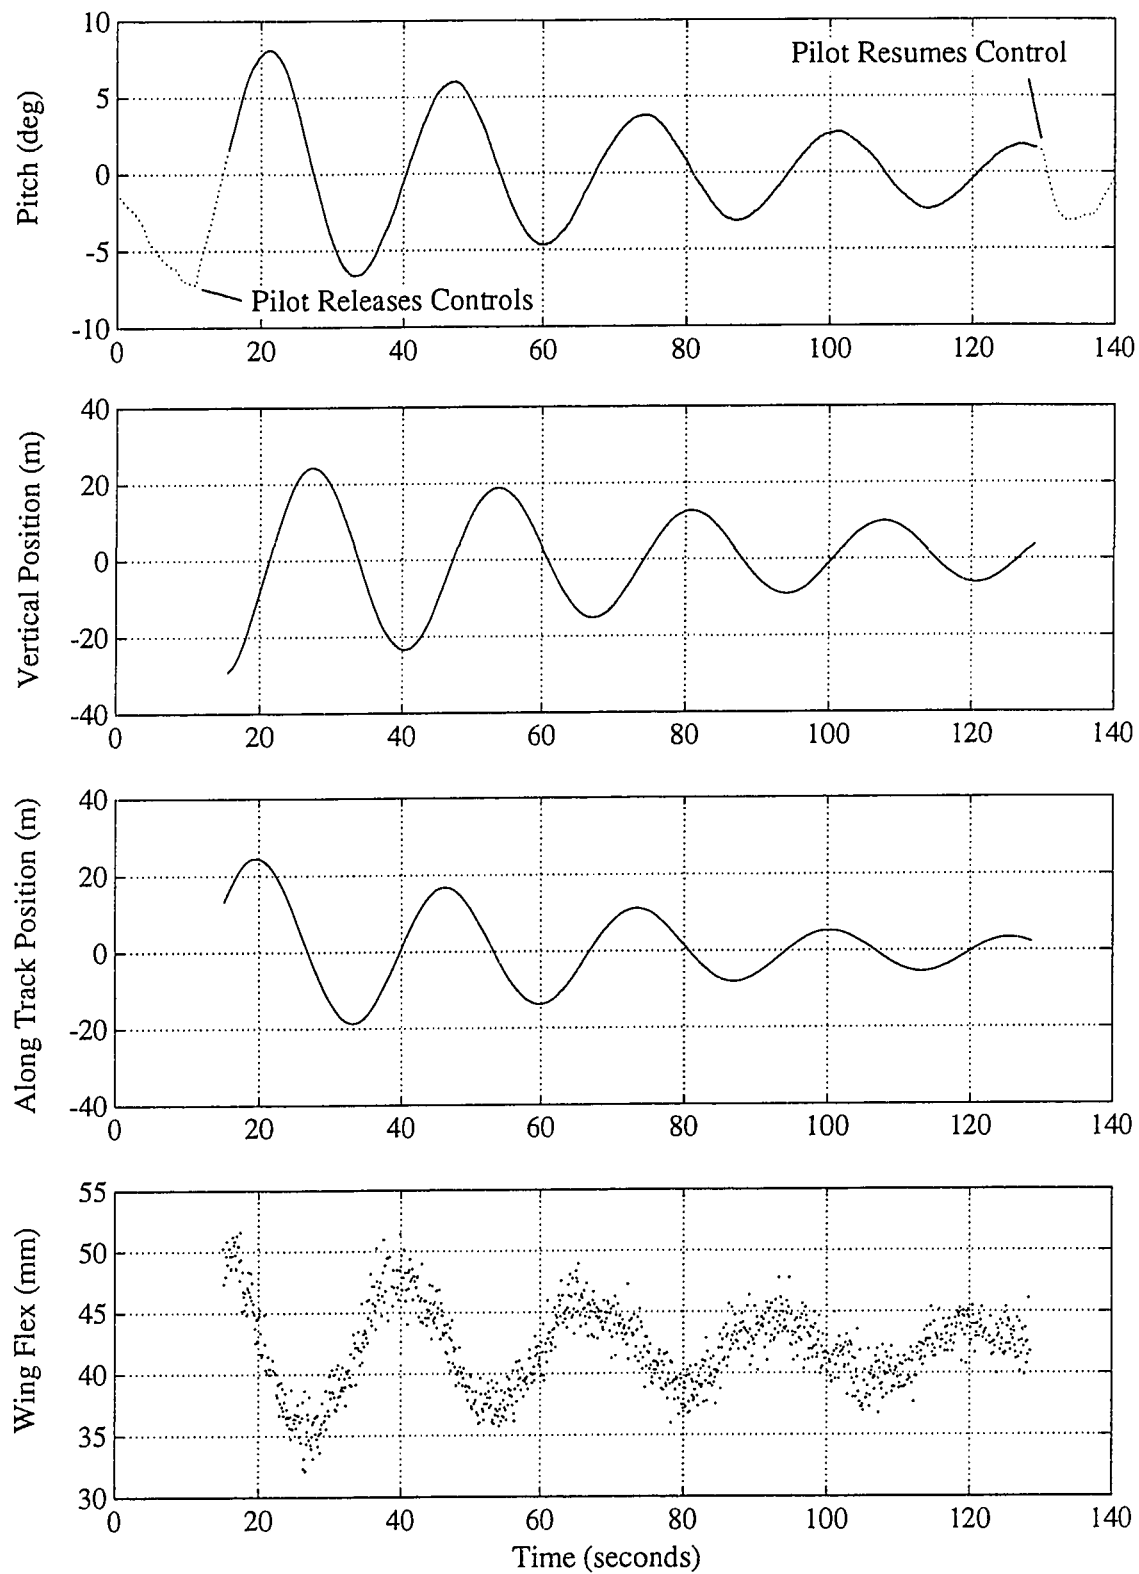


Figure 6.2.1.6: Piper Dakota Phugoid Mode State Estimates

These phugoid plots—along with the other impulse responses—contain a wealth of information regarding the aircraft dynamics. Table 6.2.1.2 provides the natural frequency, ω_n , and damping, ζ , for each mode.

A phasor (eigenvector) diagram of the measured parameters of the phugoid is shown in Figure 6.2.1.8. This




Figure 6.2.1.7: Phugoid Mode Relative Position

diagram shows the relative phases of the altitude, h , the altitude rate, \dot{h} , the along track position and velocity, x and u , respectively, the pitch angle, θ , and the wing flex.

Even from an elementary consideration of conservation of energy, one can obtain a coarse approximation for the aspect ratio of the phugoid spiral.⁴⁹ This argument suggests that the classical undamped spiral should be taller than its width by a factor of $\sqrt{2}$. The experimental value comes out to be quite reasonable at 1.3.

More accurate derivations of the dynamic response are based on the stability derivatives of the aircraft. From the aircraft equations of motion,⁴⁹ the natural frequency and damping of the phugoid mode can be expressed as

Table 6.2.1.2: Longitudinal Modes

Indicated Air Speed: 120 knots

Altitude: 7500 ft

Mode	ω_n (rad/sec)	ζ
short period	6.32	0.41
phugoid	0.236	0.063

$$2\zeta\omega_n = -X_u + \frac{M_u(X_\alpha - g)}{M_\alpha}$$

and

$$\omega_n^2 = -\frac{g}{U_0} \left(Z_u - \frac{M_u}{M_w} Z_w \right)$$

where g is the Earth's gravitational acceleration, U_0 is the nominal unperturbed airspeed, and the other parameters are stability derivatives of forces (X and Z defined in Figure 6.2.1.3) and the pitch moment, M , with respect to small changes in airspeed (u and v in Figure 6.2.1.3) and angle of attack, α .

In many cases it is possible to simplify these expressions. For the Piper Dakota, the second term in the expression for natural frequency is negligible. Furthermore, it is possible to apply the standard subsonic approximations by assuming that change in coefficient of lift is independent of u and that Z_u is given by setting the trimmed lift and weight equal. The resulting expression for natural frequency becomes⁴⁹




Figure 6.2.1.8: Phasor Diagram of Phugoid

$$\omega_n = \sqrt{2} \frac{g}{U_0}$$

This result matches the GPS experimental data to within a few percent. But what is extraordinary is that (under quasi-static conditions) through this relation, *GPS measurements alone are sufficient to make estimates of both airspeed and ground speed without the use of any other sensor* when an aircraft's phugoid mode is excited. Other information about the stability derivatives can be obtained from the relative amplitude and phase of the observable quantities. For example, the phase difference between u and θ yields the ratio of pitch moment stability derivatives:

$$\angle\left(\frac{-\theta}{u}\right) = \tan^{-1} \frac{\omega_n \sqrt{1-\zeta^2}}{\zeta \omega_n + (gM_u/M_\alpha)}$$

It is likely that the pole locations of the short period mode could also be used to extract airspeed. At slower flight regimes, the Piper Dakota short period mode becomes less damped and more prolonged.




Figure 6.2.1.9: Dutch Roll Mode Response

GPS attitude sensing provides a wealth of information on the lateral dynamic modes as well. Figure 6.2.1.9 shows the aircraft response to a rudder kick. The Dutch roll mode is clearly seen as an interchange between roll and heading angles. It also appears that the spiral divergence mode has been excited.

6.2.2 DYNAMIC TESTING WITH INERTIAL MEASUREMENT UNITS

In an effort to formally characterize its accuracy and bandwidth performance, the attitude determination receiver will be flight tested against an inertial measurement unit (IMU)

during December, 1992. Based on completely separate physical principles, this testing will provide an independent means of evaluating overall performance.

A Litton LN-93 strap-down ring laser gyro IMU will be operated in the main cabin as the independent reference. Prior to the test flights, a static calibration of both units will be performed on the ground using a surveyed azimuth marker at NASA Ames. The receiver antenna self-survey will be performed over the surveyed marker. The IMU will gyrocompass to measure the north vector. Its accelerometers will concurrently measure local vertical. This will calibrate the IMU's "gimbal angles" (which actually reside in software).

6.3 Spacecraft Experiments

Initial satellite tests of attitude determination using GPS are planned for May, 1993 on two separate unclassified Air Force Space Test Program spacecraft named APEX and RADCAL. Each of these spacecraft will carry the new attitude determination receiver.

APEX

The first test flight of the new receiver will be on APEX, the Advanced Photovoltaic and Electronics Experiment. The spacecraft, scheduled to be launched on a Pegasus, is the first PegaStar bus to be flown by Orbital Sciences Corporation. It will be injected into a 195 by 1054 n.mi elliptical orbit with an inclination of 70.5 deg. It is sun-pointing, maintaining an inertially fixed attitude.

Figure 6.3.1 shows a drawing of the vehicle. Three GPS antennas are mounted on different faces. The figure only shows one of the antennas. A second is placed exactly opposite the first on the opposing face of the spacecraft. A third is placed looking aft on the anti-sun face. Here the multiple antenna configuration provides not only excellent




Figure 6.3.1: APEX Spacecraft, Deployed Configuration




Figure 6.3.2: RADCAL Satellite Configuration




Figure 6.3.3: Top View of RADCAL Satellite

visibility for position fixes but also the opportunity for performing limited attitude determination. Note that the signal processing architecture associated with antenna multiplexing is able to track separate satellites simultaneously on separate antennas. When geometry permits, given GPS satellites will be visible simultaneously on two antennas. Three axis attitude will be available for a significant fraction of each orbit.

The telemetry stream for APEX will be limited to a fixed output format stored at one minute intervals. When the spacecraft is in contact with the ground station at Onizuka Air Force Base in Sunnyvale, this output rate will increase to 1Hz.

RADCAL

The RADCAL satellite, whose bus is being built by Defense Systems, Incorporated, is scheduled for launch on a Scout. Its primary function will be calibration of Air Force over-the-horizon radar systems. In addition to carrying the GPS receiver, the spacecraft payload will include the radar transponder, a TRANSIT receiver for navigation, and a store-and-forward mailbox.

The spacecraft will be injected into a polar circular orbit of 450 n.mi altitude. It is gravity gradient stabilized, offering the interesting prospect of incorporating a dynamic model in the GPS attitude processing.

As shown in Figures 6.3.2-3, four GPS antennas will be placed about the upper rim of the spacecraft bus. To attenuate the amount of multipath from the gravity gradient boom, each of the four GPS antennas is canted off from the zenith by 18 deg. Dual receivers will be flown for redundancy, each fed by the same four cross-strapped antennas. At Vandenberg Air Force Base prior to launch, the GPS receiver integrated into the spacecraft will be allowed to collect 24 hours worth of static calibration data. These

measurements will be used to estimate the baseline configuration and determine the nominal line biases.

The telemetry scheme for RADCAL is significantly more flexible than that for APEX. A 3MByte buffer will store *all* output coming from the receiver until the buffer is filled. The buffer is scheduled to be downloaded twice per day. Since the receiver will be reconfigurable on orbit, many GPS experiments will be possible, including precise orbit determination using carrier phase. A provision to output carrier beat phase at a 10Hz output rate has been added to the receiver, making possible initial experiments of atmospheric occultation using GPS.⁵⁰

Chapter 7:

Applications of GPS Attitude

The capability to use GPS for attitude determination opens up a new realm of applications and opportunities. The key element that makes it possible is the centimeter-level precision of GPS carrier phase. In the future, it is very likely that the integration of attitude determination into larger systems—including those that use carrier phase for positioning as well—will play a key role in realizing the full potential of GPS.

Although the scope of this work is restricted to aerospace, it is quite likely that the first widespread use of attitude determination will be in marine applications. In all likelihood, GPS will relegate the gyrocompass (used traditionally for heading determination) to the role of a backup system. GPS attitude determination may also prove to be a valuable means of antenna pointing, especially on moving platforms. This section presents a partial list of the future aerospace applications of attitude determination.

7.1 Aircraft Applications

In the near future, GPS will undoubtedly have a profound influence on aviation on a global scale. GPS has the potential to benefit many aspects of aircraft operation, raising the safety and efficiency of air travel to unprecedented levels.

With the addition of GPS-based attitude determination to the complement of sensors available to a pilot, *a single GPS sensor has the potential to perform the functions of more than half of the cockpit instruments currently in use.* A partial list of aircraft applications follows.

Heading and Attitude Sensing

Perhaps the most obvious of the aircraft applications, heading and attitude sensing based on GPS is especially attractive because the readout is immune to drift and magnetic variation. Whether ultimately used as primary or backup, GPS offers prospects of in-flight inertial alignment, cross-check, flight control, or redundant attitude determination.

In commercial passenger aviation, GPS has the potential to augment or replace expensive inertial navigation systems. All but a few airlines currently carry triple redundant INS/IMU packages, many of which cost more than \$300,000. To address potential cost savings, the questions may be asked, can GPS replace inertial navigation, become an independent backup, or become prime navigation with inertial as a backup?

Estimation of Winds Aloft/Wind Shear Detection

When combined with GPS ground velocity and pitot airspeed, good attitude information helps to provide instantaneous vector measurements of winds aloft. Subtracting the local airspeed vector (determined by vehicle attitude and pitot tube measurements and possibly augmented with angle of attack and sideslip measurements) from the inertial velocity (determined by GPS) yields the local wind velocity *vector*. This capability is useful not only for aiding route planning for the host aircraft and others (using ACARS, for example, as the data link) but also as a means for detecting downdrafts, vortices, or wind shear. Combined with the optimal wind shear avoidance strategy proposed by Chu and

Bryson,⁵¹ attitude determination offers another line of defense against this deadly phenomenon.

Precision Landing

The precision landing problem using GPS implicitly requires attitude sensing to apply the lever-arm correction between antennas and landing gear. Positioning of the landing gear—not the GPS antenna—is the pilot's focus.

The centimeter level precision of carrier phase is regarded as critically important to accomplishing Category II/III precision landing using GPS. A proposed means of resolving the cycle ambiguities by overflying a pseudolite⁵² appears promising. Should this approach prove feasible, attitude determination using GPS will play an integral role.

System Identification and Autopilot Synthesis

A new and exciting application of attitude determination is using it to identify the aircraft dynamic model. The idea is that a pilot can supply inputs to the controls that excite a dynamic response. The GPS sensor measures this response quantitatively to reconstruct an accurate model of the aircraft dynamics. The model then serves as the foundation for optimal autopilot synthesis.

Receiver Autonomous Integrity Monitoring (RAIM)

Extending the system identification idea further, it is possible that model estimation can be carried out continuously in flight. In effect, the dynamic model adds a new set of constraints to the position fixes performed by the navigation equipment. The position fix *has to agree* with the attitude solution applied to the known dynamic model or else there is something wrong, either with the GPS system or the aircraft itself. Should any off-nominal conditions arise, the system could warn the pilot immediately.

This new set of dynamic model constraints have the potential to provide an additional line of defense against a variety of GPS positioning failure modes (with the possible exception of slow drifts), leading to an overall improvement in GPS system integrity.

Formation Flying

The precise positioning and attitude provided by GPS carrier phase makes it possible to instrument formation flying. The most likely military application of this could be “blind” mid-air refueling. As another example, instrumenting the *Blue Angels* or the *Thunderbirds* would establish once and for all exactly how close these aircraft fly.

7.2 Spacecraft Applications

The flight receiver developed here is small (1300 cc), light (~1.5 kg) and low power (~3.5W) such that it can be carried on just about any spacecraft. Soon, attitude determination in space using GPS may emerge in a number of roles. A partial list of potential applications is provided here.

Maneuvering Vehicles

Since it provides unrestricted 3-axis attitude and position, the method is especially appealing for maneuverable vehicles, including the Gravity Probe B (Relativity Gyroscope Experiment), Orbital Maneuvering Vehicle, space-based robots, teleoperators, and the Space Shuttle. Once proven in orbit, attitude determination using GPS is likely to serve as a formidable sensor for closed-loop spacecraft attitude control.

The Gravity Probe B experiment—under development at Stanford—will test Einstein’s General Relativity Theory using ultra-precise orbiting gyroscopes. The spacecraft will make extensive use of GPS for initial injection, closed-loop guidance for orbit trim,⁵³

geodesy,⁵⁴ operational orbit determination, time transfer,⁵⁵ and attitude determination.⁵⁶

A single GPS receiver will be the sensor for all of these functions.

Adjunct Attitude Sensing

As an adjunct sensor, GPS can serve as backup attitude or for initial acquisition, bringing a guide star within the field of view of a star tracker.

Small Satellites

For small satellites, GPS offers significant opportunities for cost savings. As a ten state sensor, providing Position, Velocity, Time, and Attitude (PVTa), a *single GPS sensor* can augment or replace the entire traditional complement of Guidance, Navigation, and Control sensors, many of which required completely separate and independent hardware to accomplish orbit and attitude determination.

Space Vehicle Flexibility

Vehicle flexibility may yield an important new class of applications for GPS differential phase measurements. Attitude determination (the rigid body mode) is really the first of multiple modes of flexibility. Stringing GPS antennas along the length of a flexible space structure (such as the Space Station) could provide an excellent means of estimating the vibrational mode shapes and excitation (see Section 2.5). Space tethers are another area of application. GPS could be used as a means for estimating tether dynamics in real-time.

Launch Vehicle Guidance

GPS also offers new capability in booster guidance. Whether GPS is used simply for velocity cut-off or complete closed-loop guidance and control in launch vehicle applications, GPS will undoubtedly play a significant role. The high performance launch vehicle of the future with an electronically-steered “smart skin” phased array antenna is

shown in Figure 5.2.4.2. Attitude control could be based entirely on the all solid-state GPS sensor. Precise carrier phase kinematic positioning could also be applied to launch vehicles in real-time throughout their entire ascent into orbit.

Precision Orbit Determination

The precision of GPS carrier phase is now being exploited for high accuracy orbit determination.⁵⁷ Incorporating attitude sensing into the problem renders the process transparent to vehicle orientation and antenna placement. Each phase measurement from a given antenna can be projected easily to any point on the vehicle, such as the center of mass. A new technique using single-frequency code and carrier phase measurements to estimate the absolute ionospheric delay may be applicable to precision orbit determination.⁵⁸

Rendezvous and Docking

Given the precision of GPS carrier phase, it is possible that the entire rendezvous and docking operation in space could be performed using GPS as the sole sensor. With the addition of a communication link and an additional computer, the GPS sensor may be all that is required to carry out all phases, including final maneuvering in position and attitude for final contact at the docking receptacles.

Kinematic positioning in space enjoys the added benefit of significantly more satellites in view at any given time. Using the geometry of relative motion, the cycle ambiguities associated with carrier phase can be resolved quickly and reliably with the aid of the over-determined system of equations resulting from the abundance of satellites.

Another angle to precise positioning and attitude is the idea of *formation flying*, applicable both to aircraft and spacecraft. Keeping several vehicles aligned and

positioned to tight specifications in flight makes possible new flight vehicle configurations—such as a radio telescope—formed by separate, unconnected elements.

The Distributed Satellite

In the distant future, there will undoubtedly be many space applications of radio navigation which are perhaps so futuristic today that they may not yet even make sense in this era. One such possibility is the idea of the *distributed satellite*.

Could a single “chip” of hardened semiconductor free-floating in space be engineered as a complete satellite? The idea is to render the payload, power, communication, navigation and control functions of a satellite as a monolithic unit. Power provided by photovoltaic cells is just as feasible as compact communications and navigation made possible by state-of-the-art electronics and MMIC technology. It is certainly conceivable that magnetic torque could be used as one means of attitude control and that solar radiation pressure, modulated by electro-optical changes in surface reflectivity, could be invoked as a means of attitude and translation control for a spacecraft with such a large area-to-mass ratio.

Glaser⁵⁹ conceived the idea of a solar power satellite designed to transfer collected energy via a microwave link to terrestrial receiving stations. The concept is based on the construction of a massive structure in geosynchronous orbit, the economics of which are eventually justified by irrefutable trends in global environmental damage and population growth. In 1963, the United States Air Force launched Project West Ford,⁶⁰ an experiment in which 500 million tiny copper wire strands were injected into low Earth orbit. The project was architected as an array of dipole reflectors with which to form a passive communications relay.



Figure 7.2.1: The Distributed Solar Power Satellite

The idea of the *distributed satellite* combines this concept with active components to take the semiconductor spacecraft much further. Suppose that the satellite payload is a radio telescope, space-based radar, or, as the most ambitious undertaking, a solar energy conversion plant. Could a grand ensemble be constructed of single-chip satellites, each of which comprises one cell in a yet larger organism, to perform a far greater function, operating as a coordinated whole? Figure 7.2.1 shows a cloud of billions of free-flying single-chip components, collectively orbiting in an equatorial band like a ring of Saturn. Without structural connections to any of its neighbors, each chip is completely self-supporting.

Because of the minute size of the chip, the energy intercepted from the sun can only be reradiated omnidirectionally via microwaves. The grand implication of the Global Positioning System and its future descendants is that every point in space and time is uniquely addressable to millimeter-level (nearly-picosecond) accuracy. Just as two cells of heart tissue spontaneously synchronize themselves to beat in unison when placed in physical contact, a three-dimensional array of omnidirectional radiators could be designed to steer a beam of coherent energy to any destination on the Earth for distribution into the power grid.

Chapter 8:

Conclusion

Attitude determination using GPS has transitioned from a feasible concept to a proven reality. Its demonstrated performance in terms of accuracy and bandwidth qualifies it to handle an abundance of new applications. For those applications which had been addressed previously with inertial instruments, the attitude readout from GPS is completely immune to drift. But for all its performance capability, the ultimate advantage of using GPS may turn out to be the reductions in size, weight, power, and cost of the sensor hardware.

As a completely self-contained differential positioning system, attitude determination may well become the highest performance regime of GPS in existence. As such, it presents an excellent test bed for exploring the intrinsic performance limits of GPS itself.

The *performance envelope* of GPS-based attitude determination describes what accuracy is available as a function of tracking loop noise bandwidth. The state-of-the-art performance envelope indicates that pointing accuracy better than 0.1 degree is available with a usable bandwidth of several tens of Hz.

An advanced attitude determination receiver based on antenna multiplexing has been developed with a solution output rate of 10 Hz. The multiplexing design allows the receiver to be rendered in hardware at minimum size, weight, power, and cost. This receiver was installed in a single engine Piper Dakota for extensive flight testing.

Rapid, reliable cycle ambiguity resolution is key to the viability of GPS-based attitude determination. Cycle ambiguities are resolved for aircraft applications by taking advantage of the motion of the antenna array during single-axis turns on the ground or banked turns in the air. An efficient approach to estimating the integers is presented whereby relative position vectors derived from raw measurements are stacked into a system of linear equations in 9-space.

Estimation of wing flex to an accuracy of 1.4mm is demonstrated as a method for obtaining instantaneous wing loading and estimating aeroelastic effects. Such accurate positioning with GPS effectively uses the structural deformation of the aircraft under stress as an accelerometer.

The applications of GPS-based attitude determination to aircraft are explored. The aircraft dynamic model is estimated using GPS exclusively. The period of the phugoid mode is used to obtain an estimate of *airspeed* derived exclusively from GPS. The possibility of using GPS to continuously estimate dynamic parameters is explored as a means for improving efficiency and safety.

The method of attitude determination using vector observations (Wahba's Problem) has been generalized to attitude determination using GPS. The resulting implementation benchmark performs almost an order of magnitude faster than the traditional solution based on a non-linear least-squares fit.

The new GPS receiver is being prepared for flight tests in space. It will be flown on two Air Force spacecraft, APEX and RADCAL, currently scheduled for launch in May, 1993. A new technique for resolving integer ambiguity of the GPS carrier for spacecraft attitude determination has been demonstrated. The technique makes use of both *translational* and *rotational* spacecraft motion.

Multipath is the largest error source affecting attitude determination. The measurement accuracy may be improved by through antenna pattern shaping, multipath calibration, increased inter-antenna separation, solution filtering, and signal processing. The feasibility of calibrating out error due to multipath for small spacecraft in low Earth orbit has been demonstrated experimentally.

Suggestions for Future Work

Attitude determination capability using GPS opens up many new opportunities for future research. These opportunities might be divided into two categories: performance improvements and applications. Many of the applications (see Chapter 7), consist of integrating attitude determination into larger systems, such as the case of closed-loop guidance and control, precision landing of aircraft, and rendezvous and docking of spacecraft.

In areas of performance improvement, exploring ways of further reducing the effects of multipath is called for. Rejection of multipath through antenna pattern shaping (see Section 5.2) is likely to be a fruitful avenue of pursuit. Exploring phased arrays of antennas will also be a means of enhancing bandwidth performance (see Section 5.2) and addressing cycle ambiguity resolution (see Section 3.5).

Quantitative testing under dynamics will further establish the performance envelope of GPS. Test flights with an inertial measurement unit are advocated (see Section 6.2). Quantitative tests may be carried out with a receiver antenna mounted on a motorized rotary table both to measure tracking roll-off and to test signal processing algorithms to track signals through rapidly rotating antennas (see Section 2.5).

Future work in aircraft applications will contribute directly to enhancing the overall safety and efficiency of aviation. Research characterizing the dynamic models of aircraft using GPS can be used in the attitude solution itself, synthesis of optimal autopilots, and the development of more advanced receiver autonomous integrity monitoring (RAIM). The potential of GPS to measure vector wind velocity should be pursued to explore all of its possible benefits.

Potential near-term flight experiments include closed-loop attitude control for both spacecraft and aircraft, multipath calibration in space against a star sensor, and launch vehicle ascent trajectory analysis based on GPS carrier phase. Flight experiments could further extend differential phase techniques to aeroelastic and flexible space structures.

Effort directed toward miniaturization of receiver hardware will lead to the ultimate applications of GPS. In principle, the entire receiver can be realized as a monolithic unit.

Epilogue

The Global Positioning System is having a profound effect on aerospace guidance, navigation, and control. For the first time in history, almost every point in space (and time) near Earth is easily addressable in real-time to nearly millimeter-level accuracy. The return from GPS is a wealth of new capability, accompanied by enhanced safety and extraordinary improvements in efficiency in both aviation and space.

Appendix A:

Measurement Time Tags

As shown in Figure A.1, each code phase measurement is averaged over *MCOUNT* milliseconds (typically 500) and time tagged at the center of the averaging interval. For the carrier phase measurements, the instantaneous value of the carrier phase at each antenna is strobed at a receiver millisecond clock edge.

The measurement process is triggered by the receiver clock reaching an even 500 millisecond mark. The final millisecond of averaging begins at the C/A code epoch which immediately follows the millisecond mark. The time tag is then given by the midpoint of the 500 millisecond averaging interval. The time tag for the carrier phase measurement is the millisecond immediately preceding the trigger millisecond.




Figure A.1: Raw Measurement Time Tags

Appendix B:

User Spacecraft Almanac

To aid search and acquisition in low Earth orbit, a simple Keplerian orbit model was used to propagate the spacecraft orbital elements with time. The perturbation due to J_2 was accounted for using the well-known corrections to the Keplerian elements:

$$\frac{d\Omega}{dt} = -\frac{3}{2}J_2n\left(\frac{R_E}{a}\right)^2 \frac{1}{(1-e^2)^2} \cos i$$

$$\frac{d\omega}{dt} = \frac{3}{2}J_2n\left(\frac{R_E}{a}\right)^2 \frac{1}{(1-e^2)^2} \left(2 - \frac{5}{2}\sin^2 i\right)$$

$$n = \frac{dM}{dt} = \sqrt{\frac{\mu}{a^3}} \left[1 + \frac{3}{2}J_2n\left(\frac{R_E}{a}\right)^2 \frac{1}{(1-e^2)^{3/2}} \left(1 - \frac{3}{2}\sin^2 i\right) \right]$$

where Ω , ω , and M are the right ascension of the ascending node, argument of perigee, and mean anomaly, respectively, μ is the Earth's gravitational constant, R_E is the Earth's radius, a is the semi-major axis, n is the mean motion, e is the orbit eccentricity, and i is the orbit inclination.

Note that the entire GPS coordinates are referenced with respect to the Earth-Centered, Earth-Fixed (ECEF) coordinate frame. Once the user spacecraft vector position, \mathbf{r} , and velocity, \mathbf{v}_I , is calculated with respect to the inertial frame (neglecting the effect of J_2 on the osculating orbit elements), the Coriolis rule is applied to the ECEF frame, rotating about the north pole with angular velocity $\boldsymbol{\omega}_E$:

$$\mathbf{v}_{ECEF} = \mathbf{v}_I - \boldsymbol{\omega}_E \times \mathbf{r}$$

To speed the search process even further in light of antenna multiplexing, the acceleration of the user spacecraft relative to the ECEF frame is calculated.

$$\mathbf{a}_{ECEF} = -\frac{\mu}{r^2} \hat{\mathbf{r}} - \boldsymbol{\omega}_E \times (\boldsymbol{\omega}_E \times \mathbf{r}) - 2\boldsymbol{\omega}_E \times \mathbf{v}_{ECEF}$$

A set of relative position, velocity, and acceleration vectors may be constructed relative to the ECEF frame as follows:

$$\begin{aligned}\Delta \mathbf{r} &= \mathbf{r}_{GPS} - \mathbf{r}_{ECEF} \\ \Delta \mathbf{v} &= \mathbf{v}_{GPS} - \mathbf{v}_{ECEF} \\ \Delta \mathbf{a} &= \mathbf{a}_{GPS} - \mathbf{a}_{ECEF}\end{aligned}$$

The center of the search frequency window, δf , is positioned based on this predicted relative velocity and the receiver clock bias rate, \dot{b} , known approximately from previous position fixes.

$$\delta f = -\frac{f_{L1}}{c} \Delta \mathbf{v} \cdot \Delta \mathbf{r} + \dot{b}$$

The relative range acceleration seen by the GPS receiver is similar except for an additional centripetal term:

$$\delta f = -\frac{f_{L1}}{c} \left[\Delta \mathbf{a} \cdot \Delta \mathbf{r} + \left(\frac{\Delta \mathbf{v}^2 - (\Delta \mathbf{v} \cdot \Delta \mathbf{r})^2}{|\Delta \mathbf{r}|} \right) \right]$$

With a worst case search through 4 individual antennas, the relative acceleration exceeding 1g can cause the Doppler shift to drift outside the search window before the search is completed. Therefore, as discussed at the end of Section 2.4, a bias rate is applied to the search frequency based on the calculated relative range acceleration. This feed-forward arrangement subtracts off the bulk of the relative motion so that the search can proceed at maximum speed.

Appendix C:

Kinematic Carrier Phase Positioning

This appendix outlines the processing steps that were taken to determine the relative positioning for the phugoid mode shown in Figure 6.1.3.6 . The ground station set up at Stanford recorded raw carrier phase measurements, $\varphi_i^{ref}(t_{ref})$, for the i th satellite with respect to receiver clock time, t_{ref} , at a 1Hz rate. The aircraft receiver strobed its own measured instantaneous carrier beat phase, $\varphi_i^{air}(t_{air})$, with respect to its own clock time, t_{air} , at a 10Hz rate. Since the clock bias of each receiver with respect to GPS time was known from position fixes, each measurement time tag could also be fixed with respect to GPS time, t_{GPS} . To eliminate the effect of changing geometry on the relative positioning, the reference station measurements were interpolated (projected) to the aircraft measurement times now given with respect to GPS time, t_{GPS} . Given the observed stability of the TCXO in the reference receiver, the reference data was expected to be good to better than a wavelength over the one second interval.

The carrier phase single difference observable used for positioning was then constructed to eliminate the satellite transmitter clock errors and selective availability as follows:

$$\Delta\varphi_i(t_{GPS}) = \Delta\varphi_i^{air}(t_{GPS}) - \Delta\varphi_i^{ref}(t_{GPS})$$

In this particular example, the relative range between the reference station at Stanford and the aircraft en route from Palo Alto to Auburn airport was on the order of 100km. No effort was made to resolve the exact cycle ambiguities, nor was this even necessary. Instead, a second-order polynomial curve fit was subtracted out of each raw carrier phase single difference. This removed the orbital acceleration of the GPS satellites, geometrical effects due to Earth rotation, and the linear translation of the aircraft during the measurement interval.

To solve for the relative aircraft position, δx (3×1), and relative receiver clock bias, δt , the single difference observable from each of the six GPS satellites being tracked were combined together in a least-squares fit for each measurement epoch:

$$\begin{bmatrix} -\hat{s}_1^T & 1 \\ -\hat{s}_2^T & 1 \\ \vdots & \vdots \\ -\hat{s}_6^T & 1 \end{bmatrix} \begin{bmatrix} \delta x \\ \delta t \end{bmatrix} = \begin{bmatrix} \Delta\phi_1 \\ \Delta\phi_2 \\ \vdots \\ \Delta\phi_6 \end{bmatrix}$$

where \hat{s}_i is the unit vector to the i th GPS satellite in the local horizontal reference frame of the aircraft. Since this matrix equation happens to be overdetermined by two measurements, there was an opportunity for a sanity check on the carrier phase residuals. The RMS residual error was less than an L1 wavelength.

The final step was to perform a rotation on the relative position solution about the local vertical to resolve the east and north components of relative position into cross track and along track relative position.

References

¹V.W. Spinney, "Applications of the Global Positioning System as an Attitude Reference for Near Earth Uses," Institute of Navigation, Warminster, PA, April, 1976.

²R.L. Greenspan, A.Y. Ng, J.M. Przyjemski, and J.D. Veales, "Accuracy of Relative Positioning by Interferometry with Reconstructed Carrier GPS: Experimental Results", *Proceedings of the Third International Geodetic Symposium on Satellite Positioning*, Las Cruces, NM, March, February, 1982.

³A.K. Brown, W.M. Bowles, T.P. Thorvaldsen, "Interferometric Attitude Determination Using the Global Positioning System: A New Gyrotheodolite", Third Geodetic Symposium on Satellite Doppler Positioning, Las Cruces, NM, 1982, pp. 1289-1302.

⁴K.M. Joseph and P.S. Deem, "Precision Orientation: A New GPS Application", International Telemetry Conference, San Diego, CA, October, 1983.

⁵L.R. Kruczynski, P.C. Li, A.G. Evans, and B.R. Hermann, "Using GPS to Determine Vehicle Attitude: U.S.S. Yorktown Test Results", *Proc. Int. Tech. Mtg.*, ION, Colorado Springs, CO, Sept., 1989.

⁶G.H. Purcell, Jr., J.M. Srinivasan, L.E. Young, S.J. DiNardo, E.L. Hushbeck, Jr., T.K. Meehan, T.N. Munson, and T.P. Yunck, "Measurement of Aircraft Position, Velocity and Attitude using Rogue GPS Receivers", Fifth International Geodetic Symposium on Satellite Positioning, Las Cruces, NM, March, 1989.

⁷F. van Graas, and M. Braasch, "GPS Interferometric Attitude and Heading Determination: Initial Flight Test Results", *Navigation*, Vol. 38, Fall, 1991, pp. 297-316.

⁸K. Ferguson, J. Kosmalska, M. Kuhl, J.M. Eichner, K. Kepski, and R. Abtahi, "Three-Dimensional Attitude Determination with the Ashtech 3DF 24-Channel GPS Measurement System", IEEE Position Location and Navigation Symposium, March, 1991.

⁹J. Rath and P. Ward, "Attitude Estimation using GPS", Nat. Tech. Mtg., ION, San Mateo, CA, January, 1989.

¹⁰Clark E. Cohen and Bradford W. Parkinson, "Expanding the Performance Envelope of GPS-Based Attitude Determination", ION GPS-91, Albuquerque, NM, Sept., 1991.

¹¹ S.E. Moore, "PIN Diode Switch for Multiplexing TANS", Trimble Engineering Drawing, January, 1992.

¹²M. Martin-Neira, R. Lucas, and M.A. Martínez, "Attitude Determination with GPS: Experimental Results", IEEE AES Magazine, September, 1990, pp.24-29.

¹³M.E. Cannon, J.B. Schleppe, and J.P. McLellan, "Real-Time Heading Determination Using an Integrated GPS Dead-Reckoning System", ION GPS-92, Albuquerque, NM, Sept. 1992.

¹⁴R.A. Brown, A.G. Evans, "GPS Pointing System Performance", ION GPS-90, Colorado Springs, CO, September, 1990.

¹⁵Michael Dentinger and Clark E. Cohen, "Differential Phase Measurements through Antenna Multiplexing", Patent Pending.

¹⁶C.R. Trimble, "Global Positioning System Coarse Acquisition Code Receiver", United States Patent 4,754,465, 1988.

¹⁷J.J. Spilker, *Digital Communications by Satellite*, Prentice-Hall, Englewood Cliffs, New Jersey, 1977.

¹⁸J.J. Spilker, Jr., "GPS Signal Structure and Performance Characteristics", Global Positioning System papers published in *Navigation*, Vol I, 1980, pp. 29-54.

¹⁹P.E. Braisted and R.F. Eschenbach, "Global Positioning Sytem Coarse Acquisition Code Receiver; Differential Phase Measurement System", United States Patent 4,847,862, 1989.

²⁰A.E. Bryson and Y.C. Ho, *Applied Optimal Control*, Hemisphere, New York, Revised Printing, 1975.

²¹G.F. Franklin, J.D. Powell, M.L. Workman, *Digital Control of Dynamic Systems*, 2nd ed., Addison-Wesley, 1990.

²²Clark E. Cohen, "Doppler Determination of Spacecraft Attitude", United States Patent 5,070,338, 1991.

²³M. Martin-Neira and R. Lucas, "GPS Attitude Determination of Spin Stabilized Satellites", ION GPS-92, Albuquerque, NM, Sept. 1992.

²⁴R.A. Brown, "Instantaneous GPS Attitude Determination", IEEE Position Location and Navigation Symposium, Monterey, CA, March, 1992, pp. 113-120.

²⁵R.R. Hatch, "Ambiguity Resolution in the Fast Lane", *Proceedings of the Second International Technical Meeting of the Institute of Navigation*, Colorado Springs, CO, 1989.

²⁶R. Brown and P. Ward, "A GPS Receiver with Built-in Precision Pointing Capability", IEEE Position Location and Navigation Symposium, Las Vegas, NV, March, 1990, pp. 83-93.

²⁷H. Landau and J.M.F. Ordóñez, "A New Algorithm for Attitude Determination with GPS", Institute of Astronomical and Physical Geodesy (IAPG), University FAF Munich, Germany, *personal correspondence*, 1992.

²⁸G.H. Golub and C.F. van Loan, *Matrix Computations*, 2nd ed., Johns Hopkins Univ. Press, Baltimore, 1989.

²⁹Clark E. Cohen and Bradford W. Parkinson, "Integer Ambiguity Resolution of the GPS Carrier for Spacecraft Attitude Determination", *Advances in the Astronautical Sciences*, 14th Annual AAS Guidance and Control Conference, Keystone, CO, February, 1992.

³⁰R.N. Bracewell, *The Fourier Transform and Its Applications*, 2nd ed., McGraw-Hill, New York, 1986.

³¹G.J. Bierman, *Factorization Methods for Discrete Sequential Estimation*, Academic Press, Inc., New York, 1977.

³²G.H. Golub, Class Notes, CS237A Advanced Numerical Analysis, Stanford University, Fall, 1988.

³³G. Wahba, "A Least-Squares Estimate of Spacecraft Attitude", *SIAM Review*, Vol. 7, No. 3, July, 1965, p. 409.

³⁴F.L. Markley, "Attitude Determination Using Vector Observations: A Fast Optimal Matrix Algorithm", 1992 Goddard Space Flight Center Flight Mechanics/Estimation Theory Symposium, Greenbelt, MD, May, 1992.

³⁵M.D. Shuster, "Three-Axis Attitude Determination from Vector Observations", *Journal of Guidance and Control*, Vol. 4, No. 1, 1981, pp. 70-77.

³⁶F.L. Markley, "Attitude Determination using Vector Observations and the Singular Value Decomposition", *Journal of the Astronautical Sciences*, Vol. 36, No. 3, 1988, pp. 245-258.

³⁷J.R. Wertz (ed.), *Spacecraft Attitude Determination and Control*, Reidel, Boston, 1978.

³⁸C.E. Cohen and B.W. Parkinson, "Mitigating Multipath in GPS-Based Attitude Determination", *Advances in the Astronautical Sciences*, AAS Guidance and Control Conference, Keystone, CO, Vol. 74, Univelt, San Diego, 1991.

³⁹Y. Georgiadou and A. Kleusberg, "On Carrier Signal Multipath Effects in Relative GPS Positioning", *Manuscripta Geodaetica*, Vol. 13, No. 1, 1988, pp.1-8.

⁴⁰I. Nesbø, "Applications for GPS Determined Attitude for Navigation", Institute of Navigation International Technical Meeting, Colorado Springs, Colorado, September, 1988.

⁴¹J.D. Kraus, *Antennas*, 2nd. ed., McGraw-Hill, New York, 1988.

⁴²D.M. Pozar and B. Kaufman, "Design Considerations for Low Sidelobe Microstrip Arrays", *IEEE Transactions on Antennas and Propagation*, Vol. 38, No. 8, August, 1990, pp. 1176-1185.

⁴³M. Braasch and F. van Graas, "Guidance Accuracy Considerations for Real-Time Interferometric Attitude Determination", ION GPS-91, Albuquerque, NM, Sept. 1991.

⁴⁴L.L. Hagermann, "Effects of Multipath on Coherent and Non-Coherent PRN Ranging Receiver", Aerospace Corp. Report #TOR-0073 (3020-03)-3, May, 1973.

⁴⁵A.J. Van Dierendonck, P. Fenton, and T. Ford, "Theory and Performance of Narrow Correlator Spacing in a GPS Receiver", ION National Technical Meeting, San Diego, CA, January, 1992.

⁴⁶ Clark E. Cohen, H. Stewart Cobb, and Bradford W. Parkinson, "Two Studies of High Performance Attitude Determination Using GPS: Generalizing Wahba's Problem for High Output Rates and Evaluation of Static Accuracy Using a Theodolite", ION GPS-92, Albuquerque, NM, Sept. 1992.

⁴⁷Clark E. Cohen and Bradford W. Parkinson, "Aircraft Applications of GPS-Based Attitude Determination: Test Flights on a Piper Dakota", ION GPS-92, Albuquerque, NM, Sept. 1992.

⁴⁸"GPS Attitude Sensing Tested ", *Aviation Week and Space Technology*, November 30, 1992, p. 52.

⁴⁹D. McRuer, I. Ashkenas, D. Graham, *Aircraft Dynamics and Automatic Control*, Princeton Univ. Press, Princeton, 1973.

⁵⁰I. Linscott, "Earth Atmospheric Profiles Using GPS Occultation", ION GPS-92, Albuquerque, NM, Sept. 1992.

⁵¹ Peter Y. Chu, *Control of Aircraft Under Severe Wind Shear Conditions*, Ph.D. Thesis, Department of Aeronautics and Astronautics, Stanford University, December, 1986.

⁵²Clark E. Cohen, J. David Powell, Boris Pervan, David Lawrence, Bradford W. Parkinson, "Real-Time Cycle Ambiguity Resolution using a Pseudolite for Kinematic Positioning with GPS", to be presented at the Second International Symposium on Differential Satellite Navigation Systems, DSNS '93, 30 March - 2 April 1993, Amsterdam, The Netherlands.

⁵³P. Axelrad and B.W. Parkinson, "Closed Loop Navigation and Guidance for Gravity Probe B Orbit Insertion", *Journal of the Institute of Navigation*, Vol. 36, No. 1, Spring, 1989, pp. 45-61.

⁵⁴J.V. Breakwell, B.W. Parkinson, C.W.F. Everitt, M.D. Tapley, "Impact of the Gravity Probe B Mission on Satellite Navigation and Geodesy", *Int. Assoc. of Geodesy Symposia 103: Gravity, Gradiometry, and Gravimetry*, ed. R. Rummel, R.G. Hipken, Springer-Verlag, New York, 1990.

⁵⁵Todd Walter, Ph.D. Thesis, Department of Applied Physics, Stanford University, to be published.

⁵⁶Clark E. Cohen and Bradford W. Parkinson, "Attitude Determination Using GPS on Gravity Probe B", *Journal of the Astronautical Sciences*, Gravity Probe B special edition, to be published.

⁵⁷T.P. Yunck S.C. Wu, S.M. Lichten, W.I. Bertiger, U.J. Lindqwister, and G. Blewitt, "Toward Centimeter Orbit Determination and Millimeter Geodesy with GPS" Fifth International Geodetic Symposium on Satellite Positioning, Las Cruces, NM, March, 1989.

⁵⁸Clark E. Cohen, Boris Pervan, and Bradford W. Parkinson, "Estimation of Absolute Ionospheric Delay Exclusively through Single-Frequency GPS Measurements" ION GPS-92, Albuquerque, NM, Sept., 1992.

⁵⁹Peter E. Glaser, *Solar Power via Satellite*, Ellis Harwood, London, 1991.

⁶⁰B. Yenne, *Encyclopedia of US Spacecraft*, Exeter, New York, 1985.