



2806377548

**GPS Based Position and Attitude Determination
for Airborne Remote Sensing**

Kevin Francis Sheridan

Thesis submitted for the degree of
Doctor of Philosophy
of the University of London

Department of Geomatic Engineering
University College London

October 2000

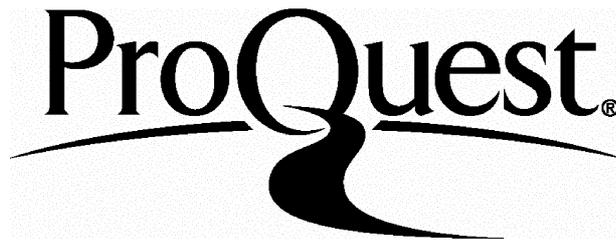
ProQuest Number: U643744

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

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



ProQuest U643744

Published by ProQuest LLC(2016). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code.
Microform Edition © ProQuest LLC.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Abstract

The design and testing of a system to determine the position and attitude of an airborne sensor from GPS observations are described. Using GPS data from a fixed array of four antennas, each connected to an independent receiver, all six exterior orientation parameters of an airborne scanner can be determined with an accuracy that meets the georeferencing requirements for a range of applications.

Airborne data collected in two flight trials have been processed using software developed during this research project, and a number of previous studies, to determine a best estimate of the aircraft trajectory. The position and attitude determination algorithms both employ single epoch ambiguity resolution. Attitude is determined directly from GPS measurements, without explicitly estimating relative antenna positions. This increases the redundancy in the solution and hence improves the reliability and robustness of the system. Attitude values derived from GPS data have been integrated with attitude data from a gyro to improve the efficiency of the ambiguity resolution process. A Kalman filter is used to obtain an optimal estimate of the trajectory by integrating all available data and generating orientation values for each scan line.

Initially the subjects of georeferencing and GPS based attitude determination are introduced. Trends within georeferencing are discussed, particularly the development of systems that directly measure the position and attitude of airborne sensors using on-board instruments. This provides the context for the specific experimental system developed in the course of this work. The direct georeferencing system is described in terms of the instrumentation and processing methodology employed. Results show that the system is capable of positioning accuracies of 5 to 10mm in plan, and 10 to 30mm in height. It determines attitude with an accuracy of 7 arc minutes in pitch, 20 arc minutes in roll, and 3 arc minutes in heading. At a nominal flying height of 600m this leads to an 'on-the-ground' positioning accuracy of below 4m. By comparing this performance to application and sensor accuracy requirements, the suitability of the georeferencing system can be evaluated.

Contents

Title Page	1
Abstract	2
Contents	3
List of Figures	9
List of Tables	13
List of Acronyms	16
Acknowledgements	19
Chapter 1 Introduction	20
1.1 Overview	20
1.2 Research Objectives and Scope	23
1.2.1 Specific Objectives	23
1.2.2 General Objectives	23
1.2.3 Scope of Research	24
1.3 Research Methodology	25
1.4 Thesis Outline	26
Chapter 2 Georeferencing	28
2.1 Introduction	28
2.2 Indirect Georeferencing	29
2.3 Direct georeferencing	33
2.3.1 Overview: Description Benefits and Evolution	33
2.3.2 Orientation Requirements for Airborne Applications	36
2.3.3 System Design	40
2.3.3.1 Overview	40
2.3.3.2 Equipment	40
2.3.3.3 Processing	50
2.3.3.4 System Performance	58
2.4 Concluding Remarks	63

Chapter 3	GPS Attitude Determination	64
3.1	Introduction	64
3.2	Principles of GPS Attitude Determination	65
	3.2.1 Co-ordinate Systems	65
	3.2.2 Rotations	70
3.3	Dedicated GPS Attitude Systems	72
3.4	Attitude from Independent GPS Receivers	75
	3.4.1 Equipment	75
	3.4.2 Processing Methods	77
	3.4.2.1 Baseline-by-Baseline Attitude Determination	77
	3.4.2.2 Direct Attitude Determination	79
3.5	GPS Observation Equations for Attitude	80
3.6	Error Sources and Mitigation	84
3.7	Concluding Remarks	87
Chapter 4	System Design and Testing	88
4.1	Introduction	88
4.2	NERC Airborne Remote Sensing Facility	89
	4.2.1 Overview	89
	4.2.2 Sensor Characteristics	93
	4.2.2.1 RC-10 Aerial Camera	93
	4.2.2.2 ATM Whiskbroom Line Scanner	93
	4.2.2.3 CASI Pushbroom Line Scanner	94
	4.2.3 Ground Spatial Resolution	97
	4.2.3.1 ATM Whiskbroom Line Scanner	97
	4.2.3.2 CASI Pushbroom Line Scanner	99
	4.2.3.3 Spatial Resolution Summary	101
	4.2.4 Direct Georeferencing Equipment	101

4.3	Flight Trials	105
4.3.1	Test Site	105
4.3.2	Equipment Specifications	107
4.3.3	Calibration	108
4.3.4	Operational Difficulties	110
4.4	Concluding Remarks	111

**Chapter 5 Position and Attitude Determination
from GPS Flight Data 112**

5.1	Introduction	112
5.2	Processing Overview	113
5.3	Position Determination	116
5.3.1	Overview	116
5.3.2	GASP	117
5.3.2.1	Methodology	117
5.3.2.2	Previous GASP Performance	120
5.3.3	SKI	121
5.4	Static GPS Attitude Determination	123
5.4.1	GRAPE Methodology	124
5.4.1.1	Determination of Antenna Body Frame Co-ordinates	128
5.4.1.2	Determination of Approximate Attitude	130
5.4.1.3	Generating and Reading NXF Data	132
5.4.1.4	Determination of a Search Space	132
5.4.1.5	Computation of the Modified Double Difference Phase Values	134
5.4.1.6	Computation of the Ambiguity Function Values	135
5.4.1.7	Least Squares Computation and Statistical Testing	136
5.4.2	GRAPE Performance Assessment: Static Testing	138
5.5	Kinematic GPS Attitude Determination: Extensions to GRAPE	143
5.5.1	Overview	143
5.5.2	Attitude Approximation for Search Volume	144
5.5.3	Search Increment and Size	146

6.4.3	Ambiguity Resolution	202
6.4.3.1	Ambiguity Validation	202
6.4.3.2	Single Baseline vs. Direct Attitude Ambiguity Resolution	203
6.4.4	Accuracy Assessment	208
6.4.4.1	Static Testing	209
6.4.4.2	Attitude Accuracy based on Positioning Accuracy	211
6.4.4.3	Attitude Accuracy Summary	214
6.5	Recording Rates	215
6.5.1	Position Rate	215
6.5.2	Attitude Rate	217
6.5.3	Summary of Recording Rates	220
6.6	Concluding Remarks	221
Chapter 7 Integration and Filtering		223
7.1	Introduction	223
7.2	GPS and Gyro Integration for Ambiguity Determination	224
7.3	Filtering	227
7.3.1	Introduction	227
7.3.2	Attitude Filter Design	229
7.3.2.1	Design Overview	229
7.3.2.2	Measurement Model	232
7.3.2.3	Dynamic Model	236
7.3.2.4	Filter Equations	239
7.3.2.5	Implementation of the Filter	242
7.3.3	Filter Results	244
7.4	Concluding Remarks	248

Chapter 8	Conclusions and Suggestions for Further Work	249
8.1	Conclusions	249
8.2	Suggestions for Further Work	254
8.2.1	Calibration	254
8.2.2	Algorithms	255
8.2.3	General Suggestions	257
	References and Bibliography	259
Appendix A	Partial Derivatives for Modified Double Differencing	273
Appendix B	GPS Data Description	278
Appendix C	GPS Phase Residuals	287
Appendix D	PRAYER Software Description	291

List of Figures

Figure 2.1	Airborne scanning	31
Figure 2.2	GPS Stand-alone horizontal accuracy	47
Figure 2.3	Feedforward configuration for an aided inertial navigation system. Linearised Kalman filter.	51
Figure 2.4	Feedback configuration for an aided inertial navigation system. Extended Kalman filter.	53
Figure 2.5	Block diagram of a decentralised Kalman filter configuration for INS/GPS integration	55
Figure 2.6	Block diagram of the AIMS tightly integrated GPS/INS system	56
Figure 3.1	Body Frame System (BFS)	66
Figure 3.2	Attitude defined by BFS with respect to LLS	67
Figure 3.3	ECEF and LLS	69
Figure 3.4	GPS phase difference measurement (interferometry)	73
Figure 4.1	NERC survey aircraft	90
Figure 4.2	NERC survey aircraft: Cross-sectional schematic	90
Figure 4.3	Winglet and GPS antenna	92
Figure 4.4	ATM with AHRS attached	92
Figure 4.5	Imaging concept of the CASI CCD pushbroom spectrograph	95
Figure 4.6	CASI spatial mode	95
Figure 4.7	CASI spectral mode	96
Figure 4.8	ATM pixel size and swath width ($\pm 45^\circ$ FOV) as a function of altitude	98
Figure 4.9	CASI pixel size and swath width (42° FOV) as a function of altitude	100
Figure 4.10	Antenna arrangement on NERC ARSF aircraft	102
Figure 4.11	Flight plan	106

Figure 5.1	Processing overview	114
Figure 5.2	GASP processing	119
Figure 5.3	GRAPE processing	125
Figure 5.4	Ambiguity Resolution in GRAPE	126
Figure 5.5	Least Squares Routine in GRAPE	127
Figure 5.6	Pitch determination from a single appropriate baseline	131
Figure 5.7	GRAPE search space	132
Figure 5.8	GRAPE search size and trial attitudes	133
Figure 5.9	GRAPE search increment	148
Figure 5.10	GRAPE search increment (selected values)	149
Figure 5.11	Satellite elevation dependent weighting function	153
Figure 6.1	GASP derived height of antenna 4 during flight line L013	173
Figure 6.2	Ambiguity validation	175
Figure 6.3	Standard deviations from SKI processing. L013, antenna 2	184
Figure 6.4	Standard deviations from GASP processing. L013, antenna 2	184
Figure 6.5	GASP and SKI agreement for antenna 2 position on L001	185
Figure 6.6	GASP and SKI agreement for antenna 1 position on L008	186
Figure 6.7	A1-A2 vector from SKI positions, L006	187
Figure 6.8	A2-A4 vector from SKI positions, L006	188
Figure 6.9	A1-A4 vector from SKI positions, L006	188
Figure 6.10	A1-A2 vector from GASP positions, L001	188
Figure 6.11	GRAPE and AHRS pitch, L008	193
Figure 6.12	GRAPE and AHRS roll, L008	193
Figure 6.13	GRAPE and AHRS heading, L008	193
Figure 6.14	GRAPE and AHRS heading, L008 – 15 second period	194
Figure 6.15	Heading derived from GASP antenna positions and GRAPE, L001	196
Figure 6.16	Pitch derived from GASP antenna positions and GRAPE, L001	197
Figure 6.17	Aircraft height. 9:45-9:50, 16/10/98	198

Figure 6.18	Aircraft plan position. 9:45-9:50, 16/10/98	198
Figure 6.19	Aircraft pitch from GRAPE. 9:45-9:50, 16/10/98	199
Figure 6.20	Aircraft roll from GRAPE. 9:45-9:50, 16/10/98	199
Figure 6.21	Aircraft heading from GRAPE. 9:45-9:50, 16/10/98	199
Figure 6.22	Deviation in GRAPE pitch solution from three minute static period	210
Figure 6.23	Deviation in GRAPE roll solution from three minute static period	210
Figure 6.24	Deviation in GRAPE heading solution from three minute static period	210
Figure 6.25	Antenna 1 to antenna 2 baseline length using GASP, L013	212
Figure 6.26	Antenna 3 to antenna 4 baseline length using GASP, L013	212
Figure 6.27	Antenna 2 height during L001, using a 10Hz and 1Hz solution	216
Figure 6.28	Antenna 2 height during a 20 second period of L001, using a 10Hz and 1Hz solution	216
Figure 6.29	GRAPE and AHRS roll. L013	218
Figure 6.30	GRAPE roll. 2Hz and 10Hz solutions	219
Figure 6.31	GRAPE heading. 2Hz and 10Hz solutions	220
Figure 7.1	Filtered solution from ADU2 and AHRS roll data. 20405	245
Figure 7.2	Filtered solution from ADU2 and AHRS pitch data. 20405	245
Figure 7.3	Filtered solution from ADU2 and AHRS heading data. 20405	245
Figure 7.4	Filtered solutions from ADU2 and AHRS pitch data. 20402	247
Figure B.1	Number of satellites from which observations were recorded at antenna 1 during the observation period on day 288	280
Figure B.2	Number of satellites from which observations were recorded at antenna 2 during the observation period on day 288	280
Figure B.3	Number of satellites from which observations were recorded at antenna 2 during the observation period on day 288	280

Figure B.4	Satellite observations recorded by receiver 1. Day 289	281
Figure B.5	Satellite observations recorded by receiver 2. Day 289	281
Figure B.6	Satellite observations recorded by receiver 3. Day 289	281
Figure B.7	Satellite observations recorded by receiver 4. Day 289	281
Figure B.8	Number of satellites tracked at the reference station on day 288	282
Figure B.9	Number of satellites tracked at the reference station on day 289	282
Figure B.10	Sky plot day 288	283
Figure B.11	Sky plot day 289	284
Figure B.12	Satellite elevations for the observation period on day 288	285
Figure B.13	Satellite elevations for the observation period on day 289	285
Figure B.14	PDOP for the observation period on day 288 (15-10-98)	286
Figure B.15	PDOP for the observation period on day 289 (16-10-98)	286

List of Tables

Table 2.1	Accuracy requirements for airborne applications	37
Table 2.2	Georeferencing accuracy requirements for airborne survey systems	38
Table 2.3	INS position, velocity and attitude accuracies	42
Table 2.4	Attitude accuracy from GPS based systems	46
Table 2.5	GPS differential positioning accuracy	48
Table 2.6	Possible sensor configurations	59
Table 2.7	Survey of direct georeferencing performance	60
Table 4.1	ADU2 accuracy specification	102
Table 4.2	Litef LCR-92H AHRS technical specifications	103
Table 4.3	Scanner specifications during test flights	107
Table 5.1	BFS co-ordinates from NERC ARSF antenna survey	128
Table 5.2	BFS co-ordinates from static GPS antenna survey	129
Table 5.3	GRAPE accuracy from static testing	139
Table 5.4	L1 and L2 double differenced phase residuals from GRAPE	140
Table 5.5	Outlier rejection comparison	142
Table 6.1	Processed flight lines	165
Table 6.2	Selected GASP ambiguity success rate with standard parameters	178
Table 6.3	Final GASP ambiguity resolution performance	180
Table 6.4	Ambiguity resolution success using a single baseline solution	205
Table 6.5	Number of possible solutions from the GASP F-test	207
Table 6.6	GRAPE attitude accuracy from static data	211
Table 6.7	Summary statistics from all baseline estimations during L013	212
Table 6.8	Attitude errors predicted from positioning errors	214
Table 6.9	Attitude accuracy from all available sources	214

Table 7.1	Search volume and processing time	225
Table B.1	GPS data logging summary. Day 288 (15-10-98)	278
Table B.2	GPS data logging summary. Day 288 (16-10-98)	279
Table B.3	Satellite elevations & azimuths at start and end of day 288 observation period	283
Table B.4	Satellite elevations & azimuths at start and end of day 289 observation period	284
Table C.1	L1 phase residuals between the reference station and antenna 2 during L003	288
Table C.2	L2 phase residuals between the reference station and antenna 2 during L003	288
Table C.3	L1 phase residuals between the reference station and antenna 4 during L003	288
Table C.4	L2 phase residuals between the reference station and antenna 4 during L003	288
Table C.5	L1 phase residuals between the reference station and antenna 2 during L006	289
Table C.6	L2 phase residuals between the reference station and antenna 2 during L006	289
Table C.7	L1 phase residuals between the reference station and antenna 4 during L006	289
Table C.8	L2 phase residuals between the reference station and antenna 4 during L006	289
Table C.9	L1 phase residuals between the reference station and antenna 2 during L008	290
Table C.10	L2 phase residuals between the reference station and antenna 2 during L008	290

Table C.11	L1 phase residuals between the reference station and antenna 4 during L008	290
Table C.12	L2 phase residuals between the reference station and antenna 4 during L008	290

List of Acronyms

ADU	Attitude Determination Unit
AFM	Ambiguity Function Method
AFV	Ambiguity Function Value
AHRS	Attitude and Heading Reference System
AROF	Ambiguity Resolution On-the-Fly
AS	Anti Spoofing
ASCII	American Standard Code for Information Interchange
ASCOT	Aerial Survey Control Tool (LH Systems)
ATM	Airborne Thematic Mapper
BFS	Body Frame System
C/A	Coarse/Acquisition GPS signal
Cs	Caesium
CASI	Compact Airborne Spectrographic Imager
cm	centimetre
DoD	Department of Defence
DOP	Dilution of Precision
DoT	Department of Transport
ECEF	Earth-Centred Earth-Fixed
FARA	Fast Ambiguity Resolution Approach
FOC	Full Operational Capability
FOG	Fibre-Optic Gyroscope
FOV	Field of View
GASP	GPS Ambiguity Searching Program (single epoch ambiguity resolution program for positioning)
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
GRAPE	GPS Routine for Attitude Parameter Estimation (single epoch ambiguity resolution program for attitude determination)
HDOP	Horizontal Dilution of Precision
Hz	Hertz

IFOV	Instantaneous Field of View
IMU	Inertial Measurement Unit
INS	Inertial Navigation System
IOC	Initial Operating Capability
ISPRS	International Society for Photogrammetry and Remote Sensing
km	Kilometre
L1	GPS L-band signal 1 (1575.42 MHz)
L2	GPS L-band signal 2 (1227.6 MHz)
LAMBDA	Least Squares Ambiguity Decorrelation Adjustment
LLS	Local Level System
LSAST	Least Squares Ambiguity Searching Technique
m	Metre
MHz	Mega-Hertz
mm	Millimetre
NERC	Natural Environment Research Council
NERC ARSF	Natural Environment Research Council's Airborne Remote Sensing Facility
NXF	Newcastle eXchange Format
OCS	Operational Control Segment
OEM	Originally Engineered by Manufacturer
OEEPE	European Organisation for Experimental Photogrammetric Research
OTF	On-The-fly
P-code	Precision code
PDOP	Position Dilution of Precision
PRAYER	Pitch Roll and Yaw Estimation Routine
PRN	Pseudo-Random Number
PVAF	Position Velocity Acceleration Filter
Rb	Rubidium
RINEX	Receiver INdependent EXchange (format)
RINtoNXF	RINEX to NXF (conversion program)
RMS (rms)	Root Mean Square (the square root of the average of the squared errors)

RTK	Real-time kinematic
SA	Selective Availability
SAR	Synthetic Aperture Radar
SEP	Spherical Error Probable (a sphere's radius centred as the true position containing 50% of the points in a three dimensional scatter)
SKI	Static Kinematic Software (from Leica Geosystems)
SV	Space Vehicle
UERE	User Equivalent Range Error
VLBI	Very Long Baseline Interferometry
WGS84	World Geodetic System 1984

Symbols

cm	centimetre
mm	Millimetre
m	Metre
nm	Nanometre
km	Kilometre
µm	Micro-metre
σ	Standard deviation
°	Degrees of arc
'	Minutes of arc
"	Seconds of arc

Acknowledgements

I would like to express my gratitude to my supervisor, Professor Paul Cross, for his support and guidance during my studies, and for proof reading this thesis. I would also like to thank my former and present colleagues at the Department of Geomatics at the University of Newcastle-upon-Tyne, and the Department of Geomatic Engineering, University College London for their help and advice.

I am also grateful for the practical assistance I received from various sources which enabled the test flights, which were an essential part of this project, to go ahead. Thanks in particular to Andrew Wilson and John Cook of the Natural Environment Research Council's Airborne Remote Sensing Facility for their help and advice. The GPS equipment used during these flights was kindly supplied by Colin Beatty (CBI Ltd), Martin Robertson (Newcastle University), and Rick Blighton and colleagues at Ashtech Europe Ltd., their assistance is gratefully acknowledged. Thanks also to Adrian Darby and Peter Smith who kindly allowed us to use Peterborough Business Airfield as our survey site, and to Joel Barnes who helped complete the ground survey and to install and operate the GPS equipment on the aircraft.

This research was made possible by financial support from the Natural Environment Research Council.

On a personal note I would like to thank my family and friends for their continued support, and especially Jo, for her love, understanding and proof-reading.

Chapter 1 Introduction

1.1 Overview

In almost all airborne remote sensing applications there is a requirement to relate features from the imagery to a mapping co-ordinate system. The process of establishing the transformation between image co-ordinates in an arbitrary, sensor dependent reference frame and a more transferable mapping reference frame is termed georeferencing. Conventional methods of georeferencing use ground control points distributed within the survey area to reconstruct the geometry of the imaging device. Establishing suitable ground control is generally expensive, time-consuming, and in some cases, is not practically possible. In remote sensing applications using airborne scanners, it is not practical to provide orientation parameters for individual lines using ground control, as the survey area is imaged as a series of thousands, or even tens of thousands, of scan lines.

By providing on-board navigation instruments it is possible to determine all six exterior orientation parameters (three components of position, and three components of attitude) of airborne sensors directly, allowing mapping co-ordinates to be determined for the imagery with minimal, or even no, ground control. These navigation instruments have included steroscopes, laser profilers, stellar cameras, inertial navigation systems (INS), and Global Positioning System (GPS) receivers. Each of these systems has the potential to resolve all or a subset of the six parameters required for individual exposures. Only recently have two of these systems, namely INS and GPS, been fully integrated such that all six orientation parameters are determinable with sufficient accuracy at each epoch (Schwarz *et al.*, 1993). The potential reductions in time and cost, and the increased operational flexibility that 'direct georeferencing' offers, have made it an increasingly active research area over the past few years.

The majority of systems that have been developed to measure both the position and attitude of airborne platforms have integrated some form of INS with GPS. Cramer and Haala (1999), Toth and Grejner-Brzezinska (1998), and Skaloud and Schwarz (1998) discuss the design and performance of a number of such systems. In general, a high quality INS comprising accelerometers and gyroscopes is used to derive all six exterior orientation parameters of the imaging device. GPS position and velocity data allow the time dependent errors of the inertial positioning to be constrained. The resultant system combines accurately derived attitude at a high measurement rate, typically 50-100 Hz, from the INS, with the long-term positional stability of GPS. The high data rate characteristic is particularly important when each individual scan line in an image requires its own orientation parameters.

An alternative approach is to determine attitude from GPS carrier phase signals measured by an array of antennas. A number of commercially available systems have been developed (e.g. the Ashtech ADU2 and Trimble TANS Vector) that consist of either three or four GPS antennas rigidly fixed together and linked to a special receiver that 'measures' the phase differences between the carrier signals received by different antennas. Essentially, kinematic phase GPS processing takes place between these antennas, and the co-ordinate differences obtained are converted to attitude components – usually pitch, roll and heading.

Attitude determination can also be achieved in a similar way, by attaching each antenna to an independent GPS receiver. GPS phase data are collected, vectors between antennas are computed, and the orientation parameters are estimated, using least squares if more than the minimum of three antennas have been deployed. In principle, any of the currently available methods can be used for the critical part of the processing – the estimation of the baseline vectors. This novel approach, using independent receivers has been used in the current research, but the relatively simple processing strategy described has been substantially modified.

One weakness of kinematic GPS baseline estimation, is the lack of redundancy in the solution, which limits the reliability (the ability to detect outlying data) and robustness (the ability to correctly resolve integer ambiguities) of the system. To overcome this, methods have been developed that estimate attitude components directly from GPS phase observations, without explicitly determining relative antenna positions. A direct attitude determination method, based on modifying the standard double differenced phase observation equation, has been adopted in this research. This approach was proposed by Cross *et al.* (1997) and has subsequently been implemented in software known as the GPS Routine for Attitude Parameter Estimation, or GRAPE (Mahmud, 1999). Results from processing GPS data collected in a static test, have demonstrated the potential benefits of this approach in terms of robustness and reliability. In the current research project, this direct attitude determination technique has been extended to allow kinematic GPS data to be processed, in order to estimate the attitude of a moving platform, in this case an aircraft.

This research has concentrated on developing the direct GPS attitude determination method as a practical solution to orientation problems. It has been used as part of a specific airborne direct georeferencing system to assess its performance under operational conditions and to assess the wider potential of the method. This direct georeferencing system not only demonstrates the potential of this GPS attitude determination technique, it also solves a more specific practical problem, i.e. to provide position and attitude information for imagery recorded by a survey aircraft.

A system has been developed to meet the exterior orientation requirements of the Natural Environment Research Council's Airborne Remote Sensing Facility (NERC ARSF). It consists of four independent dual frequency GPS receivers, each attached to an antenna on the main body of a survey aircraft, and a low-cost attitude and heading reference system (AHRS) based on three fibre optic gyros, which is rigidly fixed to a scanner inside the aircraft. Initially a solution for all six exterior parameters based only on GPS data was developed, but the benefits of combining

GPS data with higher frequency gyro data for attitude determination were also considered. This solution should provide reliable position and attitude data at a frequency and accuracy adequate for the two scanners used on this aircraft. The imaging equipment is typical of that used in a range of airborne remote sensing applications, therefore the use of the direct georeferencing system as a more generic tool can also be assessed.

1.2 Research Objectives and Scope

As the previous section has indicated, the objectives of this research can be broadly divided into two categories; specific and general.

1.2.1 Specific Objectives

- To develop an operational GPS based position and attitude determination system to provide sufficiently accurate exterior orientation values for the scanner imagery recorded by the NERC ARSF survey aircraft.
- To assess whether or not this solution can be improved by the addition of attitude data from the AHRS gyro-based device.
- To design a Kalman filter which uses all the available information to produce position and attitude values for each scan line.

1.2.2 General Objectives

- To develop the existing direct GPS attitude method implemented in the GRAPE software, into an operational tool suitable for determining the attitude of a moving platform.

- To investigate the performance of GPS attitude determination in terms of robustness and accuracy.
- To compare the performance of the direct attitude method, with its increased redundancy, to that of a conventional approach based on the determination of relative antenna positions.
- To investigate ways of improving the performance of the GPS attitude determination method by refining the existing algorithms, or by introducing supplementary data sources.

1.2.3 Scope of Research

It was necessary to place limits on the scope of this research to make these objectives realistic with the time and resources available. Some specific points to note are:

- In general, ‘direct georeferencing’ describes the whole process of relating imagery to a mapping reference frame. The overall accuracy of the final georeferenced data will depend on the accuracy of the measured image data, the accuracy with which position and attitude information is determined and then related to the scanner, the optical properties of the imaging device, and the system calibration. As no scanner imagery has been available in the course of this research, the whole direct georeferencing process cannot be assessed, instead it is only the accuracy with which the aircraft’s position and attitude is determined that will be considered. The term ‘direct georeferencing’ is not always used in this thesis in its more widely accepted sense, as these additional elements are not considered.

- Although multipath is known to be a significant error source in this application its treatment is somewhat limited. An assessment is made of the system's performance using GPS data that are contaminated to some degree by multipath, but specific attempts are not made to model multipath errors either stochastically or through the functional model.
- No attempt is made to explain general concepts of GPS. It is assumed that the reader will already be familiar with the general subject and with widely used and documented techniques such as differencing strategies and ambiguity resolution. Some relevant information sources are listed in the 'references and bibliography' section of this thesis.
- The background information given on photogrammetry and remote sensing is relatively limited. It is not the intention here to explain the fundamentals of these two broad disciplines, instead only aspects directly relevant to the current research are described.

1.3 Research Methodology

The methodology adopted in this research may be summarised as follows:

- 1) Modify the existing direct GPS attitude determination method to process kinematic data.
- 2) Collect airborne GPS data in an operational environment.
- 3) Use these data to determine the aircraft position using standard techniques implemented in existing software.

- 4) Use these data to assess the modified direct GPS attitude determination method and to compare its performance with a single baseline system.
- 5) Assess the benefits of combining attitude data from AHRS with the GPS solution.
- 6) Assess the overall performance of the direct georeferencing system in this particular application, and its potential for future use.

1.4 Thesis Outline

This thesis consists of eight chapters with additional information provided in four appendices.

Chapter 2 introduces the concept of georeferencing, showing why it is required, and describing how it can be achieved. Having identified the limitations of conventional methods it then describes the use of on-board navigation instruments for direct georeferencing. Typical orientation requirements for a range of airborne remote sensing applications are presented, and then a number of systems designed to meet these requirements are discussed in terms of the equipment used and the processing methods adopted.

Chapter 3 introduces the principles of attitude determination using GPS. A number of approaches to GPS attitude determination are then discussed, differentiating between dedicated instruments comprising a single receiver connected to multiple antennas, and systems using independent receivers each connected to a single antenna. The main discussion concentrates on independent receiver systems and identifies a trend towards direct methods in which attitude parameters are determined directly from GPS phase measurements. The direct attitude method, based on modified double differenced observations, that has been used in this research is described.

Having discussed the general subject area, chapter 4 introduces the specific direct georeferencing system used in this research project. The system is described in terms of the equipment used, and the exterior orientation requirements it has been designed to meet. The test flights in which data were gathered to test this system are described, and calibration methods are discussed.

Chapter 5 describes how the GPS data recorded during the test flights have been processed to determine position and attitude solutions. Position determination is discussed briefly as standard methods and software have been used with few significant modifications. The GRAPE software, which implements the direct GPS attitude determination method, is described in more detail. The existing software and its performance with static data is discussed before changes made in the course of this research are introduced.

Chapter 6 presents the position and attitude results that have been determined using the equipment and processing methods described. The robustness of the system is described, and the accuracy of the position and attitude solutions is assessed. The robustness of the direct attitude system is compared to that of a single baseline system. The rate at which position and attitude information must be measured is also discussed.

Chapter 7 describes how attitude data from the AHRS can be combined with GPS data to improve the efficiency and robustness of the system. Separate Kalman filters have been developed to process the position and attitude data. The attitude filter is described in some detail and results from its implementation are presented.

Finally, chapter 8 summarises the conclusions from this research and presents a number of suggestions for further work.

Chapter 2 Georeferencing

2.1 Introduction

This chapter introduces the concept of georeferencing, showing why it is required and describing how it can be achieved. Georeferencing is the process of relating data collected in the form of images to a mapping co-ordinate system. Traditionally this is achieved by connecting images to ground features by visual or digital correlation. This process requires the ground co-ordinates for a number of identifiable features in the image area to be known. This can be achieved by carrying out a dedicated ground survey or by using previously surveyed information from a map or digital database. Using these ground control points, the geometry of the sensor at an instant is reconstructed, i.e. the sensor position and attitude are determined. Once the sensor orientation is established it is then possible to find co-ordinates in the ground system for any feature in the image. This process of using ground control to establish the position and attitude of an airborne sensor can be termed ‘indirect georeferencing’.

By providing on-board navigation instruments it is possible to measure the position and orientation of airborne sensors directly, allowing ground co-ordinates to be determined for imagery with minimal, or even no, ground control. The potential reductions in time and cost and the increased operational flexibility that ‘direct georeferencing’ offers, have made it an increasingly active research area over the past few years.

Section 2.2 briefly describes how indirect georeferencing is applied in both photogrammetric and remote sensing applications. Although this research concentrates on a remote sensing application much of the following discussion is equally relevant in both disciplines and many of the geometric principles and methods applied in remote sensing have their roots in photogrammetry. As well as emphasising some of the common aspects, a number of important differences in the georeferencing requirements of remote sensing and photogrammetry are highlighted. By identifying the limitations of conventional georeferencing

techniques it is then possible to identify applications for which the alternative approach of direct georeferencing is a better, or perhaps the only, solution.

Section 2.3 discusses direct georeferencing. The ability of this technique to overcome the problems identified in section 2.2 is emphasised and its suitability for a range of applications is discussed. Typical orientation requirements for a number of common airborne remote-sensing applications, and associated imaging devices, are given. A range of positioning and attitude systems is discussed, both in terms of the navigation instruments used and the processing methodology adopted. Design features that are specific to certain applications or a particular imaging device are differentiated from the more generic aspects of direct georeferencing systems. This allows the general methodology of direct georeferencing to be described, but also gives an indication of the variety of solutions used. In section 2.4 some concluding remarks are made.

2.2 Indirect Georeferencing

The relating of image co-ordinates in an arbitrary sensor-based system to a more transferable ground co-ordinate system is essential in a wide range of geomatic applications. In conventional airborne photogrammetry, ground control points that can be identified on a single photograph can be used to determine the exterior orientation of the camera at the instant of exposure using image resection techniques. Aerial surveys are designed so that overlaps are created between photographs providing stereo-coverage. Once the exterior orientation parameters (3 position components and 3 attitude components) of two exposure stations forming a stereo pair are determined, image intersection can be used to compute the ground co-ordinates of all other features of interest recorded on the photographs.

In practice it would be unusual to reconstruct the geometry of each exposure station using ground control from the coverage area of a single photograph. Instead the survey area is composed of a series of aerial photographs forming a block. Using ground control points distributed throughout the survey area and

image features common to two or more photographs, (tie points) a combined adjustment is performed to provide all the exposure station orientations simultaneously. The end result is a mosaic of photographs whose relative orientations (with respect to each other) and absolute orientations (with respect to the earth) are known. Constructing a photo-mosaic is a labour-intensive and time-consuming process. In the final step, stereo-pairs are processed to measure the earth-referenced co-ordinates of height contours and terrain objects. Typical accuracies are in the range of 1 part in 5,000 to 1 part in 10,000, depending on map classification (Lithopoulos *et al.*, 1999). This is the principle of aerial triangulation, further details of these methods, including ground control requirements, are given in Wolf (1983).

Digital photogrammetric workstations allow more automated image matching techniques to be employed. These methods make the measurement and identification of ground control points and tie points more efficient, but the triangulation procedure remains a significant part of a photogrammetric survey. An OEEPE (European Organisation for Experimental Photogrammetric Research) and ISPRS (International Society for Photogrammetry and Remote Sensing) test (Heipke, 1999) investigated the performance of automatic aerial triangulation. It showed that a large number of tie points are necessary to achieve reliability and that in many cases manual intervention is still required.

The adjustment process relies on the strong geometrical properties provided by a block of photography. If the principal purpose of a survey is to map a linear feature such as a road, railway, waterway or stretch of coastline, additional ground control may be required to compensate for the weaker geometry.

A further weakness of indirect georeferencing using ground control is the inability to properly separate and estimate errors from different sources. The only data available during the adjustment are the co-ordinate differences between surveyed ground control points and sensor derived co-ordinates of the control points. These differences are the result of aircraft motion, sensor offsets, and misorientation. Usually an approximate interpolation method is used to model the effects of

aircraft motion on the derived co-ordinates from these differences. From a modelling point of view, a major drawback of the method is that it tries to solve a three-dimensional problem from what is essentially two-dimensional information (Schwarz *et al.* 1993).

Whether the triangulation process is carried out using analogue, analytical or digital techniques with strips, blocks or single stereo models, some form of ground control is used to transform image co-ordinates to a ground or mapping reference system.

A similar requirement to relate imagery to a ground system exists in many remote-sensing applications, both from satellites and airborne platforms. To do this rigorously the geometry of the sensing device should be reconstructed just as a camera's geometry must be reconstructed in photogrammetric applications. In the case of a line scanning device, as used in this research project, an image of the survey area is built up as a series of individual scan lines which are then linked to form a larger scene (Figure 2.1).

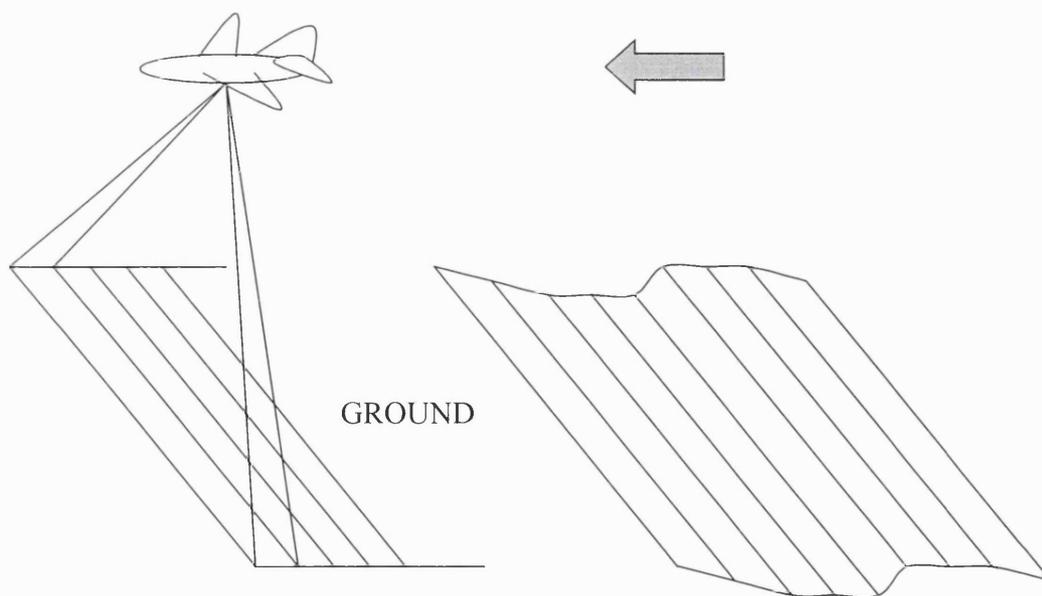


Figure 2.1: Airborne scanning

A typical airborne scanner may record 20 to 100 scan lines per second each of which may be as narrow as 1m (depending on the imaging device and the flying height and speed) and a single survey may comprise tens or even hundreds of thousands of individual scan lines. Clearly it would not be practical to provide ground control for each individual scan line, nor is it practical to identify tie points between each adjacent line.

To relate the imagery to ground co-ordinates in this situation some form of image warping takes place once a scene is formed. A number of identifiable features within the scene that have *known* ground co-ordinates are measured in the image reference frame so that a sample of points have both ground and image co-ordinates. Holding these control points fixed, the image is transformed into the ground co-ordinate system and warped to provide a best fit over the entire survey area. A number of refinements to this process have been developed to minimise the distortion, but nevertheless the ground co-ordinates of image features between the few discrete control points are a relatively crude approximation. Such an approach does not compensate for the effects of the high frequency motion of the scanning device. In reality each individual scan line has its own values for the exterior orientation elements which should be estimated in a rigorous approach. As an airborne survey will typically be carried out by a relatively small aircraft the scanner may be subject to significant high frequency attitude changes even if attempts are made to mitigate these effects with motion compensators and damping devices. The indirect georeferencing method described cannot adequately model the actual conditions experienced by the scanner as each part of the image is captured.

2.3 Direct Georeferencing

2.3.1 Overview: Description, Benefits and Evolution

In the previous section some disadvantages of indirect georeferencing in photogrammetric and remote sensing applications were described. In the case of aerial triangulation these relate primarily to the need to establish a network of ground control points, and the complexity of the adjustment process. The methods used are also optimised for large blocks of photography with good geometrical properties. The accuracy and distribution of ground control points is also an important factor when georeferencing remotely sensed imagery. Additionally, in the case of scanning devices, methods based on ground control cannot adequately account for the high frequency orientation changes experienced by the sensing device.

By measuring the exterior orientation of airborne sensors directly, using instruments on the survey aircraft, ground co-ordinates can be determined for imagery with minimal, or even no, ground control. Direct georeferencing of this kind can reduce the time and cost of a survey and increases operational flexibility. It makes it possible to work in areas where ground control cannot be established due to the terrain, or where access is restricted. This technique is also well suited to surveys of linear features or of limited extent where working with large blocks of imagery would be impractical and inefficient. Systems can also be developed that measure the sensor position and attitude at a similar rate to the acquisition frequency of the sensor, allowing the orientation of individual scan lines to be determined. The clear advantages that direct georeferencing can provide have made it an active research area in the fields of photogrammetry and remote sensing, and many systems exploiting this technology are now used routinely in a commercial environment.

In the 1980's as GPS positioning became available, GPS-supported block triangulation became possible. The application of GPS for aerial triangulation means that the camera station co-ordinates are determined by differential GPS

positioning, and these data are then used in a combined block adjustment (Ackermann and Schade, 1993). Using such a method, the ground control requirement can be significantly reduced or even eliminated. The potential of using GPS observations as constraints for the camera perspective centres in a bundle adjustment has been proven repeatedly, and GPS-supported aerial triangulation is now an accepted procedure (Skaloud *et al.*, 1996). A vast number of papers have been published on this subject outlining the method, the potential accuracies, operational considerations and the range of possible applications. For a few examples see Ackermann (1992), Colomina (1993), Cramer and Haala (1999), Curry (1993), Gruen *et al.* (1993), Lithopoulos *et al.* (1999), Lucas (1987), Mader and Lucas (1989), and Merchant (1993).

A number of commercial systems make use of GPS in block triangulation procedures, and as a tool for more efficient project planning and execution. LH Systems for example produce a series of products that use GPS technology to aid aerial surveys from the planning stages through to georeferencing the images. ASCOT (Aerial Survey Control Tool, LH Systems, 2000a) is a flight management system that can be used as a planning tool to design efficient surveys, and can then be used during the survey for aircraft guidance and pinpoint photography. Raw GPS data and mid-exposure time tags can be read into a GPS post-processing package (Flykin, LH Systems, 2000b) to compute the aircraft trajectory and interpolate the co-ordinates for each exposure station. These co-ordinates are then exported to an aerial triangulation package (ORIMA, LH Systems, 2000c) where they are used in a combined block adjustment (CBA) of GPS and image data.

Similar products are available from other manufacturers, for example the Airborne POS/DGTM from the Applanix Corporation (Applanix, 2000). There are also a number of general purpose bundle adjustment packages that handle geodetic and photogrammetric observations in combination, for example the City University Bundle Adjustment (CUBA), developed at the Engineering Surveying Research Centre at City University, London.

These systems were originally developed to position camera perspective centres at the point of exposure using GPS. In addition to estimating the three position components of exterior orientation it would clearly be advantageous if the attitude elements could also be determined by direct means to provide the complete set of exterior orientation parameters. In the case of a combined bundle adjustment these additional observations significantly increase redundancy and hence improve the reliability of the solution. This additional information is not vital though to achieve a solution, the geometric strength of interlocking bundles can be used to eliminate the attitude requirement (Schwarz *et al.*, 1994). Attitude information becomes increasingly important however when the geometric properties are weaker, for example in the case of corridor surveys of linear features. The requirements will also vary depending on the type of sensor used.

For an airborne line scanner, the ability to measure both the position and attitude is vital, as it allows the complete exterior orientation elements of single scan lines to be determined independently. The direct georeferencing system developed for the current research project has been designed specifically to meet the orientation needs of scanners, hence it measures both position and attitude.

In a combined position and attitude system the sensor attitude can either be determined using some form of Inertial Navigation System (INS) or from an array of GPS antennas on the aircraft¹. With the development of smaller, lighter, less expensive and more rugged inertial devices, and the availability of affordable purpose-built GPS attitude determination units, it has been possible to develop navigation systems that meet the georeferencing needs of a range of sensors and applications, in the fields of both photogrammetry and remote sensing.

In section 2.3.2 the exterior orientation requirements for some typical airborne applications are presented. Section 2.3.3 discusses how direct georeferencing systems can be designed, in terms of equipment and processing, to meet these criteria. The range of possible equipment configurations and processing methods

¹ GPS and INS are by far the most widely used instruments in the field of direct georeferencing but other devices which solve some, or all, of the exterior orientation parameters are available.

is described and some typical systems are discussed in more detail. The performance levels of these systems are described and throughout the discussion some common error sources and their mitigation are noted.

2.3.2 Orientation Requirements for Airborne Applications

The performance requirements of a direct georeferencing system, in terms of position and angular accuracy and recording rate, will vary according to the imaging device and the application. In general, conventional aerial surveys require the exterior orientation of a relatively small number of exposure stations to be accurately determined. Line scanners for example may require a less accurate solution but at a far higher frequency, i.e. many thousand scan lines rather than a few dozen exposure stations. In this section the georeferencing requirements for a range of commonly used sensing devices are presented. The factors that dictate these requirements are briefly discussed and the suitability of sensors for particular survey applications is considered.

To determine the georeferencing accuracy required for a given situation, the final accuracy with which imaged features need to be positioned in the mapping system must be known. These on-the-ground accuracy requirements are then used to determine the accuracy within which the exterior orientation parameters of the airborne sensor must be estimated for the design of the survey. Sensor attitude errors for example will lead to larger on-the-ground positioning errors as the flying height increases. The on-the-ground accuracy requirements will be an important element of the initial survey specifications. These survey specifications will also detail the type of information that needs to be recorded. In the case of remote sensing this may be defined in terms of spectral resolution and range. In practice there will be additional factors to consider including cost, time, equipment availability and the appropriateness of the data for the available processing methods. These factors will define how the survey will be conducted and what sensing equipment will be used. The sensor must be able to resolve ground features both spatially and spectrally at a level consistent with the survey specification.

Table 2.1 displays the sensor orientation accuracies required for different application areas, expressed as root mean square errors (RMS). It indicates that except for precise engineering applications, which require orientation at the level of half an arc minute or better, the bulk of applications require an accuracy of 1-3 arc minutes or considerably less (Skaloud, 1995).

Application Area	Attitude (RMS)	Position (RMS)
Engineering, Cadastral Mapping	15" to 30"	0.05 to 0.1m
Remote Sensing (Detailed)	1' to 3'	0.2 to 1.0m
Resource Mapping	10' to 20'	2 to 5m

Table 2.1: Accuracy for airborne applications (Skaloud, 1995).

In the case of a large or medium scale topographic survey it is likely that the only imaging medium with the required resolution to allow on-the-ground positioning accuracies at the required decimetre level will be conventional frame photography. The image resolution is a function of the photo scale, which can be set by the focal distance of the camera and the flying height. The required exterior orientation accuracy can then be determined accordingly. In applications undertaken in the environmental/earth-sciences, such as land-use surveys or resource mapping, the spatial accuracy required may be orders of magnitude lower, perhaps tens of metres. The ground pixel size must be large enough to provide sufficient radiation for classifying the land cover based on its spectral response. In such cases a balance must be struck between spatial and spectral resolution criteria. Clearly the choice of sensors and the survey design will be based on a range of factors, some that are specific to a particular application and some that are more generic. There is an increasing trend towards fully digital systems for both data capture and processing. The use of line scanners and CCD arrays instead of conventional film photography allows the image acquisition phase to fit more easily into a fully digital system. Adopting a fully digital approach may lead to a degradation in the achievable accuracy but there are many applications for which such an approach is preferable.

Table 2.2 indicates the position and attitude accuracy levels needed for georeferencing a range of sensors.

Type of Sensor	Position (m)	Orientation
aerial camera, m < 1:2000	0.05 to 0.1	15'' to 30''
aerial camera, m > 1:5000	0.75 to 1.0	50'' to 60''
CCD camera or scanner (correlated with pixel size)	0.25 to 1.0m	1' to 3'
interferometric SAR	1 to 2m	10'' to 40''
airborne gravimetry resolution 0.5 to 300km	acceleration 1 mGal	10'' to 20''

Table 2.2: Accuracy requirements for airborne survey (Skaloud and Schwarz, 1998)

The position accuracies range from sub decimetre in the case of an aerial camera to 1 or 2m in the case of interferometric SAR. This obviously has implications for the choice of positioning instruments and how they are utilised. If GPS is the primary source of position information, then a differential carrier-phase solution would be required to achieve the higher accuracy levels. This would almost certainly mean that a minimum of two dual-frequency receivers, one on the aircraft and one operating as a local reference station, would be required². The sensor attitude may need to be determined to ten arc seconds or a few arc minutes, this will dictate whether an inertial system must be used or whether a lower accuracy GPS based system will be adequate. It will also dictate the grade of INS required.

In the case of CCD cameras, the georeferencing accuracy must be correlated with the pixel size. The pixel size will be dictated by the sensing device and survey design (flying height and speed) which in turn is a function of the required spectral resolution, as discussed earlier. Obviously it is not worth assembling an expensive and complex navigation system to provide position and attitude data capable of producing on-the-ground positioning accuracies far better than the spatial resolution of the imaged data.

² There are methods that would allow relatively high accuracy positions with no local base stations (e.g. the virtual reference station concept or a multiple reference solution) but in practice it is still usual to establish a reference station in or close to the survey area.

For all direct georeferencing systems it is necessary to relate position and attitude measurements from the navigation instruments to the sensing device. In the case of GPS positioning this means determining the offset between the antenna phase centre and the principal point of the sensor, this can be achieved by conventional survey techniques³. Time differences between the instant that an image is captured and the measurement epochs of the navigation instruments must also be accounted for. If positions and attitude values are being recorded at a lower frequency than the image acquisition rate, some interpolation is necessary to provide orientation values for the sensor. The success with which an interpolation procedure can model the actual motions experienced by the sensor between measurement epochs significantly affects the achievable orientation accuracies. These points, and a number of other factors that affect how well the accuracy of the position and attitude devices can be transferred to the sensor, are discussed later in this thesis. When designing a system to meet the georeferencing requirements of a sensor it is important to consider the transfer of accuracy as well as the independent accuracy of a navigation instrument.

The overall accuracy of an acquisition system is affected by the accuracy of the measured image data, the INS/GPS position and attitude, the system calibration, the optical properties of the sensor, and the effect of image geometry (Schwarz *et al.*, 1994). A thorough description of all stages of the processing chain is beyond the scope of the chapter. The accuracy of INS/GPS position and attitude determination is the primary focus here.

³ It is not possible in practice to measure from an antenna phase centre to the principal point of a sensor as these two locations are not readily accessible, instead convenient survey marks on the antenna and sensor must be used and corrections applied for the offsets.

2.3.3 System Design

2.3.3.1 Overview

Having discussed the factors that dictate georeferencing accuracy requirements and established some general figures for different applications and sensors, the next step is to consider how georeferencing systems can be designed to meet these criteria. System design is considered here in terms of the equipment (section 2.3.3.2) and the processing methods (section 2.3.3.3) used. The wide range of possible systems is discussed, and the advantages and limitations of particular approaches are indicated. One clear trend in this field is the tailoring of equipment and processing methods for particular applications. There is increasing interest in the development of simple, low-cost systems to meet the needs of the many users whose accuracy requirements are far less stringent than those in the more traditional aerial survey field. Direct georeferencing has also helped to advance the move away from conventional frame photography to fully digital systems for data capture and processing.

2.3.3.2 Equipment

A number of sensors including steroscopes, laser profilers, stellar cameras, INS and GPS have been used for direct georeferencing. Each of these systems has the potential to resolve all, or a subset of, the six exterior orientation parameters. The majority of systems that have been developed to measure both the position and attitude of airborne platforms have integrated some form of INS with GPS.

In general, a high quality INS comprising accelerometers and gyroscopes is used to derive all six exterior orientation parameters of the imaging device. GPS position and velocity data allow the time dependent errors of the inertial positioning to be constrained. The resultant system combines accurately derived relative positions and attitude at a high measurement rate, typically 50-100 Hz, from the INS, with the long-term positional stability of GPS. The high data rate characteristic is particularly important when each individual scan line in an image requires orientation parameters. By integrating GPS and INS, positions, velocities,

accelerations and attitudes can be determined to a higher degree of accuracy than if either system were used alone (Napier and Ashkenazi, 1987; Schwarz *et al.*, 1993).

A full inertial system consists of three mutually orthogonal accelerometers and three single-degree-of-freedom (or two twin-degree-of-freedom) gyroscopes (Tait, 1990). The basic theory underlying inertial navigation is Newton's second law describing particle motion in a gravity field with respect to an inertial frame. In such a frame we could write:

$$\mathbf{f} = \mathbf{a} - \mathbf{g} \tag{2.1}$$

where;

\mathbf{f} is the specific force,

\mathbf{a} is the vehicle acceleration, and,

\mathbf{g} is acceleration due to gravity.

Of course, on the earth (a rotating reference frame) there are additional factors contributing to the sensed acceleration, this is a much simplified expression. By measuring the specific force, the vehicle acceleration can be extracted from equation 2.1 assuming a known gravity signal along the vehicle trajectory. The vehicle's velocity can be obtained by integrating the acceleration with respect to time. Relative position is then acquired by integrating velocity with respect to time. Gyroscopes measure angular velocity with respect to inertial space, to obtain attitude components the angular velocity is integrated with respect to time.

Inertial navigation systems can be divided into two main groups; stable platform systems and strapdown systems. Stable platform systems establish the internal attitude reference mechanically and provide platform orientation with respect to some prescribed co-ordinate system. In strapdown systems, the same process is done analytically, i.e. angular velocities with respect to the body frame are measured at a high rate and orientation changes with respect to the prescribed frame are computed. Although both types of system could be used in direct georeferencing applications, strapdown systems are far more common due to their

advantages in terms of data rate, attitude output, rate, price, power requirements and weight (Schwarz *et al.*, 1994).

Schwarz *et al.* (1994) conducted a comprehensive survey of the instruments available at that time, some of the findings of this work are presented below. Table 2.3 summarises INS position, velocity and attitude accuracies. The broad range of available instruments has been divided into three accuracy classes; low, medium and high.

Error in	System Accuracy Class (rms)		
	High	Medium	Low
Attitude			
1 hour	10" – 30"	1' – 3'	1° – 3°
1 min	1" – 2"	15" – 20"	0.2° – 0.3°
1 second	< 1"	1" – 2"	0.01° – 0.03°
50 Hz (noise)	0.1" – 0.2"	0.1" – 0.2"	15" – 20"
Velocity			
1 hour	0.3 – 0.5 m/s	1 – 2 m/s	200 – 300 m/s
1 min	0.01 – 0.02 m/s	0.05 – 0.1 m/s	1 – 2 m/s
1 second	0.0005 – 0.001 m/s	0.001 – 0.003 m/s	0.002 – 0.005 m/s
50 Hz (noise)	0.0002 – 0.0005 m/s	0.0005 – 0.002 m/s	0.001 – 0.003 m/s
Position			
1 hour	0.3 – 0.5 km	1 – 3 km	200 – 300 km
1 min	0.3 – 0.5 m	0.5 – 3.0 m	0.30 – 50 m
1 second	0.01 – 0.02 m	0.03 – 0.10 m	0.3 – 0.5 m
50 Hz (noise)	0.005 – 0.01 m	0.01 – 0.05 m	0.05 – 0.10 m

Table 2.3: INS position, velocity and attitude accuracies (Schwarz *et al.*, 1994)

The one hour time interval represents the long term suitability of using INS as a stand-alone georeferencing system. A one minute interval gives an indication of the short term interpolation accuracy that could be achieved during GPS outages, including periods of cycle slip detection and fixing. The one second interval characterises the interpolation time for an integrated GPS/INS system. It provides a useful guide to the rate at which GPS positions would need to be recorded to maintain a specified overall accuracy level.

The figures in table 2.3 show that only high accuracy inertial systems can be used for stand-alone georeferencing, and even with these systems, the positioning accuracy is still marginal. The position accuracies for all classes clearly illustrate the need to aid the inertial positions with some form of external measurements. It should also be noted that the attitude accuracies given are for the pitch and roll components, the errors in the heading component may be three to five times larger. This contrasts with GPS based systems in which the heading component is generally the most accurate.

So far the discussion has been limited to full inertial systems, i.e. those with accelerometers and gyroscopes that can provide values for all six elements of exterior orientation. There is also a wide range of gyroscopic instruments available, which in recent years have been increasingly incorporated into direct georeferencing systems. These instruments only measure angular rates, which provide attitude components but do not provide positions. If a gyroscope is used instead of a full inertial system, then position measurements, if required, could be provided by an external source such as GPS.

The increasing use of gyroscopes in airborne applications has largely resulted from the development of fibre optic gyroscopes (FOGs). Until the early 1990's the majority of gyroscopes were mechanical devices. These are generally large, heavy, expensive and, due to their mechanical components, are not rugged enough for airborne applications. Their slow response times also makes them unsuitable for measuring the high dynamics of a small survey aircraft. The fibre optic gyroscope however, is particularly well suited to dynamic environments, it is rugged, totally solid state, and is potentially much more precise and considerably less expensive than a mechanical instrument.

The fibre-optic gyroscope is based on an interferometer, and operates via a phenomenon called the Sagnac effect. A split light beam is sent in opposite directions through a coil of glass fibre. Any rotation around the coil axis causes one beam to travel further than the other, which induces a phase difference between the beams. This shift relates directly to the rotation rate, and is measured by a photo-detector. The sensitivity of the interferometric FOG is proportional to

the length and diameter of the fibre coil. The small size of the Sagnac effect demands long fibre lengths together with sensitive detection equipment. For example, 1 km of fibre wound around a diameter of 10cm gives a phase shift of about $1\mu\text{rad}$ for a rotation rate of $0.01^\circ/\text{h}$ (Graydon, 1997).

An alternative approach is to determine attitude from GPS carrier phase signals measured by an array of antennas. A number of commercially available systems have been developed, e.g. the Ashtech 3DF and ADU2, and Trimble TANS Vector, that consist of either three or four GPS antennas rigidly fixed together and linked to a special receiver that ‘measures’ the phase differences between the carrier signals received by different antennas. Essentially, kinematic phase GPS processing takes place between these antennas, and the co-ordinate differences obtained are converted to attitude components – usually pitch, roll and heading. These dedicated GPS attitude systems are discussed in section 3.3.

Attitude determination can also be achieved by connecting each antenna in the array to an independent GPS receiver. There are two general approaches to determining attitude from data collected in this way. In the first, GPS phase data are collected, vectors between antennas are computed, and the orientation parameters are estimated, using least squares if more than the minimum of three antennas have been deployed. In principle, any of the currently available methods can be used for the critical part of the processing – the estimation of the baseline vectors. The second method determines attitude parameters directly from the GPS observables. Both these methods are described in section 3.4.

The accuracy of GPS based attitude determination is a function of the baseline length between the antennas, the longer the separation the higher the accuracy, and multipath. Any unmodelled flexure of the system, which is assumed to be a rigid body, will also induce errors. In practice, the length of possible antenna baselines on small survey aircraft is limited, but it is still possible to use a GPS based attitude method for direct georeferencing in applications at the lower end of the accuracy spectrum. Table 2.4 presents GPS attitude accuracy as a function of baseline length from a range of sources and indicates how these figures were

derived. In each case, the values are given for the antenna baseline used in the original study. Figures in italics have been derived based on the assumption that attitude accuracy is proportional to baseline length. Giving attitude values for three common baseline lengths (one, five and ten metres) allows easier comparisons to be made between the results from different studies. Where a distinction is made in the original work between heading, roll and pitch components this is reflected in the table. The heading component will generally be the most accurate as it does not rely on the vertical component of GPS, which is the least well determined. In practice roll is usually the least well determined component in airborne applications, due to the antenna placement. Either the antennas are placed on the wing-tips, which introduces errors due to wing flexure, or if the antennas are not placed on the wings, then it is only possible to form a very short baseline. Where a distinction was not made between the accuracy in each component, a general figure for all components is given. The figures also reflect whether a single attitude accuracy value, or a range, was provided.

As the achievable accuracy of GPS based attitude determination and the factors that dictate this accuracy are key elements of this research project, this subject is dealt with in greater detail later in this thesis. At present the aim is to place GPS based systems in context by comparing their performance with more widely used inertial systems.

Table 2.4: Attitude accuracy from GPS based systems

Source	Antenna Separation	Attitude Accuracy (rms)	
Schwarz <i>et al.</i> 1994 Summary of previously published results from Schwarz and El-Mowafy, 1992 ⁽¹⁾ and Schade <i>et al.</i> , 1993 ⁽²⁾	1m	10' – 30'	
	5m	4' – 6'	
	10m	2' – 3'	
El-Mowafy and Schwarz, 1994 ⁽¹⁾ .	3m	Heading 3'	Roll/Pitch 6'
	1m	10'	20'
	5m	2'	4'
	10m	1'	2'
Ashech ADU2 Specifications, 1996 ⁽³⁾	1m	Heading 10'-15'	Roll/Pitch 20'- 25'
	5m	2'-3'	4'- 5'
	10m	1' – 1.5'	2' – 2.5'
Cohen <i>et al.</i> , 1994 ⁽¹⁾	7.9m	3'	
	1m	25'	
	5m	5'	
	10m	2.5'	
Mahmud, 1999 ⁽⁴⁾	30m	Heading 10''	Roll/Pitch 40''
	1m	5'	20'
	5m	1'	4'
	10m	30''	2'

⁽¹⁾Used a high accuracy INS as a reference.

⁽²⁾Compared GPS derived sensor attitude with values from inverse photogrammetry (i.e. based on ground control)

⁽³⁾Manufacturers specification (no reference made to specific calibration tests).

⁽⁴⁾Estimated by equating changes in attitude from a static array with attitude accuracy.

Clearly there are some significant differences in the accuracy values presented, but if only the lowest accuracy components are considered (i.e. exclude heading component where it has been considered separately) a degree of consistency is apparent. For a 1m antenna separation accuracy values range from 20 to 25 arc-minutes, this improves to 4 to 6 arc-minutes for a 5m separation and to between 2 and 3 arc-minutes for a 10m baseline. Referring to the sensor orientation requirements presented in section 2.3.2, shows that GPS-based systems can meet the attitude accuracy specifications for resource mapping applications and, if a sufficiently large antenna separation can be implemented, they can also be used for more detailed remote sensing applications. The

performance of GPS-based systems is most closely matched to the accuracy requirements of sensors using CCD arrays or line scanners, particularly those with larger pixel sizes.

It is well known that GPS positioning accuracy varies with the mode of operation. Shaw *et al.* (2000) give values for the present and future stand-alone accuracy of GPS, i.e. the instantaneous positioning accuracy that can be expected from a single GPS receiver without a reference station or other augmentation. Figure 2.2 shows the horizontal accuracy (95%) based on estimates of the UERE budget and assuming a HDOP of 1.5.

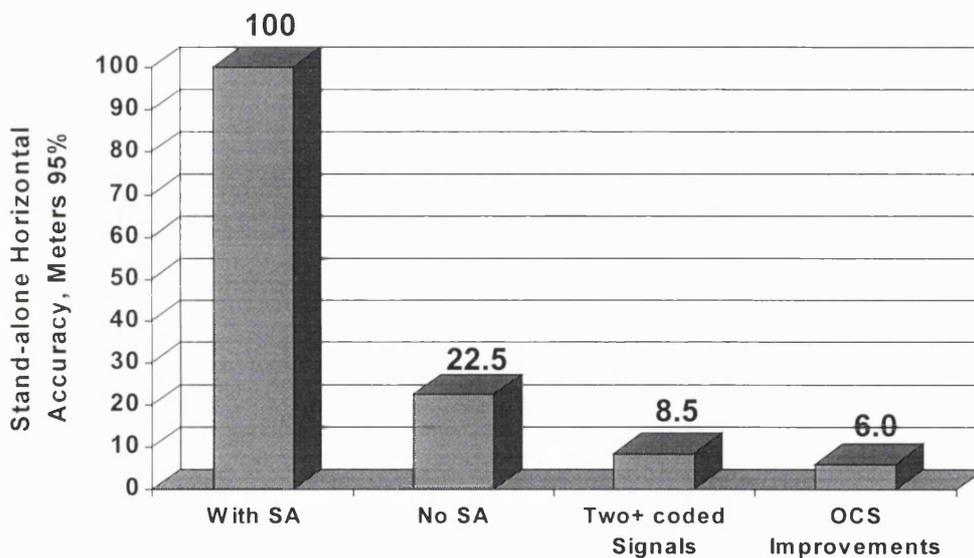


Figure 2.2: GPS stand-alone horizontal accuracy. (Shaw *et al.* 200)

Since the removal of S/A in May 2000, horizontal accuracies below 20m are achievable, and significantly higher accuracies than these have been widely reported. Hill and Moore (2000) for example, found that horizontal accuracy (95%) from a single stand-alone receiver is around 8m. With the introduction of dual C/A code satellites from 2003 onwards, and the proposed improvements to the operational control segment (OCS), these accuracies will improve still further. Additionally from 2008 it may be possible to achieve positioning accuracies of around 4m using the Galileo constellation. Despite these potential improvements, and the fact that in practice the performance of GPS tends to exceed official specifications, these levels of accuracy still fall somewhat short of the

requirements for the majority of applications and sensors discussed previously. It is worth noting however, that the 2 to 5m (rms) specifications given for resource mapping could potentially be achieved by a single receiver by 2008 (the proposed date for GPS dual code IOC, and possibly the FOC of Galileo).

In general, some form of differential positioning mode is required to meet the higher position accuracy requirements for exterior orientation⁴. Table 2.5 presents some guides to achievable accuracies when operating in a differential mode. The accuracy is related to the distance between the remote, airborne, receiver and the fixed reference station, due to spatially correlated atmospheric effects.

Mode of Operation	Accuracy (95%)	
Smoothed pseudorange differential positioning	10 km	0.5 – 3m horizontal 0.8 – 4m vertical
	500 km	3 – 7m horizontal 4 – 8m vertical
Carrier phase differential positioning	10 km	3 – 20cm horizontal 5 – 30cm vertical
	50 km	15 – 30cm horizontal 20 – 40cm vertical
	200 km (with precise orbits, same as 50km)	

Table 2.5: GPS differential positioning accuracy (adapted from Schwarz *et al.* 1994)

These figures indicate that with an appropriately designed survey, using dual frequency GPS receivers and restricting the reference/rover separation, even the highest levels of accuracy required for positioning aerial cameras can be achieved using GPS. GPS can also be used in combination with inertial systems so that the short term relative accuracies given in table 2.3 can be achieved over longer periods, and the co-ordinates can be determined in a more transferable reference frame (WGS 84).

One significant limitation of GPS for providing exterior orientation values is its recording interval. Until the last few years, GPS receivers with recording rates

⁴ Concepts such as local and wide area augmentation, precise orbit determination and virtual reference stations that can improve the accuracy from a single receiver are not dealt with here.

higher than 1Hz were unusual. Using scanning devices that may record 100 lines per second with this measurement rate clearly requires a significant amount of interpolation to attach orientation values to each scan line. No matter how sophisticated the interpolation technique, it is unlikely to adequately model the high frequency motions experienced by small survey aircraft. The orientation values from GPS will therefore depend not only on the accuracy of GPS, but also on the accuracy with which values can be transferred to scan lines. This limitation has become less significant in the past few years with the increasing availability of high data rate GPS receivers, 10Hz and 20Hz systems are now widely available, but GPS recording rates are still significantly lower than those of inertial systems. It is worth noting that with higher rate GPS receivers it is possible to constrain the inertial position drifts to a greater extent and thus improve the positioning accuracy of a combined system.

A GPS solution may also be subject to problems of discontinuity, caused by cycle slips and loss of lock. Continuous tracking can be particularly difficult to achieve in airborne applications as the signal from a satellite to an antenna may be blocked during manoeuvres. Attempts are made to prevent such problems by limiting the severity of manoeuvres and by resuming level flight well before the survey area is reached. This gives the ambiguity resolution algorithm as much continuous data as possible to try and recover integer ambiguities before the start of a flight line. If the data from a flight are post processed, as they will be in the majority of cases, forward and backward processing techniques can be used to provide a solution for the problem period using data taken from before and after the event. Another method to ensure continuity is to use an INS to provide positioning information during an outage period. INS data can also be used to aid the ambiguity recovery process by providing position estimates to define the search space. The ability to use INS data to bridge gaps in GPS positions is another example of the benefits of integrating these two systems. Depending on the level of integration, it is possible to improve the positioning accuracy from an INS using GPS range observations. Even when there are too few satellites to allow a position estimate from GPS alone, the addition of a single range to a combined solution can help to constrain time dependent inertial errors.

Having discussed the performance of INS and GPS systems in terms of position and attitude accuracy it is clear that they have the potential to meet all or some of the exterior orientation requirements presented. There is a wide range in the performance levels that can be achieved by specific GPS and INS devices. The performance will also vary depending on the survey design, e.g. mission duration, distance from reference station, aircraft size, etc. It is also apparent that neither system can independently meet all specifications for the duration of a typical aerial survey. By combining the complementary characteristics of GPS and INS it is possible to design systems capable of meeting the accuracy requirements of a far greater range of applications and sensors. Section 2.3.3.3 discusses how data from INS and GPS can be integrated to realise this potential.

2.3.3.3 Processing

Just as the choice of navigation equipment can be tailored to the needs of a particular sensor or application, so too can the processing scheme. The processing scheme will also be intrinsically linked to the instruments used and the nature of the final product required, e.g. paper map, digital terrain model, orthophoto. In this section some general comments are made on integrated processing schemes and then two specific cases are discussed in more detail. The processing schemes highlighted have been chosen as they represent some relatively typical approaches and because the models are well described in public domain literature. Comments are made on the strengths and weaknesses of some common approaches and a number of current trends in system design are discussed.

GPS and INS measurements can be combined to varying degrees in the system's processing scheme. As a simple example a GPS receiver could be used to provide the sensor position with a gyroscopic instrument providing all the required attitude information (the inertial system does not need accelerometers if it is not used for positioning). The observations from each instrument may be fed into a filter but essentially the GPS and gyro solutions are independent. Position and attitude information are only combined when orientation parameters are applied to the imagery.

In most applications however, the complementary nature of the two technologies is exploited in an integrated system. A standard way to do this is to use a Kalman filter. “A Kalman filter is a multiple-input, multiple-output, digital filter that can optimally estimate, in real time, the states of a system based on its noisy outputs. These states are all the variables needed to completely describe the system behaviour as a function of time”, Levy (1997). Essentially, it produces a best estimate of the system’s state based on all past and present information. Details of the theory and application of the Kalman filter, including the mathematics, are given in numerous texts, for example, Brown and Hwang (1992) and Gelb (1974).

In the majority of integrated systems the INS is the primary navigation device and provides all six elements of exterior orientation. In section 2.3.3.2 it was noted that when using inertial systems significant position errors accumulate over time. To achieve the accuracy requirements discussed in section 2.3.2 position data from an external source must be available, this process can be termed ‘inertial aiding’ (Brown and Hwang, 1992). GPS positions are now the most commonly used form of external data, due to the accuracy and applicability of GPS. Figure 2.3 is a block diagram of an aided inertial navigation system.

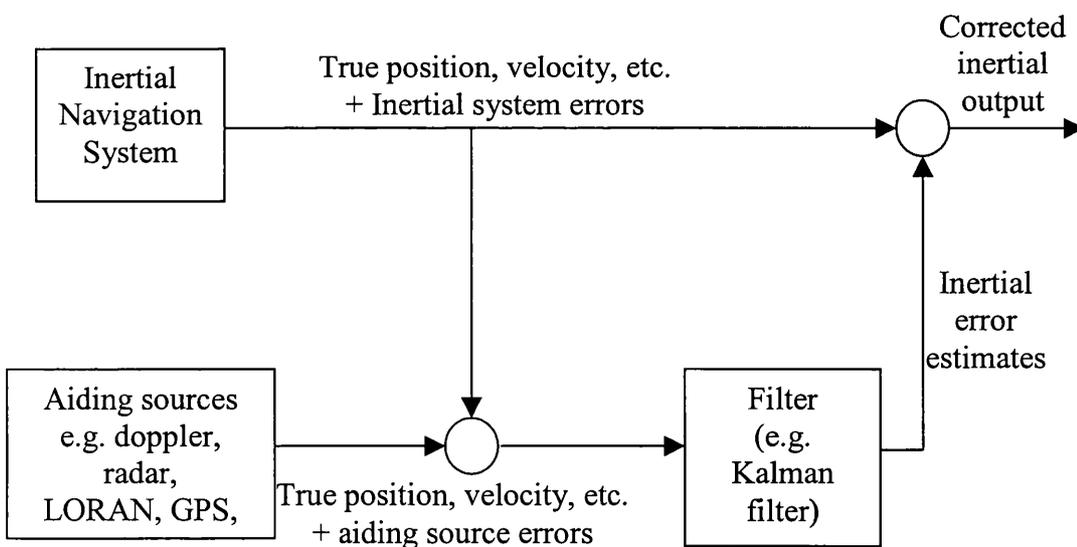


Figure 2.3: Feedforward configuration for an aided inertial navigation system. Linearised Kalman filter (Brown and Hwang, 1992)

This diagram describes the typical situation in which the filter estimates the deviation of a reference trajectory from the true trajectory. When an inertial system is used, a nine dimensional reference trajectory (three components of position, velocity and attitude) will normally be determined. The filter states are the system errors from the inertial and aiding instruments rather than the actual position and attitude values. This is important in the navigation problem, because vehicle position and velocity are usually *not* properly modelled as random processes (Brown and Hwang, 1992). Estimates of the inertial errors from the filter are then used to correct the original position, velocity and attitude values, which have passed through the system without distortion or delay. Values from the aiding device, which may only be available intermittently, are used as filter inputs in the measurement model to provide better estimates of the system errors. The integration scheme shown in figure 2.3 is a regular linearised Kalman filter.

It should be noted that within the filter scheme only *like* quantities can be differenced (e.g. GPS *velocity* can be compared with INS *velocity*) and they must be compared in a common reference frame. This will often mean that inertially derived quantities need to be converted to an appropriate co-ordinate system before comparisons are made. The INS is chosen as the primary system as it is able to measure all the required quantities. Any additional devices that measure some of these same quantities can be described as aiding. Aiding devices do not need to measure all the quantities recorded by the inertial system, in practice it would be rare to have a supplementary instrument measuring all the INS derived quantities, even one common measurement provides some degree of aiding.

Figure 2.4 presents a variation to the previous integration scheme. In this scheme the inertial system errors are fed back, and the corrections are made internally in the INS, i.e. the system outputs are not raw uncorrected values. This is the principle of a feedback configuration. The main difference between the two approaches is in the computation of predictive measurements. In the previous feedforward scheme this is computed based on the nominal reference trajectory, which may slowly diverge from the true trajectory over time. In the feedback scheme the predictive measurements are based on a continually updated

trajectory, based on the filter's best estimate of the inertial system's errors. This produces an extended Kalman filter mode of operation rather than an ordinary linearised filter. The extended filter is preferable in applications with a long mission time in which the reference trajectory can diverge from the truth beyond acceptable limits.

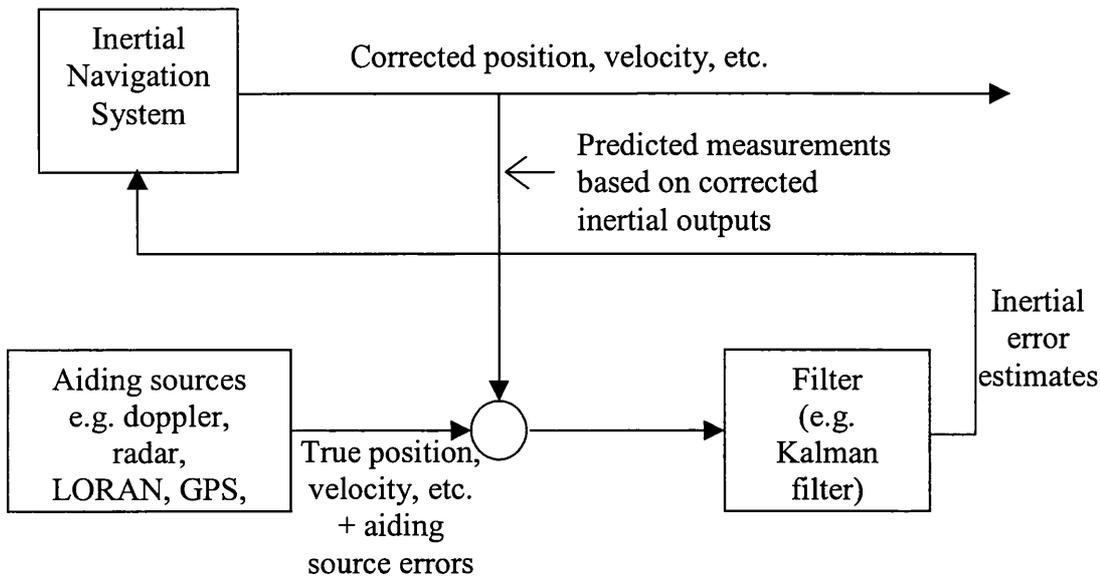


Figure 2.4: Feedback configuration for an aided inertial navigation system. Extended Kalman filter (Brown and Hwang, 1992)

Having briefly described one common integration process, that of *inertial aiding*, at a general level, two specific processing schemes from the vast array of different designs put forward are introduced. Figure 2.5 illustrates the decentralised Kalman filter configuration for INS/GPS integration used in various projects undertaken at the Department of Geomatic Engineering, University of Calgary (Skaloud, 1995; Skaloud *et al.*, 1996; and Skaloud and Schwarz, 1998). In a decentralised configuration there are two filters working independently: an INS filter and a GPS filter. Double differenced pseudorange, carrier-phase and phase rate observations, form the measurement vector in the GPS filter. The GPS error states are the position and velocity errors augmented by double differenced ambiguity parameters. It outputs position and velocity error information, which is used to update the INS master filter. The INS error model has 15 states made up of navigation errors (attitude, position, velocity) and sensor noise terms (gyro drifts and accelerometer biases). Updated error states from the INS master filter are fed back so that corrections can be applied to the raw inertial measurements in

the INS processor, i.e. this is an extended Kalman filter. Thus the output from the INS processor contains integrated INS/GPS position, velocity and attitude information. This information can be used to check the validity of the GPS measurements and to aid carrier phase ambiguity resolution following cycle slips or loss of lock. The output of the GPS processor can be checked for reliability by comparing the states derived from measurements at an epoch with the predicted states from the dynamic model. Detecting outliers at this stage will prevent any erroneous GPS data contaminating the master filter.

The second example (figure 2.6) shows the tightly coupled GPS/INS processing scheme of the Airborne Integrated Mapping System, (AIMSTM) developed at the Centre for Mapping at The Ohio State University (Da *et al.*, 1996; Da, 1997; Grejner-Brzezinska *et al.*, 1998; Toth and Grejner-Brzezinska, 1998). Figure 2.6 is an overview of the integration process. In the Kalman filter, raw inertial measurements are combined with double differenced GPS observations to estimate errors in the inertial navigation solutions, the IMU measurements, and GPS double differenced observations. Estimates of the inertial error states are fed back to the inertial computation module to improve the inertial solution accuracy, while estimates of GPS observation errors and inertial navigation solutions are sent to the GPS computation module to support cycle-slip detection and fast and reliable ambiguity resolution. Processing of raw inertial and GPS observations allows optimal utilisation of the observations provided by both systems (Da, 1997).

A tightly integrated system can provide precise positioning even if the number of satellites drops below four. In particular, the total loss of lock of about 60 seconds can be properly bridged and ambiguities recovered instantly after tracking is recovered (Grejner-Brzezinska *et al.*, 1998). Despite the potential performance improvements of tightly integrated systems they have been less prevalent than loosely coupled systems mainly due to the increased computation load demanded by the feedback loop. The rapid increase in computational power is facilitating a move towards tightly coupled systems (Burman, 2000).

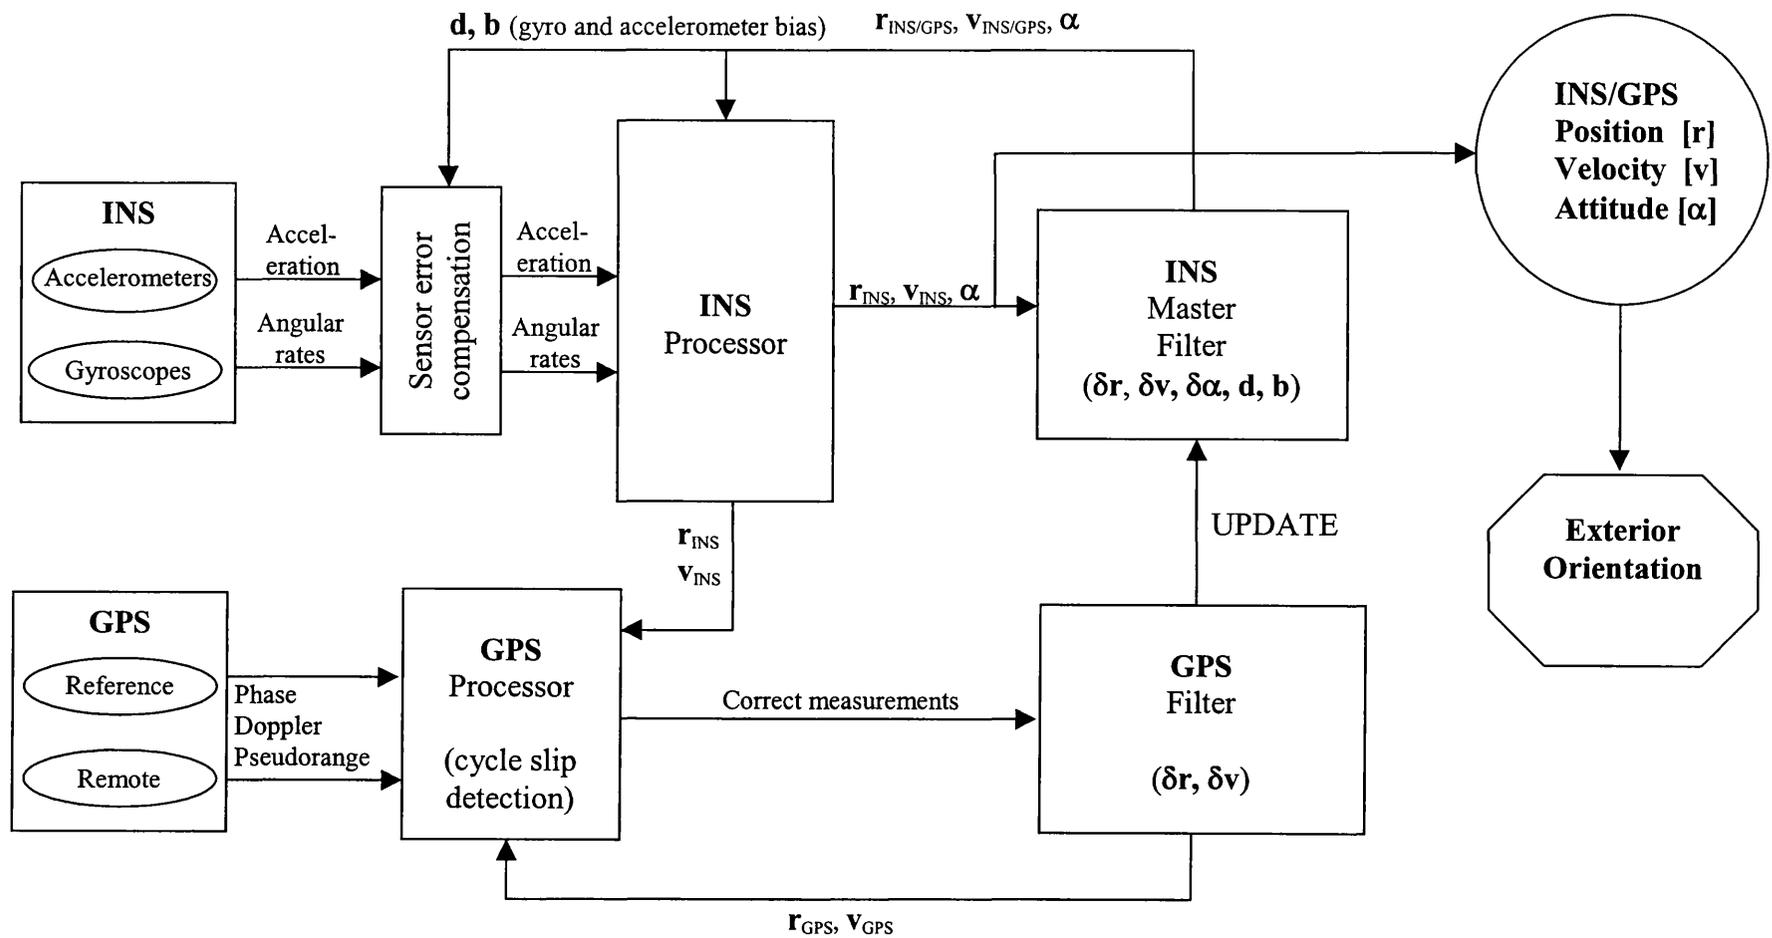


Figure 2.5: Block diagram of a decentralised Kalman filter configuration for INS/GPS integration (adapted from Skaloud, 1995).

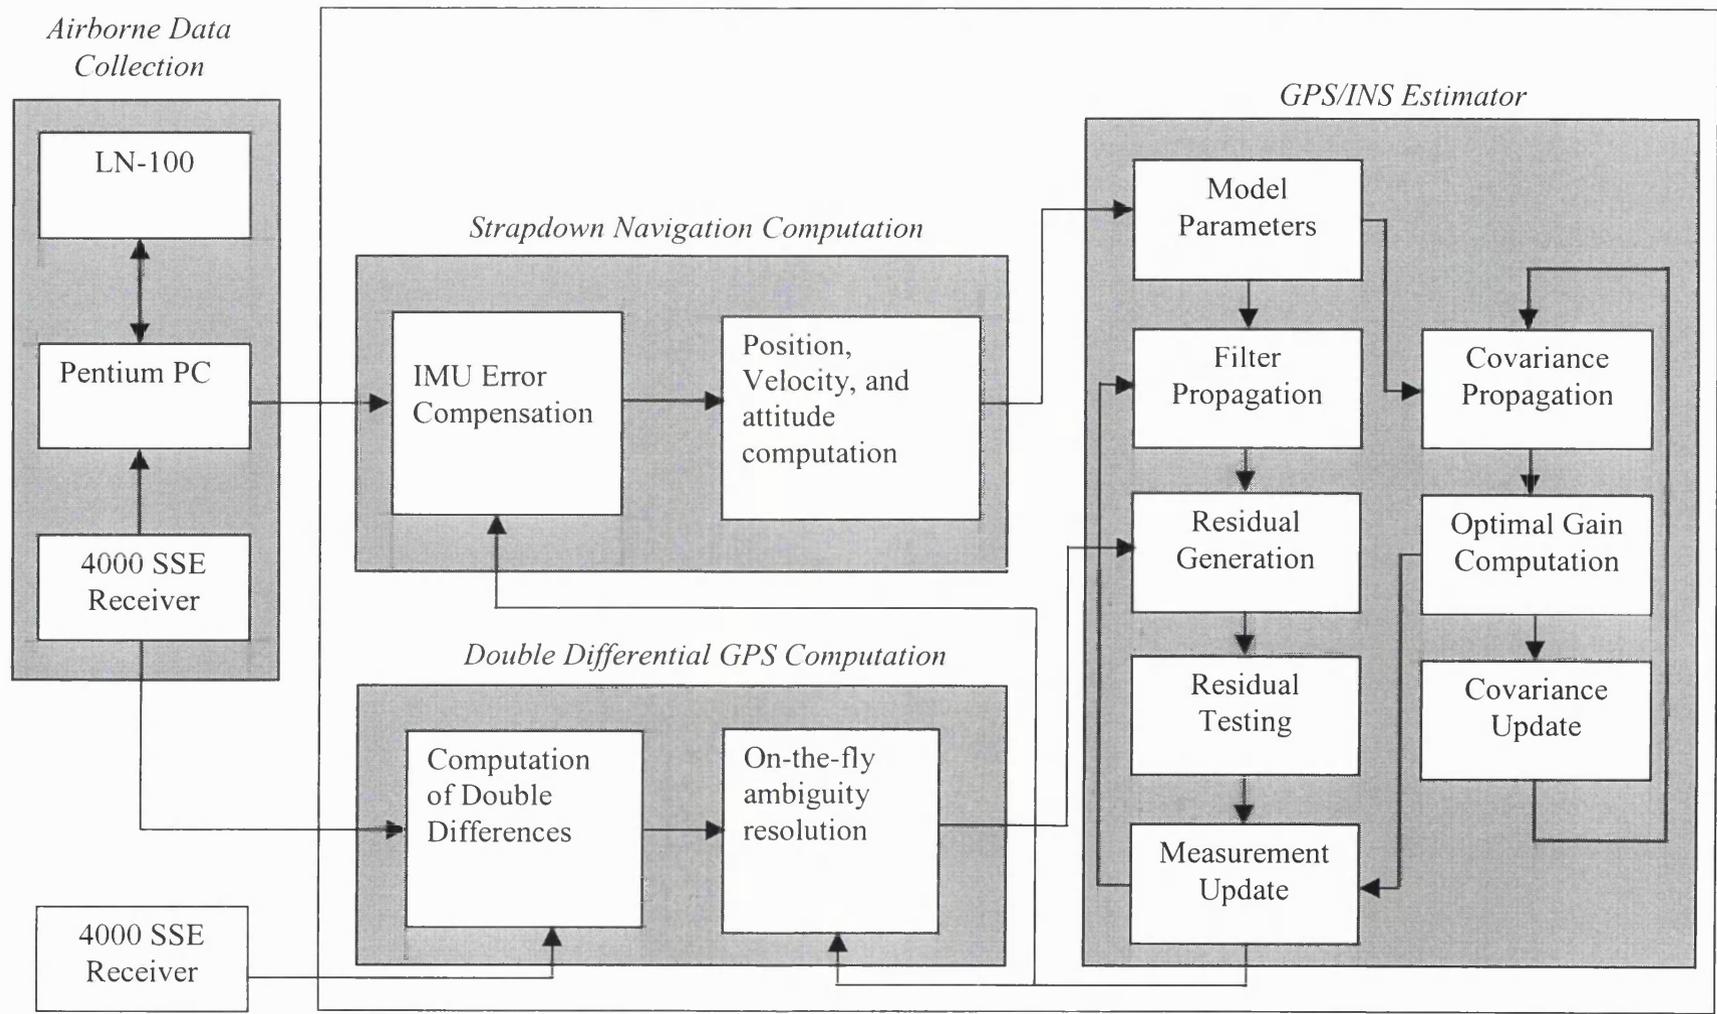


Figure 2.6: Block diagram of the AIMS tightly integrated GPS/INS system (Da, 1997)

In order to estimate the state, a Kalman filter depends on a measurement and a dynamic model. The quality of the final estimates of the states depend therefore on the quality of both the measurements being made and the models being used (Bruton, 1999). Ultimately then, parameter estimation relies on the estimated measurement covariance matrix and the process covariance matrix. One method of improving the estimates of these covariance matrices is to adopt an *adaptive* or *self-learning* procedure. In an adaptive filtering procedure residuals from the previous segment of processing are used to improve the estimated covariance matrices. Within a single segment the covariance matrices for each epoch are assumed to be the same. An adaptive stochastic modelling approach therefore has two main steps: firstly the derivation of appropriate formulae to determine the covariance matrices, and secondly the choice of segment or window size. Wang *et al.* (1999) discuss the benefits and application of adaptive filtering in GPS, GLONASS and INS integration. They conclude that in the case of GPS/INS integration, adaptive procedures can improve the accuracy of the final attitude components. Mohamed and Schwarz (1999) have also demonstrated that the adaptive Kalman filter outperforms the conventional filter by tuning either the measurement covariance matrix, the process covariance matrix, or both. Similarly, an adaptive procedure was chosen by Campana *et al.* (1999) for GPS attitude determination as again it was able to provide more appropriate stochastic modelling.

The trend towards tightly integrated systems and the use of adaptive techniques both demonstrate the desire to fully exploit all the information provided within the processing scheme. By feeding as much information as possible back into the modelling and estimation stages the overall system performance can be improved. The system's ability to adapt becomes increasingly important as more flexible georeferencing configurations are developed. A variety of measurement devices may provide inputs to the same processing engine depending on the application, or the trajectory characteristics may vary if the navigation system is flown on a helicopter rather than a fixed wing aircraft.

This section has attempted to give a general picture of integrated processing schemes. The main objective in all these systems is to exploit the complementary characteristics of INS and GPS technologies to obtain optimum performance. The best way to achieve this will vary with the application and the navigation instruments used. Just as the choice of equipment can be tailored to a particular application, so too can the processing scheme.

2.3.3.4 System Performance

Having discussed the characteristics of GPS and INS instruments, and then described how they can be integrated, the performance of some integrated systems is now described. By mapping performance to requirements, the suitability of a direct georeferencing system for a particular application and sensor can be assessed. Initially the types of navigation instruments required to meet particular accuracy criteria are discussed. Following this general outline, the performance of some specific direct georeferencing systems is discussed. It should be noted once again that there are a wide range of direct georeferencing systems currently in use and under development, and that the following selection describes a small but representative sample.

Table 2.6 (from Schwarz *et al.*, 1994) presents possible sensor and georeferencing configurations for particular accuracy ranges. The costs were based only on hardware costs (in 1994) and do not include any costs associated with interfacing, developing dedicated processing software, or installing such systems on an aircraft.

Accuracy Required	System Configuration	Airborne Sensor	Cost (K\$ = \$1000)	Characteristics
0.05 – 0.1m 15” –30” 0.0002 – 0.0005 m/s	High accuracy INS plus DGPS (carrier phase)	All airborne sensors	K\$250	Suitable for all applications
0.05 – 0.01m	DGPS (carrier phase)	Photogramm- etric camera	K\$50	Block adjustment only
2 - 5m > 10’ 0.01-0.02 m/s	DGPS (psuedo- range) + GPS multi-antenna system	Line scanner, CCD frame imagers	K\$50	Low data rate. Attitude transfer to sensor is problematic. (stability)
2 – 5m > 1’ 0.0002 – 0.0005 m/s	DGPS (psuedo- range) plus medium accuracy INS	Line scanner, CCD frame imagers, some SAR applications	K\$100 -120	Long term velocity may be marginal

Table 2.6: Possible sensor configurations (Schwarz *et al.*, 1994).

The table shows that it is possible to develop direct georeferencing systems to meet the full range of accuracy requirements. The higher accuracy configurations may impose some operating constraints, for example the use of carrier phase GPS positioning will restrict the distance that the aircraft can travel from the reference station to around 15km. In all cases the position and alignment of the navigation devices must be related to the sensors.

If high accuracy applications are excluded, it is possible to find economic configurations suitable for most applications in the resource mapping/land-use sector. Schwarz *et al.* (1994) conclude that two solutions seem particularly attractive: one is completely based on GPS technology, the other integrates GPS with a low cost INS.

Table 2.7 presents details of the design and performance of some selected direct georeferencing systems. It gives the position and attitude accuracies of the system, with details of the equipment (navigation and imaging) and processing methods where possible.

Source	Accuracy		Configuration
	Position	Attitude	
University of Calgary (Skaloud <i>et al.</i> 1996; Skaloud and Schwarz, 1998)	15cm horizontal 20cm vertical (rms)	40" Heading (σ) 2' Pitch/Roll (σ) 10–20" Heading(rms) after 1998 pre-filtering	Medium accuracy INS [Litton LTN-90-100] DGPS (carrier phase) [Trimble 4000 SSE, Ashtech Z12] Aerial Camera [Zeiss RMK A] Processed in a decentralised Kalman filter (figure 2.5)
AIMS™ Centre for Mapping The Ohio State University (Da <i>et al.</i> , 1996; Da, 1997; Grejner-Brzezinska <i>et al.</i> 1998; Toth and Grejner-Brzezinska, 1998)	4-7cm (rms)	< 10" (estimated/potential accuracy only – no results from calibration tests published to date, to the best of this author's knowledge)	Medium accuracy INS [Litton LN-100] DGPS (carrier phase) [Trimble 4000 SSI] CCD Camera [Lockhead Martin Fairchild Semiconductors] Tightly coupled Kalman filter (figure 2.6)
POS/AV Applanix (Specification from Applanix website http://www.applanix.com)	5-30cm (rms)	20" Pitch/Roll 30" Heading (510, highest accuracy version, rms) 2.5' Pitch/Roll 5' Heading (210, lowest accuracy version, rms)	Compact IMU DGPS (carrier phase) Suitable for a range of sensors Tightly coupled filter
POS/DG Applanix (Specification from Applanix website) (Test results in Lithopoulos <i>et al.</i> 1999; Burman 2000)	5-10cm system specification (rms) 20cm test results	20-30" system specification (rms) 30-45" test results	Compact IMU DGPS (carrier phase) Suitable for a aerial cameras Tightly coupled filter
C-MIGITS II Boeing (Martin and Detterich, 1997a, & 1997b) (Skaloud <i>et al.</i> , 1997)	4.5m (SEP)	8–9' Pitch/Roll 10' + δ Heading (δ = growth rate based on time since horizontal accelerations, typically 1-3°/h) 10' Pitch/Roll 20' Heading (U. of Calgary tests)	Low-cost [Digital Quartz IMU] Low power 5 channel L1 GPS [Rockwell MicroTracker™] Tightly coupled 28 state Kalman filter (flexible)

Source	Accuracy		Configuration
	Position	Attitude	
FLI-MAP® from John E Chance & Associates Inc. (Liu <i>et al.</i> 1997)	<10 cm	<6'	4 Independent GPS receivers on aircraft plus a reference station [Trimble 4000 SSE] Low cost vertical gyro Laser ranging scanner Helicopter mounted Kalman filter for integration
Saab Topeye (Burman, 2000)	1.2m	5' The accuracy of aerial triangulation values proved inadequate to properly calibrate the GPS/INS system	DGPS (carrier phase) [2 x Trimble 4000] High Accuracy INS [Honeywell H-764 INS] Laser Scanner
Institute of Geodesy and Navigation, IfEN University FAF Munich (Wolf <i>et al.</i> 1997) (static test)	<10cm	2' – 3' Pitch/Roll 10-15' Heading	Low accuracy INS [Systron Donner MotionPak] Multi-antennas GPS [Trimble TANS Vector] Tightly coupled 27 state Kalman filter

Table 2.7: Survey of direct georeferencing performance.

The completeness of the information in this table reflects the varying levels of detail in the source material. Where possible the accuracy measure is indicated (i.e. RMS, σ , or SEP) and a distinction is made between estimated and calibrated performance results. In the majority of calibration tests an aerial camera has been flown with the direct georeferencing system over a test site of surveyed targets. Camera orientations from inverse photogrammetry are then used as reference values for the directly derived orientation values. In the case of the AIMSTM and Topeye system, the sensor accuracy from photogrammetric processing has not been good enough to provide a suitable reference. Some of the lower accuracy systems, e.g. the C-MIGITS II tested at the University of Calgary, have been calibrated by flying them alongside much higher accuracy direct georeferencing systems, and treating the results from the higher accuracy system as a truth. As the accuracy of the exterior orientation provided by direct georeferencing

approaches the best achievable accuracies of conventional photogrammetry, it becomes increasingly difficult to find a suitable reference solution.

The instruments used in particular systems are subject to change, in line with the trend to tailored systems. In general use, most of the sensing devices flown will be digital (CCD cameras or scanners), but when calibrating systems it is often necessary to use a conventional camera as the resolution and geometric stability of the standard digital sensor may not be high enough to reliably transfer ground control accuracies to sensor orientation parameters. A number of the systems have been developed to meet the georeferencing needs of a variety of sensors, and most are of a modular nature to allow higher or lower accuracy navigation instruments to be used for particular applications.

The range of performance figures presented clearly cover the full range of sensor and application requirements discussed in section 2.3.2. In the commercial airborne remote sensing and mapping sector, direct georeferencing systems (particularly the range of Applanix products) are being used routinely for a variety of applications and the economic benefits of this method are already being realised. Surveys are being carried out with little or no ground control and products are being delivered more rapidly and at a lower cost. Some of the other systems described have been developed for a specific application, and in recent years operational experience has demonstrated that the development objectives are being realised. FLI-MAP for example, is being operated on low flying helicopters for corridor surveys of roads and pipelines. By operating at low altitudes, typically around 50m, this system allows relatively low accuracy attitude determination of the laser scanner, but still produces on-the-ground accuracies of 15 to 20cm.

The systems described paradoxically demonstrate both the diversity and similarity of direct georeferencing solutions. The generic tools of GPS and INS equipment and Kalman filtering are common to most systems, but these tools are implemented in such a way as to produce ad-hoc solutions to meet an ever increasing range of tasks.

2.4 Concluding Remarks

This chapter has introduced the concept of georeferencing and demonstrated that a key stage in the process of relating imagery to a mapping reference frame is the determination of the sensor position and attitude, i.e. exterior orientation. The conventional methods for finding orientation from ground control were discussed and some of the inherent limitations were identified. Direct georeferencing offers a solution to many of these problems and makes aerial surveys more cost-effective, less time-consuming and more flexible.

It is clear that direct georeferencing systems have evolved to meet the full spectrum of accuracy requirements for photogrammetry and remote sensing. These systems generally incorporate INS and GPS technology, and integrate the measurement data in some form of Kalman filter. Although many of the systems have much in common, they are increasingly being adapted to meet the needs of particular applications, providing ad-hoc solutions. One of the major trends in direct georeferencing over the past few years has been the move towards low cost systems suitable for lower accuracy remote sensing applications. These systems no longer use a medium or high quality INS as the primary navigation instrument, instead they incorporate GPS with a low cost INS, or rely on GPS alone for both position and attitude data. The system developed in the course of this research project fits well with this trend as it uses GPS for determining both position and attitude. Although it can operate with GPS only, it has the flexibility to incorporate additional high rate attitude data from a low cost gyro. The following chapter discusses the general subject of deriving attitude from GPS observations and then introduces the particular method used in this project. Chapter 4 then describes the system in terms of navigation and sensing equipment, and describes how the system has been calibrated. As more information is provided on this specific system, its place in the range of direct georeferencing solutions should become clear.

Chapter 3 GPS Attitude Determination

3.1 Introduction

Attitude information is required in a range of applications on land, sea, air and in space. In hydrographic surveying for example, the orientation of an acoustic sounder needs to be determined to relate depth measurements to geographic positions. The importance of attitude information in airborne photogrammetry and remote sensing was emphasised in section 2.3.2, and as this is the current field of study the majority of the following discussion will be related to the airborne case. The increasing interest in GPS based attitude systems as an alternative to inertial instruments is not restricted to the aerial survey market, this trend is repeated to varying degrees across all disciplines reliant on orientation information. The decreasing size and cost of GPS equipment over the last decade has made it an increasingly attractive option. At present a variety of techniques and equipment are being used for GPS attitude determination.

Initially this chapter introduces the principles of attitude determination with GPS (section 3.2). Having discussed the underlying theory, a number of approaches to GPS attitude determination are discussed in terms of equipment (section 3.3) and processing methods (section 3.4). Dedicated attitude instruments comprising a single receiver connected to multiple antennas are described, and so are more flexible solutions where a number of antennas each connected to an independent receiver are used to compute attitude. Trends towards low-cost systems using OEM GPS receiver cards, and the development of two antenna systems, that determine only two attitude components, are briefly discussed. The main discussion in terms of processing, concentrates on independent receiver systems, as this is the approach used in the current research. The trend towards more *direct* methods is highlighted, describing some of the advantages that this offers. Details of the particular processing method used in this research project are given in section 3.5. This algorithm computes attitude components directly from GPS phase observations by modifying the standard double difference observation

equations. Section 3.6 describes the main error sources in GPS attitude determination, and shows how the adopted method can improve the overall performance. Finally some concluding remarks are made.

3.2 Principles of GPS Attitude Determination

Attitude is the orientation of a rigid body with respect to an external reference frame. In the case of GPS attitude, this rigid body is related to the external reference frame using antenna positions. The co-ordinates of each antenna are determined in a rigid Body Frame System (BFS) prior to a flight and held fixed throughout¹. At any epoch during the flight the antenna positions can also be determined in an Earth Centred Earth Fixed (ECEF) system which in this case is WGS84. These ECEF co-ordinates can be converted to a Local Level System (LLS) which then acts as the external reference frame for the rigid body. Three rotation angles between the axes of the BFS and the reference system (LLS) can be used to describe attitude, typically these are referred to as pitch(ψ), roll(θ), and yaw(γ) [the Euler angles].

Firstly, the co-ordinate systems will be described, then the rotations that relate these systems, and therefore define attitude, will be introduced. The following definitions and the associated discussion follow the description given by Kleusberg (1995) and later repeated in Cross *et al.*(1997) and Mahmud (1999).

3.2.1 Co-ordinate Systems

The co-ordinate frames used within this application are the Earth Centred Earth Fixed (ECEF) system, the Local Level System (LLS) and the Body Frame System (BFS).

The BFS is fixed with respect to the vehicle whose attitude is to be determined. In general the BFS origin should be the centre of gravity of the subject vehicle, but

¹ In a permanent installation of course it would not be necessary to determine co-ordinates in the body frame prior to every flight. Values from a single survey could be used repeatedly, but should be checked occasionally.

for convenience the phase centre of a specified GPS antenna attached to the vehicle is often used. In this research the origin is chosen as the phase centre of antenna 1 (figure 3.1). The axes are defined so that the rotations describe the vehicle motion in terms of pitch, roll and yaw.

- The *roll angle* measures the rotation of the aircraft about the fuselage axis (Y)
- The *pitch angle* measures the rotation of the aircraft about the wing axis (X)
- *Yaw* measures the rotation of the aircraft about the vertical axis (Z)

If antennas cannot be arranged on the BFS axes, which will usually be the case in practice, then some angular offsets need to be applied to the measured attitude. The antenna locations in the BFS are determined prior to the attitude determination through standard GPS static surveying techniques or by conventional survey methods.

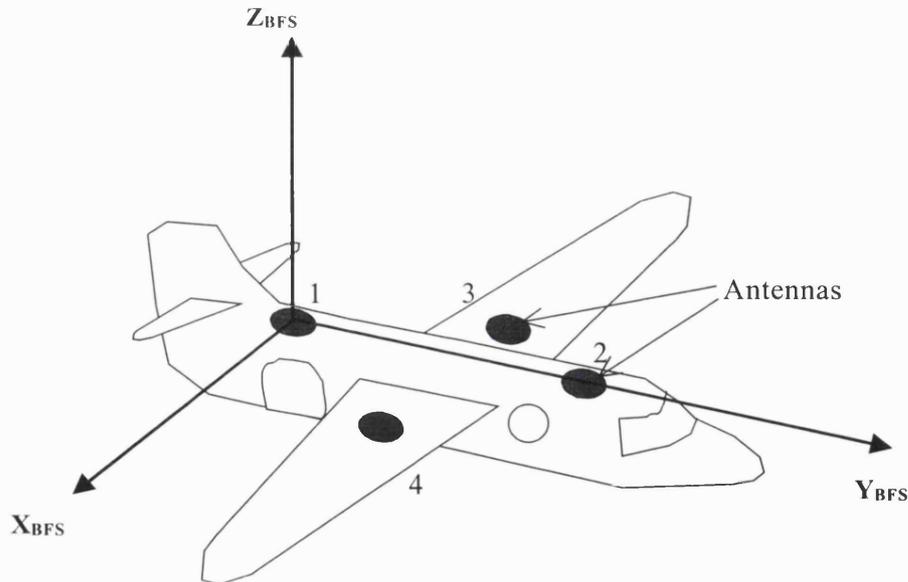


Figure 3.1: Body Frame System (BFS)

In this application the attitude is defined as the orientation of the BFS with respect to the LLS, as shown in figure 3.2.

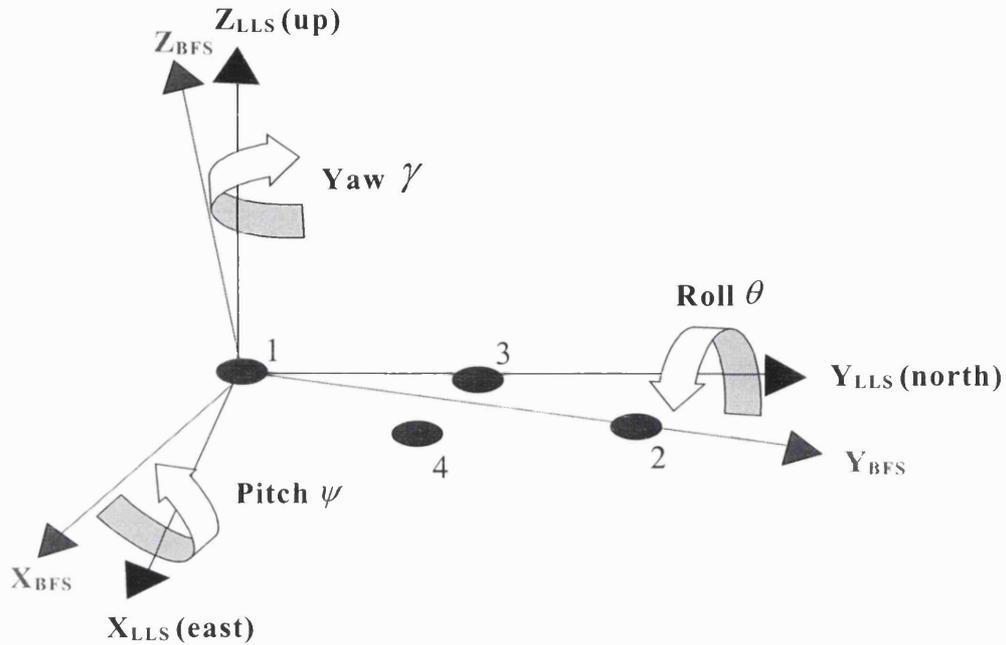


Figure 3.2: Attitude defined by BFS with respect to LLS

The LLS is defined as;

- Origin is at a chosen point on or near the earth's surface (X_0),
- Z-axis normal to the reference ellipsoid, pointing up,
- Y-axis pointing towards geodetic north
- X-axis completes a right handed cartesian co-ordinate system by pointing east.

In attitude applications the LLS origin is usually chosen to be coincident with the BFS origin, hence, in this research it is also located at the phase centre of antenna one.

At any epoch the mathematical relationship between the co-ordinates of any number of antennas (A, B, C, D, .. to N) in the two systems is expressed as

$$X^{LLS} = R_{\psi\theta\gamma} X^{BFS} \quad (3.1)$$

where

$$X^{LLS} = \begin{bmatrix} x_A^{LLS} & x_B^{LLS} & x_C^{LLS} & x_D^{LLS} & \dots & x_N^{LLS} \\ y_A^{LLS} & y_B^{LLS} & y_C^{LLS} & y_D^{LLS} & \dots & y_N^{LLS} \\ z_A^{LLS} & z_B^{LLS} & z_C^{LLS} & z_D^{LLS} & \dots & z_N^{LLS} \end{bmatrix} \quad (3.2)$$

$$X^{BFS} = \begin{bmatrix} x_A^{BFS} & x_B^{BFS} & x_C^{BFS} & x_D^{BFS} & \dots & x_N^{BFS} \\ y_A^{BFS} & y_B^{BFS} & y_C^{BFS} & y_D^{BFS} & \dots & y_N^{BFS} \\ z_A^{BFS} & z_B^{BFS} & z_C^{BFS} & z_D^{BFS} & \dots & z_N^{BFS} \end{bmatrix} \quad (3.3)$$

and

$R_{\psi\theta\gamma}$ is a composite rotation matrix containing nine direction cosine elements made up of the three attitude components.

Equation (3.1) is concerned with rotation only and not translation. This is because the origin of the LLS is coincident with the BFS origin. Also, because the body is assumed rigid, no scale factor is considered.

Co-ordinates determined from the measurements of a GPS receiver, will usually be in the WGS84 co-ordinate system (an ECEF system), not a local level system. A preliminary transformation must therefore be performed before attitude computations can take place. This is a standard transformation required in many applications as the final co-ordinates of a survey are often required in some local system rather than WGS84.

The ECEF used in this research is the World Geodetic System 1984 (WGS84). It provides “a single, common, accessible three-dimensional co-ordinate system for geospatial data collected from a broad spectrum of sources”, Malys *et al.* (1997). A consistent global set of three-dimensional station co-ordinates infers the location of an origin and the orientation of an orthogonal set of Cartesian axes.

As far as possible, the inferred WGS84 system has the following characteristics:

- Origin is at the centre of mass of the earth,
- Z-axis coincident with the mean spin axis of the earth,
i.e. positive towards the north celestial pole,
- X-axis is in the equatorial plane pointing towards the Greenwich meridian
- Y-axis completes a right handed Cartesian co-ordinate system.

Figure 3.3 illustrates the Earth Centred Earth Fixed system, and its relation to the Local Level System.

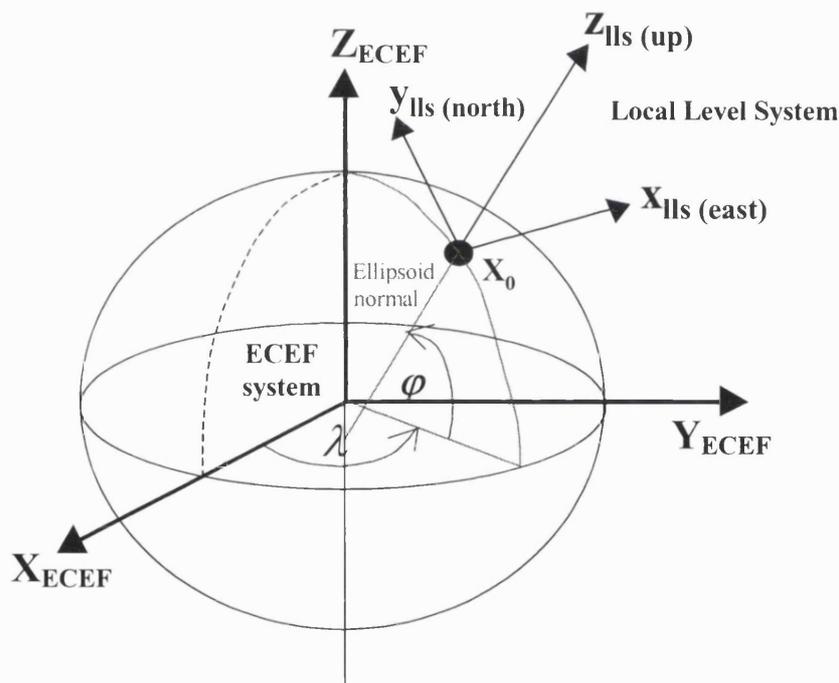


Figure 3.3: ECEF and LLS

With respect to the ECEF, the LLS is shifted along the position vector of the vehicles centre of gravity (X_0), and then rotated twice (Kleusberg, 1995). The relationship between co-ordinates in the LLS (X_i^{LLS}) and ECEF system can be expressed as,

$$X_i^{LLS} = R_x(90^\circ - \phi)R_z(90^\circ + \lambda)\Delta X \quad (3.4)$$

where,

- φ is the geodetic latitude of the LLS origin (X_o);
- λ is the geodetic longitude of the LLS origin (X_o);
- $R_x(90^\circ - \varphi)$ is the rotation matrix with respect to the vehicle's latitude by the angle $(90^\circ - \varphi)$ of a right-handed co-ordinate system around the x-axis;
- $R_z(90^\circ + \lambda)$ is the rotation matrix with respect to the vehicle's longitude by the angle $(90^\circ + \lambda)$ of a right-handed co-ordinate system around the z-axis; and
- ΔX is the position vector between an unknown point and the LLS origin ($X_{ECEF} - X_o$).

Equation (3.4) expands to

$$\begin{bmatrix} x_i^{LLS} \\ y_i^{LLS} \\ z_i^{LLS} \end{bmatrix} = \begin{bmatrix} -\sin \lambda & \cos \lambda & 0 \\ -\sin \varphi \cos \lambda & -\sin \varphi \sin \lambda & \cos \varphi \\ \cos \varphi \cos \lambda & \cos \varphi \sin \lambda & \sin \varphi \end{bmatrix} \begin{bmatrix} x_{ECEF} - x_o \\ y_{ECEF} - y_o \\ z_{ECEF} - z_o \end{bmatrix} \quad (3.5)$$

3.2.2 Rotations

In equation 3.1 the BFS and LLS are related by a composite transformation matrix, $R_{\psi\theta\gamma}$. It has been stated above that this matrix contains nine direction cosines, composed of the three attitude components. The derivation of this matrix is now described showing how the attitude components are explicitly related to the co-ordinate systems.

The three parameters that are commonly used to define attitude components are the well-known Euler angles (ψ, θ, γ) for pitch, roll and yaw. These are shown in figure 3.2. Equations (3.6), (3.7) and (3.8) give the explicit form of the primitive rotation matrix for right handed rotations about the x, y and z axes, respectively.

$$R_{\psi} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \psi & \sin \psi \\ 0 & -\sin \psi & \cos \psi \end{bmatrix} \quad (3.6)$$

$$R_{\theta} = \begin{bmatrix} \cos \theta & 0 & -\sin \theta \\ 0 & 1 & 0 \\ \sin \theta & 0 & \cos \theta \end{bmatrix} \quad (3.7)$$

$$R_{\gamma} = \begin{bmatrix} \cos \gamma & \sin \gamma & 0 \\ -\sin \gamma & \cos \gamma & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (3.8)$$

For this research the rotation order has been defined as follows:

- Firstly a rotation, R_{γ} , of the yaw angle γ about the z-axis (equation 3.8)
- Secondly a rotation, R_{θ} , of the roll angle θ about the y-axis (equation 3.7)
- Finally a rotation, R_{ψ} , of the pitch angle ψ about the x-axis (equation 3.6).

The product of the three rotation matrices produces $R_{\psi\theta\gamma}$,

$$R_{\psi\theta\gamma} = \begin{bmatrix} \cos\theta\cos\gamma & \cos\theta\sin\gamma & -\sin\theta \\ \sin\psi\sin\theta\cos\gamma - \cos\psi\sin\gamma & \sin\psi\sin\theta\sin\gamma + \cos\psi\cos\gamma & \sin\psi\cos\theta \\ \cos\psi\sin\theta\cos\gamma + \sin\psi\sin\gamma & \cos\psi\sin\theta\sin\gamma - \sin\psi\cos\gamma & \cos\psi\cos\theta \end{bmatrix} \quad (3.9)$$

As the BFS co-ordinates of the antennas are known from an initial survey, by measuring antenna positions at an epoch, and transforming from ECEF to LLS co-ordinates, as described above, the rotation matrix $R_{\psi\theta\gamma}$ can be formed (equation 3.1). Attitude components can then be extracted using the elements as follows;

$$Pitch, \psi = \arctan\left[\frac{2,3}{3,3}\right] = \arctan\left[\frac{\sin\psi\cos\theta}{\cos\psi\cos\theta}\right] \quad (3.10)$$

$$Roll, \theta = -\arcsin[1,3] = -\arcsin[-\sin \theta] \quad (3.11)$$

$$Yaw, \phi = \arctan\left[\frac{1,2}{1,1}\right] = \arctan\left[\frac{\cos \theta \sin \phi}{\cos \theta \cos \phi}\right] \quad (3.12)$$

The 3-2-1 sequence of rotations adopted in this research allows a simple conversion between the heading and yaw angles because both are in the same plane. Therefore

$$\text{heading} = \gamma \text{ (for positive yaw)} \quad (3.13)$$

or

$$\text{heading} = 2\pi - \gamma \text{ (for negative yaw)} \quad (3.14)$$

The sign convention is defined so that a positive pitch angle corresponds to an upward movement of the aircraft nose, a positive roll angle corresponds to left wing up, and the yaw angle increases counter-clockwise.

These are the basic principles of GPS attitude determination. The antenna locations are determined in two co-ordinate systems (BFS and LLS) at an epoch, and the rotations between these two systems represent the combined affect of the three attitude components. These mathematical principles have been applied to derive attitude from GPS measurements using a range of instruments and processing methods.

3.3 Dedicated GPS Attitude Systems

Since the early nineties a range of dedicated GPS attitude determination units have been commercially available. These include the Trimble TANS Vector, Ashtech's 3DF and later the ADU2, the Sercel NR230 and the LABEN/Loral Tensor. These receivers integrate four (or three in the case of the Sercel NR230) antennas into one self-contained unit and operate all the tracking channels from a single receiver oscillator (Han *et al.*, 1997). An array of fixed antennas, separated by known baselines, directly measure carrier phase difference (figure 3.4). It is an application of the same interferometric principle that has been used by geodesists

working in VLBI for many years. GPS single differencing applies the same principle, but in the computational stage rather than at the measurement stage. Processing the phase differences leads to the determination of the components of the baseline vectors between the pairs of antennas and division by length yields directly the direction cosines (the orientation) of the lines joining them – and therefore the attitude of the object on which the antennas are mounted (Corbett and Short, 1994).

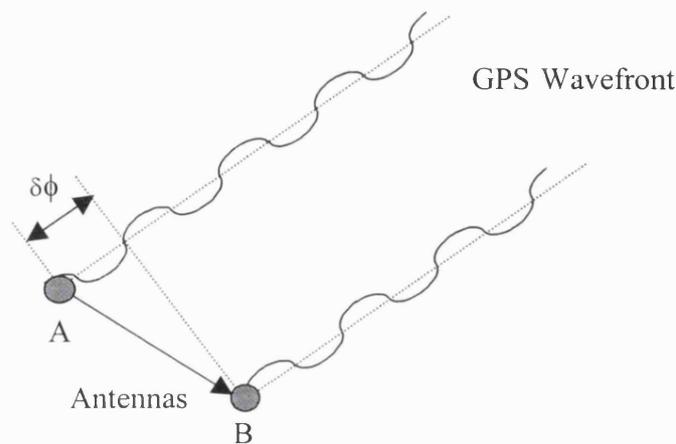


Figure 3.4: GPS phase difference measurement (Interferometry)

Several algorithms have been introduced in the past for attitude determination using GPS differential phase measurements. They are based on different approaches, one of which uses a least squares fit of the attitude matrix to the basic GPS phase measurements. The cost function of the fit is expressed in the literature in the most straight-forward formulation as a function of the attitude matrix (Nadler and Bar-Itzhack, 1998). Campana *et al.* (1999) describe the GPS least squares estimation used in the LABEN/Loral Tensor receiver where the three attitude components are computed from raw differential phase measurements. In this case, the measured relative range Δr , to a GPS satellite between a pair of antennas is expressed as:

$$\Delta r_{ij} = S_i^T A^T b_j + v_{ij} + m_{ij}$$

where:

GPS channels, $i = 1, 2, \dots, n$;

Antenna baselines, $j = 1, 2, 3$;

S_i is the line of sight direction from the GPS satellite;

b_j is the baseline distance in the body frame from the reference antenna to the j^{th} antenna;

v_{ij} is the receiver phase measurement noise;

m_{ij} is the phase multipath error; and,

A is the angular transformation matrix from the chosen reference frame to the body frame system.

Given an initial attitude matrix, a solution estimate can be obtained by solving for a correction matrix δA of small angular rotations. Attitude determination is achieved by improving estimates of the transformation matrix from successive measurement of the carrier phase (and hence Δr). This solution assumes that integer ambiguities are solved as an initial step and that integers are kept updated by lower level processing in the receiver. As attitude calculations are generated, these integer values are continuously updated and checked.

Such devices have been widely used on aircraft, ships and land vehicles, and many assessments of their performance have been reported. The accuracy of these dedicated systems is in the order of 5 to 15 arc-minutes for the heading component and 20 to 30 arc-minutes for the pitch and roll components, using a 1m antenna separation. In theory the accuracy increases in proportion to the antenna separation (i.e. a 10m baseline provide ten times the accuracy of a 1m baseline) although in practice longer baselines can introduce flexure errors. After testing an Ashtech 3DF, Toth and Grejner-Brzezinska (1998) concluded that the accuracy of GPS-derived attitude components is highly competitive compared to their counterparts determined from the stand-alone inertial navigation system.

3.4 Attitude from Independent GPS Receivers

3.4.1 Equipment

An alternative to using a dedicated system is to use a number of antennas, each connected to an independent receiver. The accuracy of these non-dedicated systems is comparable to that of the dedicated systems (Mahmud, 1999). This approach offers greater flexibility as each receiver can be used for other positioning applications when it is not required as part of an attitude system. It is also more convenient to connect each antenna to a local receiver in cases where antennas are widely distributed, i.e. on a ship. In this situation it would be possible to connect each antenna to one common receiver but this would require significant lengths of cabling and the signal from the antenna would almost certainly need some additional amplification before being processed in the receiver.

The use of independent receivers has further advantages in the field of airborne direct georeferencing. Firstly, the use of a dual-frequency receiver allows positions to be determined to a high accuracy, using a fixed ambiguity carrier-phase solution, in conditions when such a solution may not be possible with a single frequency dedicated attitude system. Secondly, using independent receivers, a system with a recording rate of 10 or 20Hz is possible, whereas the highest available rate of a dedicated system is 2Hz (using the Ashtech ADU2). A high recording rate is important as the high frequency motions of the airborne sensor need to be estimated. In an integrated GPS/INS solution the higher recording interval allows a better constraint of inertial angular drifts. As part of the current research project an assessment has been made of the attitude data-rate required for a direct georeferencing application. The principal aim is to establish whether a high rate gyro is needed to model high frequency motions, or whether a 10Hz GPS solution would be adequate. Results from this analysis are presented in section 6.5.1.

Independent receivers have been used to measure attitude for a range of applications, a few examples are;

- The FLI-MAP system which uses four independent receivers on an helicopter to directly georeference a laser scanner (Liu *et al.*, 1997).
- Corbett and Short used independent receivers to determine the attitude of an aerial camera in a number of test flights (Corbett, 1993; Corbett and Short, 1994).
- Han *et al.* (1997) used the same approach to determine the attitude of a car.

Schleppe (1997) describes a trend towards the use of low-cost GPS attitude systems. With the recent and rapid development of low-cost OEM GPS receiver cards capable of accurate phase measurement, coupled with the availability of low-cost processors, the cost of GPS attitude hardware has fallen dramatically. Using a 40cm antenna array, four OEM receivers and a laptop computer the reported accuracies (RMS) were 13 arc-minutes in heading, 1.5 degrees in pitch and 52 arc-minutes in roll.

A number of *partial* GPS attitude systems have also been developed which use two antennas to obtain two attitude components (see for example, Ford *et al.*, 1996; Vinnins and Gallop, 1997; Harvey, 1998; Keong and Lachapelle, 1998; and Vinnins and Gallop, 1998). The antennas are usually arranged along the ‘direction of travel’ axis of the subject vehicle, e.g. the fuselage axis of an aircraft, so that they record heading and pitch measurements. This arrangement can be particularly productive in an integrated GPS/INS configuration (as described in chapter 2) using low-cost gyros for attitude determination. The heading information from GPS allows the gyro’s azimuth to be determined and it can then produce true heading measurements. In general, the gyro measurements alone will provide adequate pitch and roll information, although having established two antennas to measure heading it doesn’t require much additional effort to bring GPS pitch data into the integrated system.

Kornfeld *et al.* (1998) have even developed a system for aircraft attitude determination based on a single GPS antenna. In this approach GPS velocity measurements are used to derive a flight path angle in the local horizontal plane, and a roll angle about the aircraft velocity vector. These attitude variables have been termed *pseudo-attitude* to distinguish this approach from traditional roll and pitch attitude. Flight tests have shown that for co-ordinated flight (level flight on a constant heading, as is the case during flight lines of an aerial survey) that pseudo-roll corresponds closely to the traditional roll angle.

3.4.2 Processing Methods for Independent Receivers

There are two general approaches to attitude determination from independent receivers. The first is a *baseline-by-baseline* approach in which the vectors between antennas are determined explicitly and are then used to determine attitude. In the second approach, attitude components are determined directly from GPS observations, hence this is termed the *direct* approach. Within these two general approaches many specific processing methods have been developed. The following description is intended to serve as an overview of the general approaches only, it does not cover the range of possible algorithms in any detail. In both methods ambiguity resolution is a key issue and will govern the success of the attitude determination.

3.4.2.1 Baseline-by-Baseline Attitude Determination

The baseline-by-baseline approach borrows most of its algorithms from kinematic GPS positioning. Holding one antenna position fixed, the vectors to any additional antennas in the array are computed. Once ambiguities are fixed, baseline determinate accuracy is only dependent on the carrier phase noise. This is a function of the receiver type and the multipath environment, typically accuracies of a few centimetres are achievable. Using *a priori* knowledge of the antenna positions in a fixed body reference frame, the orientation of the platform can then be computed. In the case where redundant baselines exist, the attitude parameters can be estimated using a least squares process. There are two distinct

parts then to this approach, the determination of baseline vectors from GPS observations, and the determination of attitude from the baseline vectors. These two parts will now be dealt with in turn.

i) Baseline Determination

Baseline determination can be achieved by any of the available standard positioning methods using GPS carrier phase observations. The main task in this process is ambiguity resolution which can be achieved using a multiple or single epoch method. A single epoch approach effectively removes the continuity problems associated with cycle slips and loss of lock, but at the cost of increasing the computational load. Most of the single-epoch ambiguity techniques employed in practice are adapted to make use of additional information such as known baseline lengths. Additional information of this kind can make the algorithms more efficient, something that becomes increasingly important in real-time applications.

Liu *et al.* (1997) describe the baseline determination method used by the FLI-MAP system, it is based on a method developed by Remondi (1991). Possible candidates for double differenced carrier phase ambiguities are formed which can meet the known baseline constraint within a searching volume centred at the DGPS position. Ambiguity candidates are propagated over time using triple differencing techniques along the trajectory formed, and statistics are generated in terms of the L1 norm (the mean of the absolute computed deviations). The successful candidate can be chosen if it gives the best statistics and can be distinguished from the other candidates. In the FLI-MAP system, since baselines are short and known *a priori*, and a vertical gyro provides approximate roll and pitch measurements, the number of possible ambiguity candidates is very small.

Han *et al.* (1997) use an instantaneous ambiguity resolution technique based on the LAMBDA² method. Carrier-phase and pseudo-range observations are used with fixed baseline length constraints, and gyro data. Euler and Hill (1995) also

² LAMBDA is the Least squares AMBiguity Decorrelation Adjustment developed by P.J.G. Teunissen and first described in Teunissen (1993).

discuss how exploiting redundant GPS information and *a priori* knowledge of baseline components can aid ambiguity resolution.

ii) Attitude Determination

Having determined the baselines, this information is then used to determine attitude with reference to the *a priori* antenna positions in the body frame system. To determine initial values of the attitude parameters, a method which uses only two vectors can be used. More precisely it uses two direction parameters from the first baseline and one direction parameter from the second baseline. The bias from these three direction parameters will directly affect the accuracy of the attitude parameters. Multipath effects will also directly affect the attitude solution. If there are more than two baselines then a least squares estimation can be carried out using the initial attitude values from the two vector method as starting approximations.

The least squares method will use all the available baseline information to derive the three attitude parameters. In addition, the least squares method can easily implement the constraints from the constant baseline lengths and data from other sources such as gyros, at this stage. This will reduce the error in the baseline vectors, especially due to multipath effects. Although the multipath effect on a baseline may be a systematic bias between neighbouring epochs (depending on dynamics and recording rate) it is a random feature for independent baselines. Hence the least squares estimation will have the ability to reduce the multipath effect.

3.4.2.2 Direct Attitude Determination

An alternative approach is to directly estimate the attitude parameters from the GPS observations. Axelrad and Ward (1994) presented a technique using carrier phase single difference observations to directly estimate attitude. Cross *et al.* (1997) developed modified double difference observation equations that also allow attitude components to be determined directly from GPS observations. Both these techniques incorporate the *a priori* knowledge of the antenna body frame

co-ordinates into the measurement models. Theoretically, using the *direct* approach to attitude estimation, the attitude can be determined using three antennas and only three single or double differences, one of which must be on a second non-collinear remote antenna. This contrasts to the baseline-by-baseline approach in which at least six single or double difference measurements (i.e. 3 observations for each of 2 non-collinear baselines) must be used to determine three dimensional attitude (Schleppe, 1997).

Obviously the reduced number of unknown parameters not only allows attitude determination when fewer satellites are available, but it also the means that given the same number of antennas and satellites, the direct solution will have greater redundancy. Increased redundancy improves the systems ability to detect outliers amongst observations, i.e. its reliability. In the case of the double difference approach, once ambiguities have been fixed (either correctly or incorrectly) a single baseline only provides a redundancy of $f(p-1)-3$, where p is the number of satellites in view, and f is the number of frequencies used. Making use of all the available observations, but only solving for the three attitude components leads to a redundancy of $n-1[f(p-1)] - 3$, with n antennas. For example, with four antennas, and five satellites in view, there would be a redundancy of twenty-one. Section 3.6 describes how direct attitude determination can help to mitigate the major error source in GPS attitude systems - multipath. The equations needed to relate the GPS double-differenced phase observables directly to the three attitude parameters are derived in section 3.5.

3.5 GPS Observation Equations for Attitude

The direct attitude determination method based on the double difference carrier phase observable proposed by Cross *et al.* (1997) and Mahmud (1999) is used in this research. In order to implement this approach, the standard double difference phase observation equation has to be modified to include the attitude parameters and integer ambiguities. The following derivation of the modified double difference phase observation equation is taken from Mahmud (1999).

Firstly the composite rotation matrix $R_{\psi\theta\gamma}$ (3.9) is re-written as

$$R_{\psi\theta\gamma} = \begin{bmatrix} m_{11} & m_{12} & m_{13} \\ m_{21} & m_{22} & m_{23} \\ m_{31} & m_{32} & m_{33} \end{bmatrix} \quad (3.15)$$

where, m_{11} to m_{33} are the nine direction cosine elements. Substituting equations (3.2), (3.3) and (3.15) in (3.1) results in

$$x_A^{LLS} = m_{11}x_A^{BFS} + m_{12}y_A^{BFS} + m_{13}z_A^{BFS} \quad (3.16)$$

$$y_A^{LLS} = m_{21}x_A^{BFS} + m_{22}y_A^{BFS} + m_{23}z_A^{BFS} \quad (3.17)$$

$$z_A^{LLS} = m_{31}x_A^{BFS} + m_{32}y_A^{BFS} + m_{33}z_A^{BFS} \quad (3.18)$$

$$x_B^{LLS} = m_{11}x_B^{BFS} + m_{12}y_B^{BFS} + m_{13}z_B^{BFS} \quad (3.19)$$

$$y_B^{LLS} = m_{21}x_B^{BFS} + m_{22}y_B^{BFS} + m_{23}z_B^{BFS} \quad (3.20)$$

$$z_B^{LLS} = m_{31}x_B^{BFS} + m_{32}y_B^{BFS} + m_{33}z_B^{BFS} \quad (3.21)$$

⋮

$$x_N^{LLS} = m_{11}x_N^{BFS} + m_{12}y_N^{BFS} + m_{13}z_N^N \quad (3.22)$$

$$y_N^{LLS} = m_{21}x_N^{BFS} + m_{22}y_N^{BFS} + m_{23}z_N^{BFS} \quad (3.23)$$

$$z_N^{LLS} = m_{31}x_N^{BFS} + m_{32}y_N^{BFS} + m_{33}z_N^{BFS} \quad (3.24)$$

The double difference phase observation equation for each pair of satellites (i and j) and receivers (A and B) is given in Hofmann-Wellenhof *et al.*(1994) as

$$\Phi_{AB}^{ij} = \frac{1}{\lambda} \rho_{AB}^{ij} + N_{AB}^{ij} \quad (3.25)$$

where

- Φ_{AB}^{ij} is the double differenced phase observable;
- ρ_{AB}^{ij} is the double differenced geometric distance;
- N_{AB}^{ij} is the double differenced initial integer ambiguity; and
- λ is the wavelength.

Multiplying (3.25) by λ leads to

$$\lambda \Phi_{AB}^{\ddot{ij}} = \rho_{AB}^{\ddot{ij}} + \lambda N_{AB}^{\ddot{ij}} \quad (3.26)$$

which expands to

$$\rho_A^i - \rho_B^i - \rho_A^j + \rho_B^j + \lambda N_{AB}^{\ddot{ij}} = \lambda (\Phi_A^i - \Phi_B^i - \Phi_A^j + \Phi_B^j) \quad (3.27)$$

Now we write

$$\rho_A^i = \left[(x_A^i)^2 + (y_A^i)^2 + (z_A^i)^2 \right]^{1/2} \quad \text{etc.} \quad (3.28)$$

where

$$x_A^i = x_i^{\text{LLS}} - x_A^{\text{LLS}} = x_i^{\text{LLS}} - (m_{11} x_A^{\text{BFS}} + m_{12} y_A^{\text{BFS}} + m_{13} z_A^{\text{BFS}}) \quad (3.29)$$

which leads to the complete double differenced phase observation equation in terms of attitude parameters and integer ambiguities :

$$\begin{aligned} & \left\{ \left[x_i^{\text{LLS}} - (m_{11} x_A^{\text{BFS}} + m_{12} y_A^{\text{BFS}} + m_{13} z_A^{\text{BFS}}) \right]^2 + \left[y_i^{\text{LLS}} - (m_{21} x_A^{\text{BFS}} + m_{22} y_A^{\text{BFS}} + m_{23} z_A^{\text{BFS}}) \right]^2 + \left[z_i^{\text{LLS}} - (m_{31} x_A^{\text{BFS}} + m_{32} y_A^{\text{BFS}} + m_{33} z_A^{\text{BFS}}) \right]^2 \right\}^{1/2} - \\ & \left\{ \left[x_i^{\text{LLS}} - (m_{11} x_B^{\text{BFS}} + m_{12} y_B^{\text{BFS}} + m_{13} z_B^{\text{BFS}}) \right]^2 + \left[y_i^{\text{LLS}} - (m_{21} x_B^{\text{BFS}} + m_{22} y_B^{\text{BFS}} + m_{23} z_B^{\text{BFS}}) \right]^2 + \left[z_i^{\text{LLS}} - (m_{31} x_B^{\text{BFS}} + m_{32} y_B^{\text{BFS}} + m_{33} z_B^{\text{BFS}}) \right]^2 \right\}^{1/2} - \\ & \left\{ \left[x_j^{\text{LLS}} - (m_{11} x_A^{\text{BFS}} + m_{12} y_A^{\text{BFS}} + m_{13} z_A^{\text{BFS}}) \right]^2 + \left[y_j^{\text{LLS}} - (m_{21} x_A^{\text{BFS}} + m_{22} y_A^{\text{BFS}} + m_{23} z_A^{\text{BFS}}) \right]^2 + \left[z_j^{\text{LLS}} - (m_{31} x_A^{\text{BFS}} + m_{32} y_A^{\text{BFS}} + m_{33} z_A^{\text{BFS}}) \right]^2 \right\}^{1/2} + \\ & \left\{ \left[x_j^{\text{LLS}} - (m_{11} x_B^{\text{BFS}} + m_{12} y_B^{\text{BFS}} + m_{13} z_B^{\text{BFS}}) \right]^2 + \left[y_j^{\text{LLS}} - (m_{21} x_B^{\text{BFS}} + m_{22} y_B^{\text{BFS}} + m_{23} z_B^{\text{BFS}}) \right]^2 + \left[z_j^{\text{LLS}} - (m_{31} x_B^{\text{BFS}} + m_{32} y_B^{\text{BFS}} + m_{33} z_B^{\text{BFS}}) \right]^2 \right\}^{1/2} \\ & + \lambda N_{AB}^{\ddot{ij}} \\ & = \lambda (\Phi_A^i - \Phi_B^i - \Phi_A^j + \Phi_B^j) \quad (3.30) \end{aligned}$$

In practice, there will be p antennas (A,B,... p) and q satellites (i,j,... q) leading to $(p-1)(q-1)$ double difference equations. These are linearised to

$$\frac{\partial \rho_{AB}^{ij}}{\partial \psi} d\psi + \frac{\partial \rho_{AB}^{ij}}{\partial \theta} d\theta + \frac{\partial \rho_{AB}^{ij}}{\partial \gamma} d\gamma + \frac{\partial \rho_{AB}^{ij}}{\partial N} dN = [\Phi_{AB}^{ij}(\text{observed}) - \Phi_{AB}^{ij}(\text{computed})] + v \quad (3.31)$$

where

$\frac{\partial \rho_{AB}^{ij}}{\partial \psi}, \frac{\partial \rho_{AB}^{ij}}{\partial \theta}, \frac{\partial \rho_{AB}^{ij}}{\partial \gamma}, \frac{\partial \rho_{AB}^{ij}}{\partial N_{AB}^{ij}}$ are the partial derivatives computed at the provisional values;

$d\psi, d\theta, d\gamma, dN$ are the corrections to the provisional values of the attitude parameters and ambiguity term;

$\Phi_{AB}^{ij}(\text{observed})$ is the observed double differenced phase for each pair of satellites; and

$\Phi_{AB}^{ij}(\text{computed})$ is the computed double differenced phase for each pair of satellites using the provisional values;

and then assembled to the well-known matrix form

$$Ax = b + v \quad (3.32)$$

where

- A is the matrix of partial derivatives;
- x is the vector of unknown parameters;
- b is the vector of observed minus computed phase double differences;
- v is the vector of residuals.

The partial derivatives needed to construct the design matrix A are given in appendix A.

Associated with equation 3.32 will be a weight matrix, based on the correlations between the observations. Assuming all observations should be assigned equal weighting, the correlation matrix Cl for the case of three antennas and four satellites using double differenced phase observations is,

$$Cl = \begin{bmatrix} 4 & 2 & 2 & 2 & 1 & 1 \\ 2 & 4 & 2 & 1 & 2 & 1 \\ 2 & 2 & 4 & 1 & 1 & 2 \\ 2 & 1 & 1 & 4 & 2 & 2 \\ 1 & 2 & 1 & 2 & 4 & 2 \\ 1 & 1 & 2 & 2 & 2 & 4 \end{bmatrix} \quad (3.33)$$

with a corresponding *a priori* weight matrix, $W = Cl^{-1}$ (Hofmann-Wellenhof, 1994). Mahmud (1999) suggests that attitude results may be improved by introducing a satellite elevation dependent weight matrix. Section 5.5.5 describes the introduction of such a model.

3.6 Error Sources and Mitigation

Multipath is the single largest error source which limits the accuracy of GPS based attitude determination systems (Ashtech, 1996). Multipath can prevent ambiguities from being fixed, and can lead to incorrect ambiguity resolution. It can also degrade the accuracy of the final solution even when the correct ambiguities have been fixed. The period and amplitude of multipath oscillations vary significantly depending on a vehicle's dynamics and its environment. In static conditions the multipath may exhibit a predictable pattern, allowing it to be averaged out over time. A dynamic vehicle will experience less predictable multipath errors that become more noise-like.

Unlike other GPS errors (e.g. satellite and receiver clock errors) multipath cannot be eliminated or reduced by double-differencing, in fact the combination of observations increases the error. A variety of techniques have been introduced to mitigate the affects of multipath in order to get closer to realising the potential

positioning accuracies from GPS. Generally, these techniques can be classified as hardware and software techniques. In terms of hardware, special antenna designs, like choke ring antennas and antennas with ground planes, try to prevent multipath signals entering a receiver. Developments in receiver design, like narrow correlation architecture, attempt to reduce the impact of any multipath errors that are present. Mitigation techniques using software include parametric approaches and digital filtering techniques that deal with the observations during the data processing stage. Ray (1997) and Jai *et al.* (1999) both provide a review of some of these techniques.

It is not possible to take advantage of improved antenna design (choke-ring antennas and ground planes) in an airborne environment. Antennas installed on aircraft must be small and aerodynamic so that they do not cause unwanted drag. Nor is it possible simply to site antennas in favourable multipathing locations, more practical considerations such as the ease of installation, cabling and aerodynamics must take precedence. None of the mitigation techniques implemented in the receiver architecture can effectively reduce low frequency multipath caused by nearby (< 30m) reflecting surfaces (Ray, 1999). This leaves only software based mitigation techniques, the majority of which can only be applied to static sites, where predictable patterns of reflectance can be modelled over time.

One software based approach to reducing problems due to multipath is to identify particular satellite-to-antenna combinations where the error is large and remove them from the solution. Cross *et al.* (1997) found that by removing satellites on a trial and error basis it was possible to identify the source of biased data resulting from multipath affects. By excluding particular satellites from the solution the success rate of single epoch ambiguity resolution improved significantly, however, removing satellites on a purely trial and error basis is obviously not an efficient way to identify *problem satellites*, particularly in dynamic applications where the satellite-receiver geometry is changing more rapidly. Therefore an automated approach to detect noisy data was introduced. Outlier detection of this kind requires significant redundancy, something that does not exist in standard

GPS baseline determination. The method of direct GPS attitude determination proposed by Cross *et al.* (1997) and used in this research, leads to far greater redundancy and therefore has the potential to improve the system's reliability by allowing automated outlier detection.

Mahmud (1999) describes how this automated method has been implemented in a software package known as the GPS Routine for Attitude Parameter Estimation (GRAPE). Results from static tests have shown that using the direct attitude approach, the automatic outlier detection process is consistently able to identify and reject biased data that would go undetected in a baseline-by-baseline approach. The direct attitude system used in this research therefore has the potential to mitigate the effects of multipath.

A second source of errors in GPS attitude determination is caused by flexure in the body frame. In the mathematics presented, it is assumed that the antenna array is a rigid body and that baseline vectors measured *a priori* will be accurate at any instant of time. This condition cannot always be met in practice, particularly when antennas are installed on an aircraft. To allow the greatest possible antenna separation, two of the antennas would need to be located on the wing-tips, which are subject to significant vibrations (or flutter) during flight.

A series of test flights were carried out in 1993 and 1994 in which antennas were installed on the Natural Environment Research Council's Piper Chieftain survey aircraft (Corbett 1994; Corbett and Short 1994 & 1995). Results of these tests suggested that the wing flexure was around 150mm and that as this could not be adequately modelled it was a limiting factor on the roll accuracy. As a result the two wing-tip antennas were relocated on small and stable *winglets* attached to the fuselage. The increased stability however was achieved at the expense of antenna separation, the original 12m baseline was reduced to 1.2m. This particular aircraft was used for the current research and is described in more detail in the following chapter.

Models have been developed to try and account for wing flutter by assuming some degree of regular motion, but they do not completely remove the problem. One of the obstacles to more comprehensive modelling is that the recording rate of the antennas is too low to detect the higher frequency flutter characteristics. It is quite common in GPS based direct georeferencing to include wing flexure parameters as states in a Kalman filter and then to estimate them along with the other error states (see for example Creamer *et al.*, 1998; and Cohen *et al.*, 1994).

A more specific discussion of the error sources associated with the equipment and processing methods used in this research is presented in chapters five and six.

3.7 Concluding Remarks

This chapter has given an overview of GPS based attitude determination from the underlying principles to specific applications. It was shown that essentially attitude describes a set of three rotation angles that relate a fixed body to a reference co-ordinate system. Systems in which these principles have been put into practice were then described. The most commonly used systems are dedicated single receivers, which make differential carrier-phase measurements between a fixed array of antennas. It is also possible to determine attitude using an array of antennas each connected to its own independent receiver.

Attitude has generally been derived from these independent receiver systems in two steps. Firstly, baselines are measured in a reference system using standard GPS techniques at an epoch. Then, using *a priori* knowledge of the antenna positions in a fixed body reference frame, the orientation of the vehicle is computed. An alternative approach, which directly estimates attitude parameters from raw GPS observables was then introduced. The modified double difference approach presented by Cross *et al.* (1997), and adopted in this research, has been described and its advantages in terms of improved reliability were discussed. A detailed description of how this method has been employed in the current airborne application is given in sections 5.4 and 5.5.

Chapter 4 System Design and Testing

4.1 Introduction

This chapter describes the design and testing of the direct georeferencing system used in this research project. It has been designed to reflect a number of the trends in direct georeferencing and GPS attitude determination described in chapters two and three. The system has been tested during two trial flights on an aircraft equipped with appropriate remote sensing instruments. This chapter focuses on the navigation and sensing equipment used, and the design of the flight tests. Details of the data processing algorithms and results from the system are presented in chapters five and six respectively.

In chapter two a number of trends were identified in the field of direct georeferencing, these included:

- the widespread integration of GPS and INS
- the design of orientation systems to meet the specific requirements of sensors and applications
- an increasing interest in low-cost systems, particularly those based on GPS attitude determination
- the use of low-cost gyros to aid GPS attitude systems

In line with these trends, a system has been developed for this research that is based primarily on GPS for both position and attitude determination, but which can also incorporate attitude measurements from a low-cost gyro. The GPS attitude determination is based on four independent receivers each attached to an antenna. GPS observables are processed using the *direct* attitude determination method outlined in section 3.5, the advantages of which were described in sections 3.4 and 3.6. The predicted accuracy of this system makes it appropriate for georeferencing airborne line scanners used for resource mapping and other environmental applications (see section 2.3.2). To test this system in an airborne operational environment two flight trials were undertaken using the Natural Environment Research Council (NERC) Piper Chieftain survey aircraft.

4.2 NERC Airborne Remote Sensing Facility

4.2.1 Overview

The proposed direct georeferencing system was installed and flown by the NERC Airborne Remote Sensing Facility (NERC ARSF), its support on this project is gratefully acknowledged. Much of the following information relating to specific instruments, has been extracted from the 'user guide handbook' (Wilson, 1995).

The mission statement of the NERC ARSF is:-

To provide airborne remotely sensed data to environmental scientists in support of their research and survey/monitoring and to promote awareness and application of airborne remote sensing techniques by dissemination to the wider scientific community. By carrying out its mission it is expected that the NERC Airborne Remote Sensing Facility will contribute to an improvement in the United Kingdom's industrial competitiveness and quality of life.

To achieve its mission the ARS Facility maintains a Piper PA31 350 Chieftain aircraft which has been converted for airborne survey and research operations. Figures 4.1 and 4.2 show a photograph and cross-sectional schematic of this aircraft, respectively. A survey to position the sensors and GPS antennas in a fixed body frame system has been conducted using conventional techniques.

The three imaging instruments used are:-

- a Wild RC-10 survey camera;
- a Daedalus 1268 Airborne Thematic Mapper (ATM); and,
- a Compact Airborne Spectrographic Imager (CASI).

The characteristics of these systems and their spatial resolutions are discussed in sections 4.2.2 and 4.2.3.



Figure 4.1: NERC Survey Aircraft

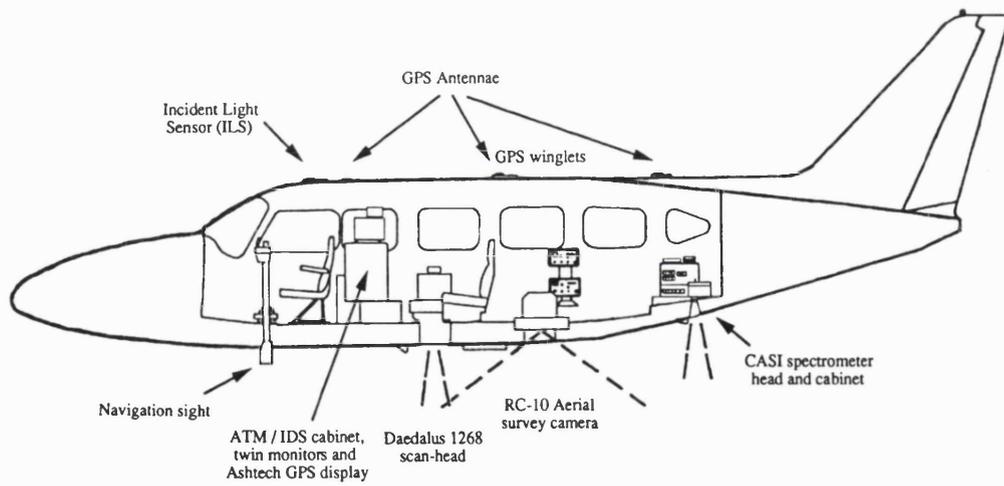


Figure 4.2: NERC Survey Aircraft: Cross-sectional schematic (from Wilson, 1997)

The NERC ARSF is currently implementing a programme which aims to supply users with geometrically and radiometrically corrected data products in a convenient mapping system. A key part of the so-called Integrated Data System (IDS) is the provision of on-board navigation instruments to determine the exterior orientation parameters of the sensors (i.e. they are developing a direct georeferencing system). The principal instrument that is normally used for determining both position and attitude information is an Ashtech ADU2 dedicated GPS attitude system of the type described in section 3.3. The receiver, which is located behind the co-pilot/navigator's seat, is connected to four, low drag, military specification, dual frequency GPS antennas, made by Sensor Systems. Two antennas are installed on top of the fuselage, one above the cockpit and the second in front of the tail plane. The remaining two are located on winglets, small supports built onto the side of the fuselage above the wings (figure 4.3).

In the Summer of 1998 a Litef LCR 92-H Attitude and Heading Reference System (AHRS) was purchased and installed. At present it is rigidly fixed to the top of the ATM scan-head to ensure it records the same motions as this sensor (figure 4.4). This is a relatively inexpensive and rugged fibre optic gyroscope device, or FOG, as described in section 2.3.3.2. It was added to the ADU2 as it was considered that higher frequency attitude output (the gyro has a maximum rate of 64Hz in contrast to the ADU2's 2Hz) was required to adequately georeference each scan line from the CASI or ATM.

To test the independent receiver approach outlined previously, four high data rate dual-frequency GPS receivers were connected to the four existing antennas during two test flights conducted in October 1998. Signals from each antenna were split to allow the ADU2 to operate in parallel to the test configuration thus allowing it to act as a comparative solution. Further details of these flights and the specific position and attitude system being tested are given in section 4.3.



Figure 4.3: Winglet and antenna (above)



Figure 4.4: ATM with AHRS Attached

4.2.2 Sensor Characteristics

In this section the three sensors installed on the survey aircraft, and used during the test flights for this project are described. The two direct georeferencing solutions (the IDS programme implemented by NERC and the specific system tested in this project) have been designed to meet the needs of the two scanners, and are not sufficiently accurate to meet the exterior orientation requirements of aerial photography. In addition the camera has no means of digital input or output, to reference instants of exposure to navigation data. The characteristics of the camera will therefore not be described in any detail here, but its usefulness in the calibration process will be discussed later.

4.2.2.1 RC-10 Aerial Camera

The survey camera occupies the rear camera port and can be automatically controlled using a laptop control box and a Zeiss navigation site, mounted beside the co-pilot (navigator). It can take black and white, colour or false-colour infra-red photography. The standard mount has been designed for near-vertical photography but it can be adjusted for oblique photography.

4.2.2.2 ATM Whiskbroom Line Scanner

The Daedalus 1268 scan head occupies the forward camera port with the processing cabinet immediately behind the pilot's seat. This instrument is an 11 channel multi-spectral whiskbroom¹ scanner. These channels cover the visible and near-infrared (ch 1-8), shortwave infrared (SWIR) (ch 9 & 10) and thermal infrared (TIR) (ch 11) and include channels that closely match the seven spectral channels of the Landsat Thematic Mapper. The scan mirror has three synchronised speeds (12.5, 25 and 50Hz) to optimise the scan rate to more closely match data acquisition and coverage over the ground at various altitudes. This avoids both under-sampling (coverage gaps) and too much over-sampling

¹ In a whiskbroom system, incident light is captured by a rotating mirror which then filters it into a number of paths which are imaged onto detectors.

(repetition of ground data) in the along-track direction. The actual pixel size, or ground spatial resolution, depends on the aircraft altitude since the ATM has a fixed Instantaneous Field of View (IFOV) of 2.5mrad (0.14°).

4.2.2.3 CASI Pushbroom Line Scanner

The CASI instrument is much smaller than the ATM or RC-10 and does not require a full size camera aperture. It has been mounted towards the rear of the aircraft in a custom built mounting. Although both the data acquisition and recording system are at the rear, the operator has access to the keyboard and display monitor, allowing complete control. The CASI, produced by Itres Research of Canada, is a two-dimensional CCD array based pushbroom² imaging spectrograph. One dimension of the array element is used to obtain an image frame of 512 spatial pixels of the scanned surface, which then builds into a scene as the aircraft moves forward (figure 4.5). After passing through a 15µm wide spectrographic slit, a reflection grating disperses the light from each pixel over the 400 to 915nm spectral range and records it using the 288 detectors on the orthogonal dimension of the CCD. Although it is possible to image all 512 spatial pixels and 288 spectral channels simultaneously, in practice a compromise is required to avoid excessive smearing in the along-track direction caused by the length of time taken to process data from each array element.

² A pushbroom system has fixed linear optics (i.e. no rotating mirror) and is therefore more geometrically stable.

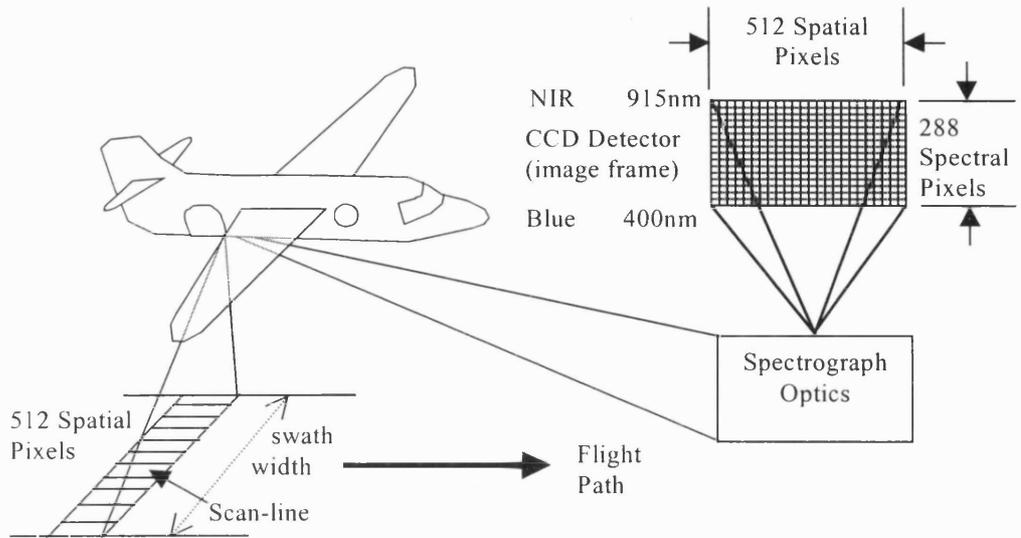


Figure 4.5: Imaging concept of the CASI CCD pushbroom spectrograph

The compromise is reached by offering two alternative operating modes, *spatial* and *spectral*. In spatial mode, data is recorded by all 512 spatial pixels, but for a limited number of programmable bands (maximum of 18) by selection and summation of the resolution detector elements (figure 4.6). A spatial bandset can be formed from single detector elements or unique summations of two or more adjacent detector elements throughout the spectral band.

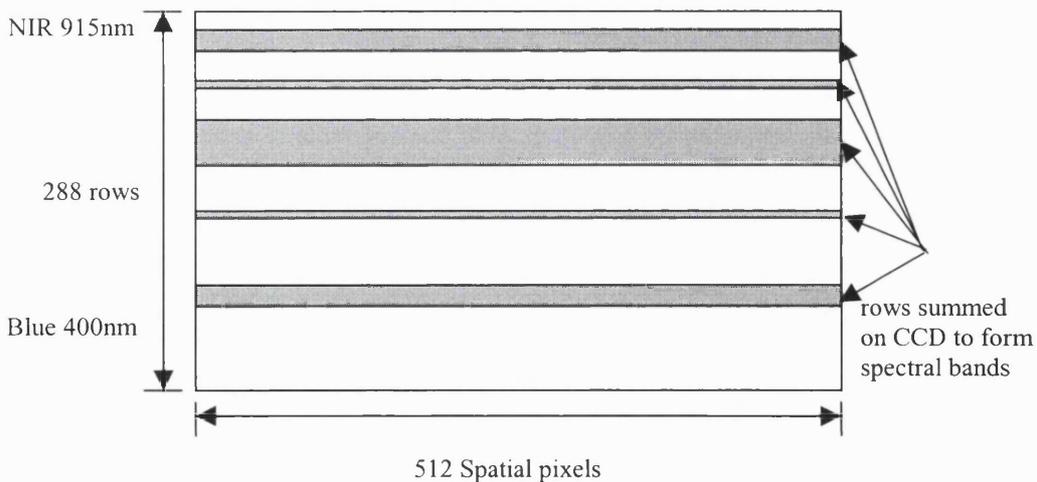


Figure 4.6: CASI spatial mode

In spectral mode the full spectral profile of 288 channels can be recorded, but for a limited number of *look directions* or pixel positions (figure 4.7). These look directions (maximum 39) must be separated by blocks of 4, 8, 12, or 16 non-imaged pixels. Minimising the spacing concentrates pixels for higher spatial resolution, whereas increasing the spacing spreads the *looks* out over the entire swath to give good spatial coverage. Due to the missing spatial pixels it sometimes becomes difficult to precisely locate the strand of spectral pixels within a scene. To offset this problem, the spectral mode produces one single full spatial band, formed using a single detector element, which can then be used to register *look directions*, this is termed the Scene Recovery Channel (SRC).

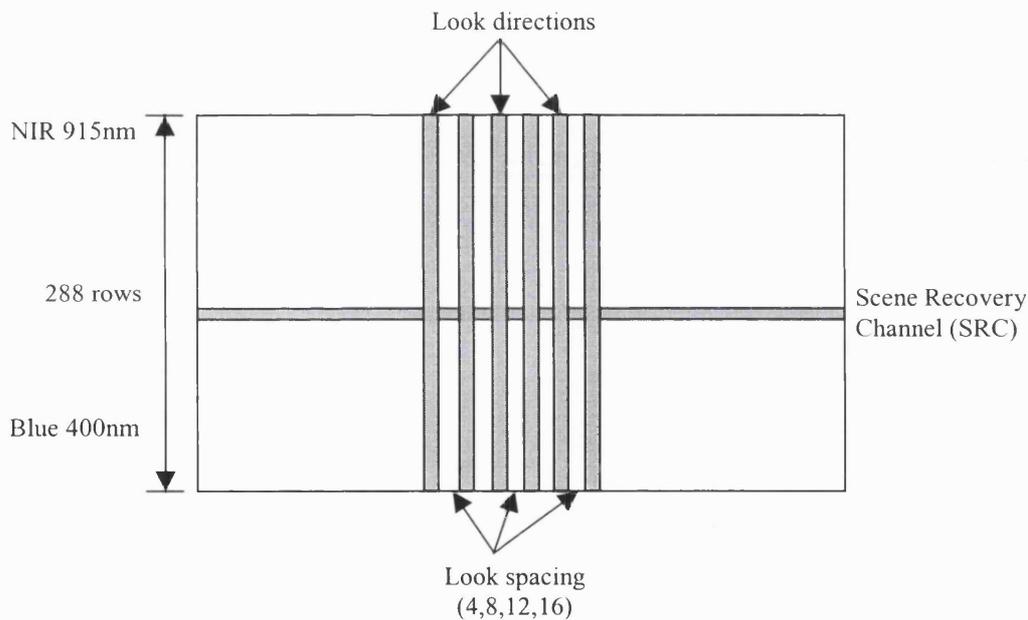


Figure 4.7: CASI spectral mode

The spectral mode has a distinct drawback in that integration times are much longer than for spatial mode. This precludes the general use of these high data rate modes at low altitudes where fast integration times are required to maintain the squareness of pixels (pixel widths being smaller at low altitude). If this is not attainable at the required altitude, an aspect ratio (pixel length to width) of 2:1 or more has to be implemented.

4.2.3 Ground Spatial Resolution

Typical ground spatial resolutions, or pixel sizes, obtainable with these multi-spectral instruments are in the range 1 to 10 metres. Spatial resolution is determined by the optical characteristics of the scanners, the spectral range and resolution required for the application, and the aircraft flying height and speed. The spatial resolution and corresponding on-the-ground accuracy requirements are of particular importance in this research as they dictate the accuracy requirement for the direct georeferencing system (see section 2.3.2). Each scanner will now be described in terms of its spatial resolution.

4.2.3.1 ATM Whiskbroom Line Scanner

The ground spatial resolution of the ATM is defined firstly by the optical characteristics of the scan head, which provides an Instantaneous Field of View (IFOV) of 2.5mrad ($\sim 0.14^\circ$), and secondly, by the aircraft altitude. Pixel length and width are both defined by the IFOV, but ground dimensions will vary as a function of the scan angle away from the nadir. Figure 4.8 illustrates the change in pixel dimension and swath width with attitude. It also shows the change in pixel dimensions at the maximum off-nadir scan angles, where square nadir pixels have become enlarged rectangles. The Y dimension is in the direction of flight, or along-track, X is the along-scan direction.

The scan geometry is further complicated by the forward motion of the aircraft during the scan mirror rotation, since adjacent pixel centres will advance as the aircraft moves forward during a scan line. However, the short time required for signal integration of an individual pixel means that there is negligible smearing of the pixel and the pixel length is dominated by the IFOV. There is a clear relationship between the pixel length determined by the aircraft altitude, and the distance over the ground that the aircraft moves between successive scan lines. For example, if the aircraft flies at a nominal speed of 120kts ($\sim 62\text{m/s}$) it will move forward approximately 1.24, 2.48 or 4.96 metres during one scan-line at the selected scan speed of 50, 25 or 12.5 scans per second. If the aircraft flies at an

altitude of 1000m, the nadir pixel length will be approximately 2.5m. At this altitude a scan-rate of 50Hz would significantly over-sample the ground, whilst a 12.5Hz rate would under-sample the image. The most appropriate scan rate needs to be chosen by the operator. The exact scan geometry used during data acquisition of a particular flight line will be determined by the geometric sensor model and the aircraft attitude, altitude and velocity. Using the ATM, all 11 spectral channels can be recorded at all times. Unlike the CASI, no specific compromise needs to be made between spectral and spatial resolution in terms of the sensor operating mode.

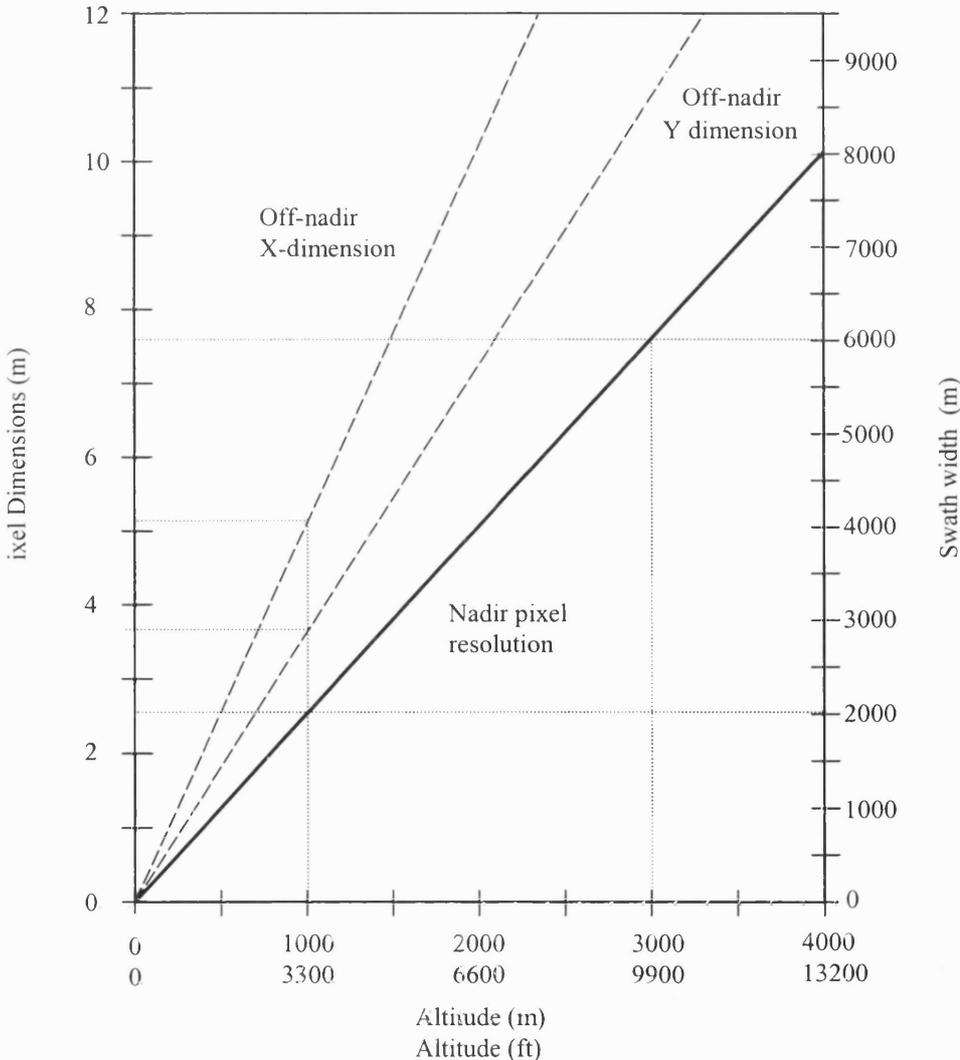


Figure 4.8 ATM pixel size and swath width (+/- 45° FOV) as a function of altitude

4.2.3.2 CASI Pushbroom Line Scanner

The CASI sensor produces a ground spatial resolution where pixel width and length are essentially independent of each other. The along track spatial resolution is a function of the integration time (the time taken for detected light to be processed and a digital state recorded), and the speed of the instrument over the ground. The integration time is itself a function of the signal level and number of channels requested by the user (spectral resolution and range). For a given light level input to the instrument, the number of channels requested by the user determines the minimum integration time and, for an average aircraft speed, the equivalent pixel length in metres. Along-track resolution can be calculated using the following formula:

$$PL = (IT * GS)/1000 \quad (4.1)$$

where,

PL = pixel length in metres,

IT = integration time in milliseconds, and,

GS = ground speed in metres/second

At 120 knots the minimum pixel length attainable is approximately 1m using only one spectral band, and increases to 4m using all 18 bands. Typically the CASI will operate with around 14 bands, leading to an integration time of 51ms, which produces a pixel length of 3.15m.

The across-track spatial resolution, or pixel width, is a function of the CCD detector element size, the focal length of the imaging lens (both of which are constant for a particular instrument), and the altitude of the aircraft. Clearly, as the altitude increases so does the pixel width and the swath width. The NERC CASI has a total field of view of 42°, produced by a 10mm lens focal length and a CCD detector size of 15µm. The across-track spatial resolution can be computed from:

$$PW = \frac{CCD \text{ Detector Width}}{Focal \text{ Length}} * Altitude \quad (4.2)$$

where,

- PW = pixel width in metres,
- CCD Detector Width = 0.015mm,
- Focal Length = 10.0mm, and,
- Altitude = metres above ground.

For example, if the aircraft flew at 3000m, the pixel width would be 4.5m. Due to the optical properties of this 'pushbroom' sensor, all pixels in a scan line are of equal width. The swath width of a spatial mode band (described in 4.2.1) will therefore be 4.5m x 512 pixels = 2304m. Figure 4.9 shows the CASI pixel width and total swath width as a function of aircraft altitude.

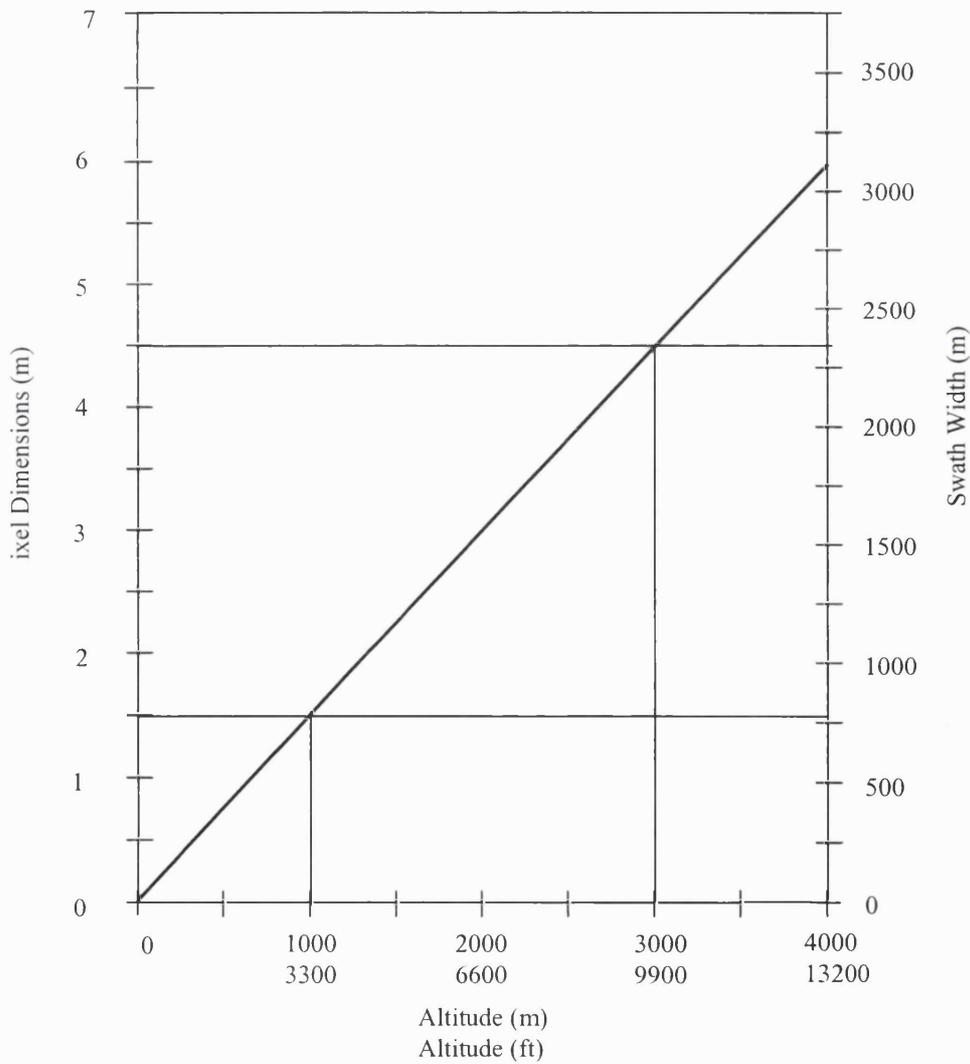


Figure 4.9: CASI pixel size and swath width (42° FOV) as a function of altitude

4.2.3.3 Spatial Resolution Summary

For both scanners there are a range of factors that will determine the spatial resolution of the imagery, and if both scanners are required to operate simultaneously an optimum solution becomes even more difficult to define. The aircraft altitude will dictate the ATM pixel size, the CASI pixel width, and the total swath width for both scanners. The aircraft speed, together with the spectral range, determines the CASI integration time, which in turn determines the pixel length. Aircraft speed also dictates the choice of ATM scan speed.

To cover a reasonably large survey area without recording unmanageably large data sets, the spatial resolution must be fairly low. In most of the environmental monitoring applications undertaken by this survey aircraft the pixel size will be in the 4 to 8m range. The on-the-ground positioning accuracy is obviously related to the image resolution, but there is no set formula that can be applied in all situations. Generally, a positioning accuracy of 10 to 20m is sufficient for the survey tasks carried out by this facility for the environmental monitoring community, although there are some applications where sub-pixel accuracy is required.

4.2.4 Direct Georeferencing Equipment

In this section, firstly the navigation equipment that is permanently located on the survey aircraft is described, and then details are given of the specific system installed for this research project.

The principal source of both position and attitude information on the NERC aircraft is an Ashtech ADU2, a dedicated GPS attitude system. The receiver is connected to four antennas arranged in a cruciform pattern, as shown in figure 4.10. The baseline between the front and rear fuselage antennas is 3.55m, and across the winglet axis the separation is 1.2m. As discussed previously, the attitude accuracy obtainable from GPS is a function of the antenna separation. It was also noted that the accuracy of the heading component is generally twice as

good as the pitch and roll. In table 4.1 the manufacturers accuracy figures, based on a 1m antenna separation, have been extrapolated to produce predicted accuracies for the configuration used on this aircraft. Clearly, the accuracy of the roll component will be the limiting factor on the overall system accuracy.

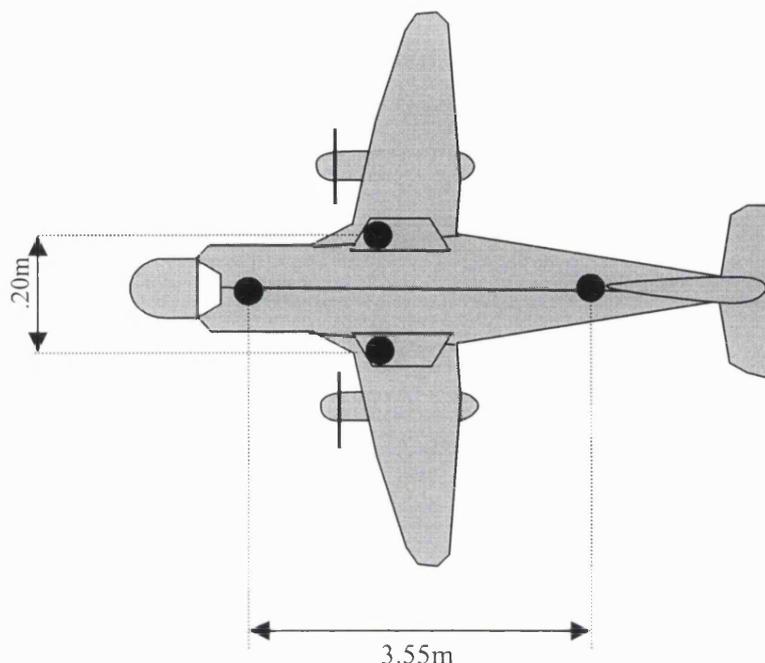


Figure 4.10: Antenna arrangement on NERC ARSF aircraft

Position	Accuracy
Autonomous	22.5m 2dRMS(95%)*
Differential	2m 2dRMS(95%)
Attitude	
Heading	3.5' (RMS)
Pitch	7' (RMS)
Roll	20' (RMS)

* corrected for S/A removal

Table 4.1: ADU2 Accuracy Specification

To obtain positioning accuracy at the 1 to 2m level a differential solution is used, either using a local temporary reference station on or near the survey site, or using the permanent reference station located at the Institute of Terrestrial Ecology (ITE), Monks Wood, Cambridgeshire. The maximum recording rate of the ADU2

is 2Hz, which is well below the data acquisition rate of the two scanners. To provide attitude data at a higher rate, and hence fill in the gaps between GPS recording epochs, a gyro system has been installed on the aircraft.

The gyro instrument used is a Litef LCR-92H Attitude and Heading Reference System, AHRS, (Litef, 1997). The system consists of three elements: the Attitude and Heading Reference Unit, AHRU, a Compass Control Unit, and a remotely mounted Magnetic Sensor Unit, MSU, (a flux valve). The AHRU is a strapdown inertial measurement unit based on three fibre optic gyro (FOG) sensors. Raw incremental angles measured by the three FOGs are accumulated and processed along with compensation software to remove known error sources. Magnetic heading from the MSU and tilt sensor inputs are used to bound azimuth, roll and pitch errors in low dynamic flight conditions. In high dynamic conditions the system acts as free inertial sensor until conditions allow the use of tilt sensor and MSU data. After applying any internal error compensation it provides angular rates, attitude and magnetic heading data. This device is typical of the class of instruments described in section 2.3.3.2, i.e. it is a small, light, rugged, low-cost unit with no moving parts. The instrument specifications are given in table 4.2.

	Accuracy (95%)
Pitch and Roll	1°
Heading	2°
Angular Rates	0.1°/s
Resolution in relative mode	0.01° (36")
Maximum Recording Rate	64Hz

Table 4.2: Litef LCR-92H AHRS technical specifications

The absolute accuracy of this unit is well below that of the ADU2 (particularly the heading component) but its resolution of 36" and recording rate of 64Hz mean it can pick up relative variations in attitude between the ADU2 epochs (two per second). This allows attitude parameters to be determined for each scan line with little or no interpolation. The AHRS can also continue to record attitude data during periods where no ADU2 solution is available due to loss of lock or cycle slips. The AHRS also has an advantage over the ADU2 in that it can be rigidly

fixed to a scanner, in this case the ATM (figure 4.4). It is therefore experiencing the same motions as the scanner at all times. The ADU2 records the mean attitude of the airframe at an epoch based on data from the four antennas. Any ADU2 attitude measurement can only be applied to the scanner imagery on the assumption that the sensor experiences exactly the same attitude changes as the airframe. In practice this is not the case, damping devices on the ATM mean that it should not experience much of the higher frequency airframe vibrations. This is a general limitation of using GPS to record sensor attitudes, it is not physically possible to mount the antennas directly to the sensor or to maintain a completely rigid connection between them.

In the standard NERC ARSF installation the ADU2 provides position and absolute attitude, and the AHRS measures relative attitude changes between ADU2 epochs. In effect, the AHRS instrument acts as an attitude interpolator. Theoretically, all the scanner, GPS and gyro data should be synchronised using GPS time so that each scan line is provided with unique position and attitude values, in practice though this integrated solution has proved difficult to achieve. The accuracy of the exterior orientation parameters is determined by the absolute accuracy of the ADU2 values and the accuracy of their transfer, in time and space, to the scan lines. The manufacturer's specifications for ADU2 accuracy given in table 4.1, have been shown to be reliable in previous tests. To quantify how successfully this accuracy can be applied to any scan line requires a full calibration of both the navigation and imaging data.

The GPS attitude determination method proposed in this research uses an independent receiver approach. To test this system, four dual-frequency high data-rate GPS receivers were connected to the four existing aircraft antennas. The signal from each antenna was split to allow the ADU2 and the independent receivers to operate simultaneously. The attitude accuracy of this system should be as good as, or better than, the ADU2. This system allows positions to be determined with centimetre level accuracy using fixed ambiguity carrier-phase processing methods, provided a local reference station is available. Details of the trial flights in which this system was tested are given in the following section.

4.3 Flight Trials

4.3.1 Test Site

Two test flights were undertaken on October 15th and 16th, 1998 around Peterborough Business Airfield, Conington, Cambridgeshire, UK. This site was chosen for:

- its convenience in being able to land and install equipment on the aircraft,
- its proximity to ITE Monks Wood where a permanent GPS reference station is operated,
- its suitability as a survey area due to the well defined shapes of buildings and hard surfaces, and,
- the good relations that had already been established between the airfield operators and personnel at the NERC ARSF.

The survey site was approximately 4 x 3km and the original flight plan (figure 4.11) was designed so that there were large areas of over-sampling to aid calibration. The degree of overlap varies for the two different scanners as their swath widths vary. Five flight lines were flown in an east-west direction and four north-south. A number of targets were placed on the run-off areas around the runway and surveyed prior to the flights using GPS stop-go methods³. The accuracy of these control points is of the order of a few centimetres. In addition, aerial photography was taken during the flight (1:4,000 scale) to allow a more extensive network of ground control to be built up if required. The control could also be extended using ground survey methods to position features imaged by the scanners, once the scanner imagery is available.

³ Using Leica SR-399 GPS equipment and the SKI processing software.

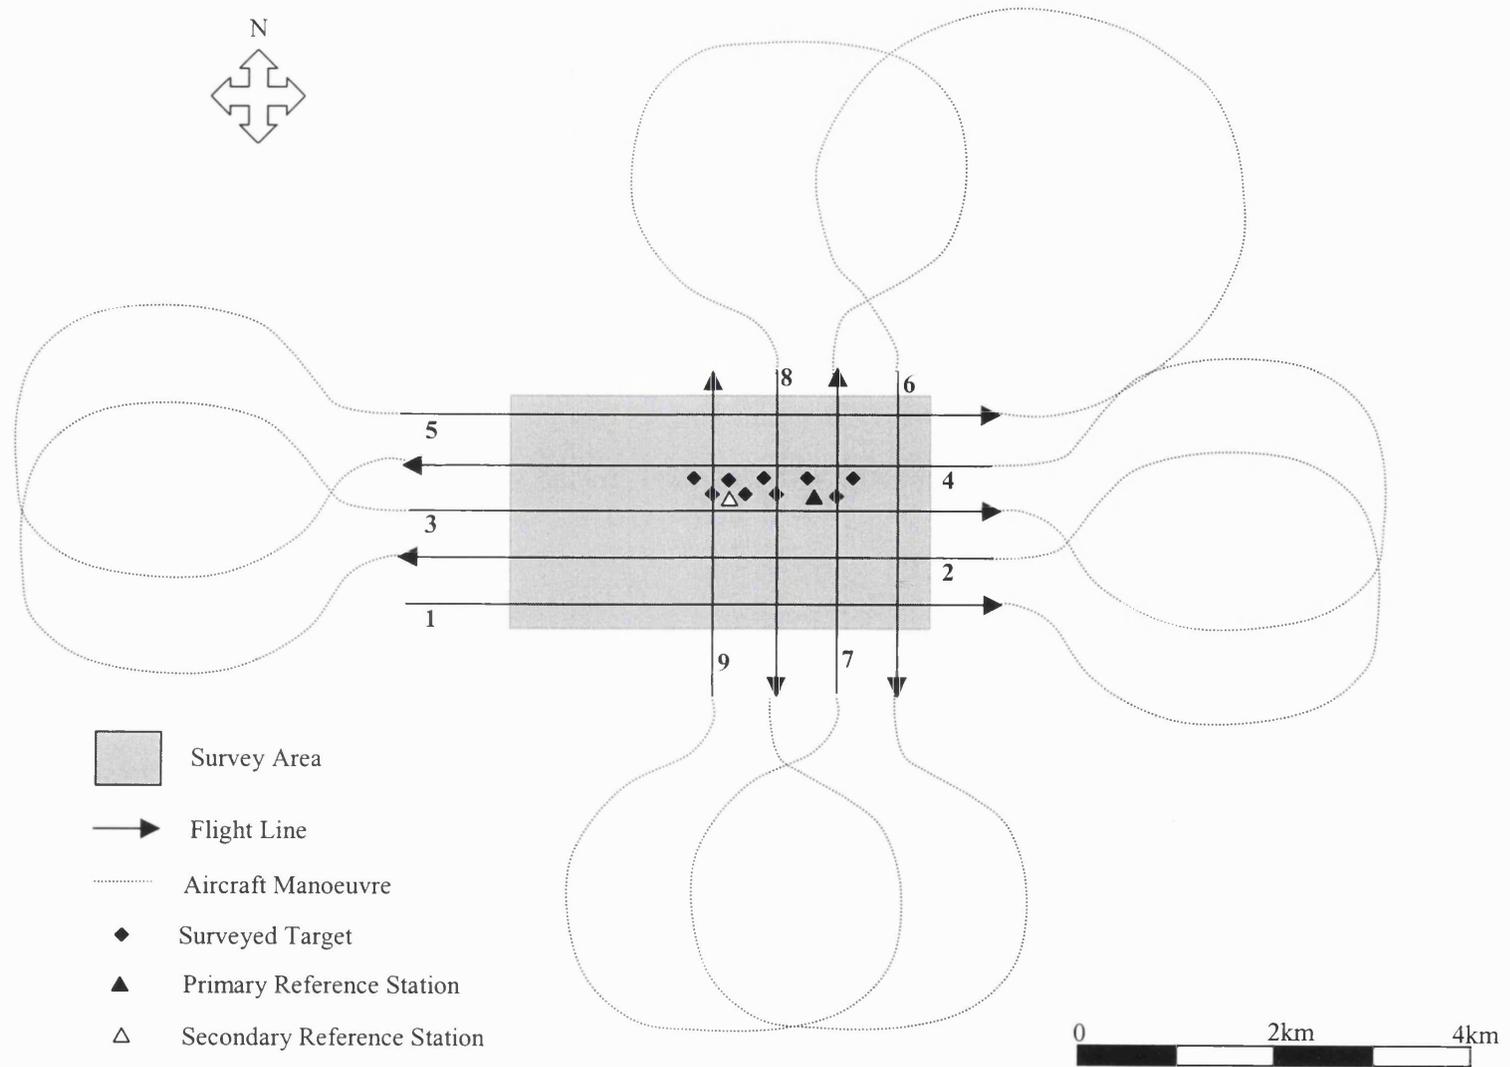


Figure 4.11: Flight Plan

4.3.2 Equipment Specifications

The survey was designed to provide the best possible spatial resolution from the scanners (i.e. smallest pixel sizes). To achieve this, the aircraft was flown at its minimum survey altitude of 600m, and at its minimum air speed of 100 knots (51m/s). The spectral range of the CASI was minimised to allow rapid integration times. Details of the flight specification are given in table 4.3.

	CASI	ATM
Pixel size (m)	1 x 1	1.5 x 1.5 near nadir (3 x 4.5 scan edge)
Scans per second	43	50
FOV (°)	42	90
Swath Width (m)	460	1,200

Table 4.3: Scanner specifications during test flights

The independent GPS receivers⁴ installed on the aircraft were:

- Two Ashtech Z-Surveyors.
- One Ashtech Z-Sensor.
- One Ashtech Z-X12.

Two local reference stations were established in the survey area using a Leica MC1000 receiver with a choke-ring antenna and a Leica SR-399 with integrated antenna. Additionally the permanent reference station at ITE Monks Wood was recording during the test flights. This station is within 10km of the test site (7.3km from the local reference).

Three of the aircraft receivers (the Z-Surveyors and Z-Sensor) and one of the reference receivers (the Leica MC1000) have a maximum recording rate of 10Hz, which makes it possible to obtain values for all six position and attitude components at 10Hz. The Z-12 has a maximum recording rate of 5Hz, which therefore defines the rate at which a four antenna attitude solution is available.

⁴ Supplied by Ashtech Europe Limited and CBI (Colin Beatty International) Ltd. Their support is gratefully acknowledged.

The Leica SR-399 reference station acts as ‘back-up’ in case of problems with the 10Hz reference, and allows the possibility of using some multiple reference solutions for positioning the aircraft at 1Hz. For the duration of the flights, the ITE station was also operating at 10Hz. This provides another source of reference data in case of problems, and also allows solutions based on multiple reference stations, or single reference stations at different distances from the survey area.

During the flight, the Z-Sensor and Z-12 logged to lap-top PCs as the Z-Sensor has no internal storage device and the Z-12 has insufficient internal memory to record for the two hours from the start of recording to the end of the survey. The Z-Surveyors each had a 20Mb memory card, which can theoretically record for around 90 minutes at 10Hz under normal observing conditions. The primary local reference station, the MC1000 records data to a 20Mb internal storage device using a compressed binary format and could therefore operate throughout the test period.

4.3.3 Calibration

The aim of the two test flights was to assess the performance of the proposed direct georeferencing system, in terms of accuracy, reliability and continuity. Ultimately the system has been designed to position imaged features in a ground co-ordinate system without using ground control points. The accuracy with which the on-board system accomplishes this can be determined by comparing the co-ordinates of ground targets derived from scanner imagery, processed with directly derived orientation parameters, against the target co-ordinates from the ground survey. From this it is possible to establish the accuracy with which the scanner position and attitude has been estimated.

As the scanners were operating at their highest resolutions, it is possible to assess how accurately a feature has been positioned with a certainty of around one metre. The general tasks undertaken by this aircraft require an accuracy of around ten metres, so it will be possible to test if this is being achieved, and hopefully to assess the degree to which it is exceeded. If the ADU2 attitude accuracies are translated into on-the-ground displacements based on a flying height of 600m, the combined position error is 3.75m (most of this error is contributed by the roll component). In addition, the sensor positioning error could contribute another 1 to 2m, giving a final on-the-ground accuracy of around 5m. The proposed system, based on independent receivers, should deliver comparable, or better, attitude accuracy and the sensor positioning error should have no significant impact on the final accuracy. It should therefore allow positioning accuracies of 4m or better. The attitude solution from the AHRS, and the position and attitude solution from the ADU2, also provide further data with which to compare the test solution.

During the test flight preparation, it was intended that an aerial camera with digital input and output would be flown on-board the aircraft. The exterior orientation parameters of each exposure station could be determined from standard photogrammetric techniques (i.e. a resection using ground control points) and from the on-board instruments. This would provide accurate exterior orientation parameters at a limited number of discrete points. Digital input or output is required to synchronise the instant of exposure to the relevant position and attitude data. As it was not possible to install a camera with this capability for these test flights, due to the considerable time and expense it would have required, this form of calibration data is not available.

4.3.4 Operational Difficulties

A number of problems were experienced during the flight and in the subsequent processing stages that have restricted the available data and therefore altered the emphasis of this project.

- Logging problems with the independent GPS receivers were experienced during the flight, due to a combination of hardware and software problems. As a result, the availability of data from each antenna varies significantly. During the first flight almost no GPS data were logged from the starboard antenna, and there were also data gaps from the other antennas. After changes made between flights, GPS data were recorded by all four antennas on the second flight, although there were still significant gaps.
- Due to the wind conditions, it was only possible to fly from east to west, and not from west to east . This increased the flight duration and as a result the GPS receivers recording to internal memory devices stopped recording before the survey was completed.
- As the weather conditions deteriorated significantly during the second test flight it was aborted halfway through.
- Due to operational difficulties experienced by the NERC ARSF, no scanner imagery recorded during the flight has been provided. This has prevented the full calibration of the georeferencing system.
- Similarly, no ADU2 data from the flight has been provided, therefore the proposed comparison of a dedicated GPS attitude system and an independent receiver system has not been possible.

4.4 Concluding Remarks

This chapter has described the equipment and test procedure used to calibrate the proposed direct georeferencing solution. A discussion of the NERC ARSF sensor characteristics, and in particular their ground spatial resolutions, has helped to define the attitude and position accuracy requirements. Provided the sensor position can be established using standard carrier-phase solutions, the on-the-ground accuracy can be defined in terms of sensor attitude accuracy at a given altitude, only. The system is required to deliver accuracies of 3 to 5' in heading, 5 to 10' in pitch, and 10 to 20' in roll. Ultimately the accuracy with which the roll component can be determined over the shortest baseline (1.2m) will limit the overall system accuracy.

The on-board equipment used by NERC to fulfil these criteria, i.e. the Ashtech ADU2 and AHRS gyro, has been described. The independent receiver solution proposed in this project was then outlined in terms of the equipment used.

Two test flights to gather the data needed to validate the proposed system were undertaken. The proposed system was going to be calibrated using ground control information and reference solutions from the NERC ARS facility's ADU2 and AHRS instruments. Unfortunately, the only reference solution available is the AHRS data, due to the operational difficulties indicated. This lack of calibration data has made it impossible to give a full assessment of the overall system accuracy. Instead, the degree of agreement between the gyro and GPS solutions, and the internal consistency of the GPS solution is used to make inferences about the likely system accuracy.

Chapter 5 Position and Attitude Determination from GPS Flight Data

5.1 Introduction

This chapter describes how the GPS data recorded during the test flights, from the four independent GPS receivers on the aircraft and the primary reference station, have been processed. The algorithms used for position and attitude determination are described, together with the software in which these algorithms are implemented.

Positions have been determined from SKI, a commercially available GPS processing package from Leica, and from GASP, the GPS Ambiguity Searching Program, developed initially at the University of Newcastle upon Tyne and later University College London. Attitude values have been computed directly from raw GPS observations using software known as the GPS Routine for Attitude Parameter Estimation (GRAPE). The GRAPE software implements the direct attitude estimation procedure based on modified double differenced observation equations, as described in section 3.5. Both GRAPE and GASP resolve ambiguities using GPS data from a single epoch. By computing positions or attitude values on a single epoch basis this method is immune to cycle slips or loss of lock. The relevant features of previously developed software are described before introducing the changes and extensions that have made in the course of the current research project.

Section 5.2 gives an overview of the complete data processing strategy, describing in general terms the main steps in the transition from raw GPS data to position and attitude values. In section 5.3 position determination is discussed in terms of algorithms, software and previously documented performance. Section 5.4 describes the design of GRAPE, showing how the direct attitude determination algorithms have been applied in existing software. An assessment of the attitude solution generated by GRAPE in a static environment is given to establish what has been achieved during previous research. In section 5.5 a number of changes

and extensions to the previously developed algorithms and software are described. These features are novel to the current work and have been implemented both to improve the performance of the existing software and to extend its functionality to allow it to process the airborne GPS data collected in the course of this research. Some concluding remarks are made in section 5.6.

5.2 Processing Overview

Figure 5.1 illustrates the data flow from raw GPS data, recorded on the aircraft and at the primary local reference station, through to filtered position and attitude values. The filtering stage of this process, including the integration of gyro data, is dealt with in chapter 7. These steps have been included in the diagram for completeness. All rectangles represent data sources, in various formats. Ellipses represent processing software.

The starting point for the processing scheme is the set of five raw GPS files recorded during each flight. The four files from the aircraft receivers are all in an Ashtech proprietary binary format ('b' and 'e' files). The file from the reference station is in a compressed Leica binary format ('lb2' format). All these files are converted to the widely used GPS data exchange format, RINEX (Receiver Independent Exchange, see Gurtner (1994)), using the appropriate manufacturer's software ('Ashtorin' for Ashtech data, 'LB2R' for Leica data). An observation file (*.98o) and a navigation file (*.98n) are created for each data-set.

For use with GASP and GRAPE software, these RINEX files must first be converted to NXF (Newcastle Exchange Format) files, using the RINtoNXF conversion software. The NXF format and conversion software were developed at the University of Newcastle (Corbett, 1994) to provide a more compact and legible exchange format than RINEX. It was also designed to assist the development of subsequent processing packages by carrying out a number of time consuming and iterative calculations that are required as a preliminary step in most GPS positioning algorithms. These calculations are primarily concerned with the calculation of WGS84 satellite co-ordinates from Keplerian elements.

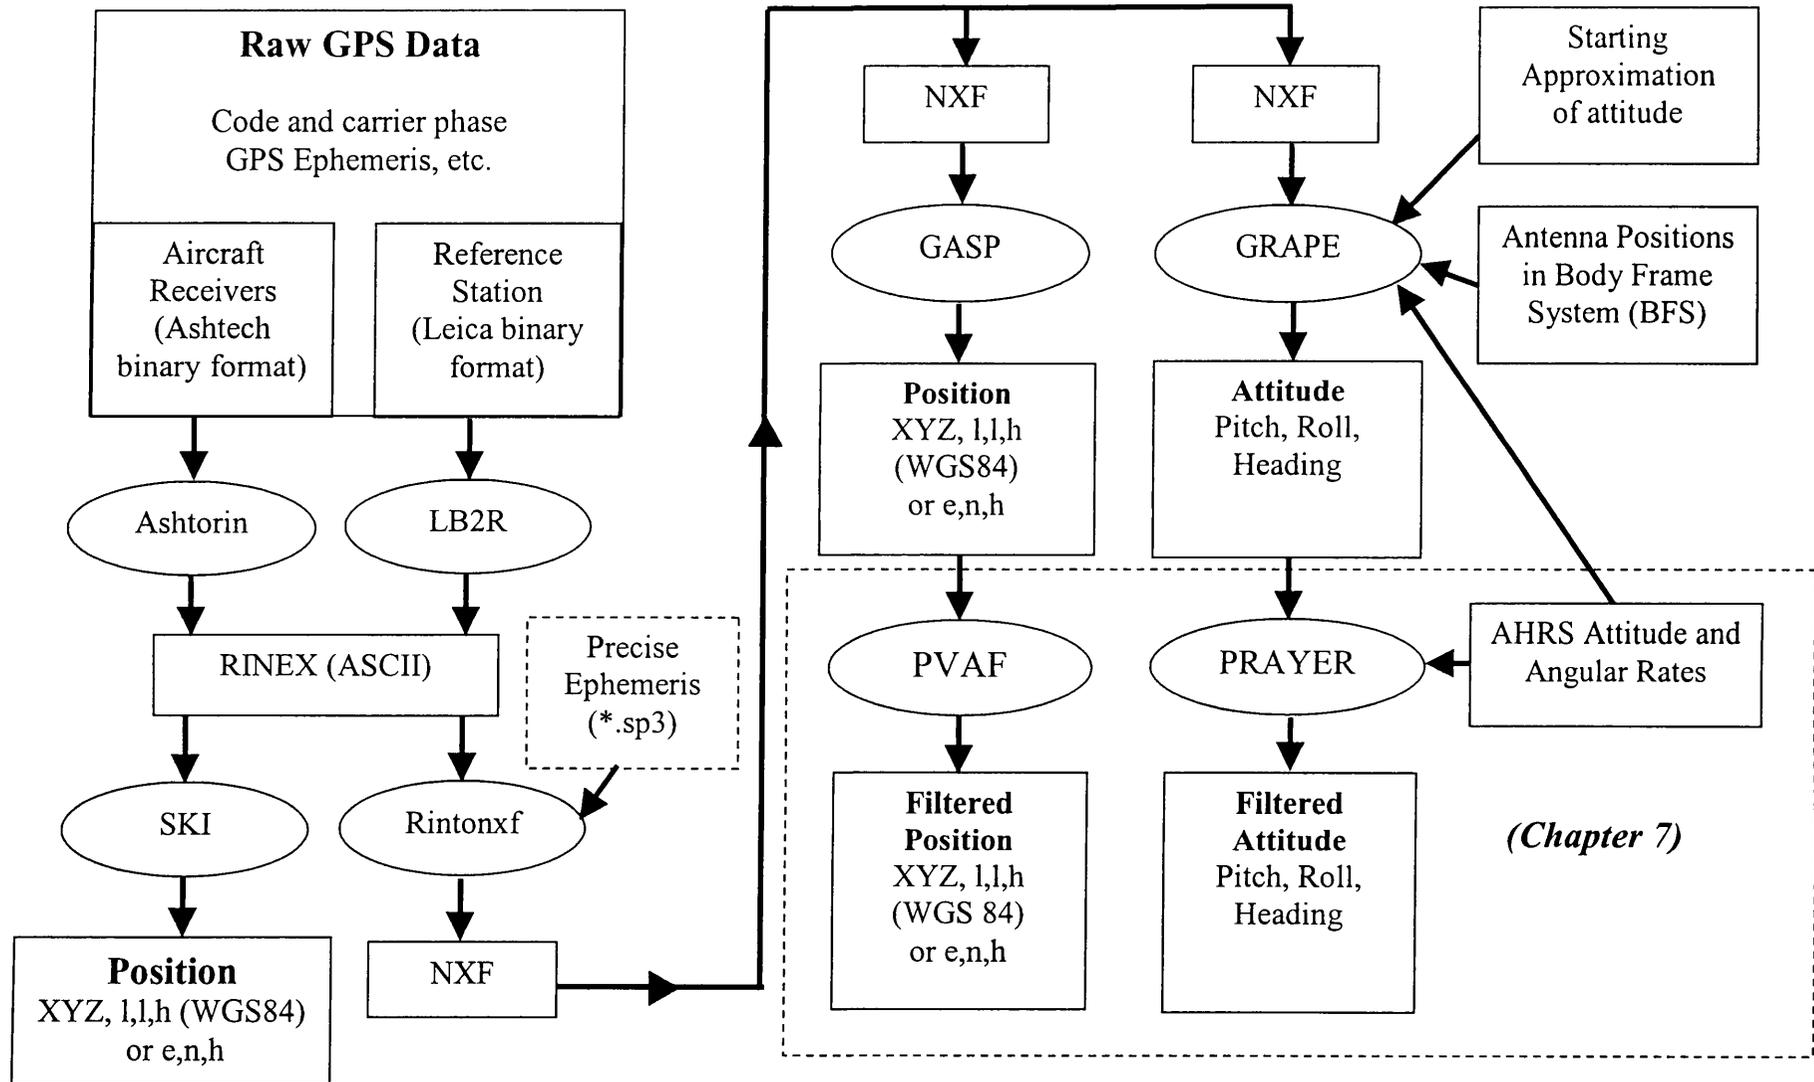


Figure 5.1: Processing overview of GPS position and attitude determination

The program also computes an approximate receiver position using a single point, single epoch code solution. The RINtoNXF program allows precise ephemeris files or broadcast orbits to be used. Although precise ephemeris data was used instead of broadcast data, it should not produce a significantly better solution in this application due to the limited antenna baselines. The aircraft stayed within 8km of the reference station throughout the whole flight, and was considerably closer than this during the flight lines on which imagery was recorded. Once the conversion from RINEX to NXF is complete, all the original GPS data from a single receiver is stored in a single file. These files are the basic input for both the position and attitude determination programs, GASP and GRAPE respectively.

Once the data is in RINEX files, it is also possible to process the GPS data in a standard commercial software package to determine positions. With a RINEX observation file for the reference receiver and any one of the aircraft receivers, and appropriate ephemeris data, either from a RINEX navigation file or a precise ephemeris file, a solution for the aircraft trajectory can be determined. This has been done using Leica's SKI processing software. A brief description of this GPS processing software is given in section 5.3.3.

Positions can be determined with GASP using the NXF files from the reference station and any one of the aircraft receivers. Section 5.3.2 gives some details of this software and the algorithms it employs. To compute attitude values using GRAPE, GPS data, in the form of NXF files, must be available for at least three of the aircraft antennas. In addition the antenna co-ordinates in a fixed body frame system (BFS) and a starting approximation for each attitude component must be available in an initialisation file. The GRAPE program and additions made to it during the course of this research are described in detail in sections 5.4 and 5.5.

Once attitude and position values at each epoch have been determined by GRAPE and GASP these values can be entered into a Kalman filter if required. The PVAF (Position, Velocity, Acceleration Filter) program filters position components and the PRAYER (Pitch Roll And Yaw Estimation Routine) program filters attitude

values. In the case of attitude, the GPS derived solution can also be integrated with attitude values and angular rates from the AHRS gyro, within the filtering mechanism. The design and operation of these filters are described in chapter 7.

5.3 Position Determination

5.3.1 Overview

The position determination methods used in this research project are relatively standard and have not been adapted significantly in the course of this work. The description of the techniques will therefore be brief. However, as position is obviously an important part of the overall direct georeferencing system being developed, some detail of its determination must be included.

For the current georeferencing application, position and attitude values are required for each scan line. If GASP and GRAPE can both successfully resolve ambiguities, and therefore accurately compute position and attitude respectively using a single epoch of GPS data, then all six exterior orientation parameters can be determined on a single epoch basis. Such a solution may be preferable to one using multiple epochs to resolve ambiguities, as is used in SKI, because it is immune to cycle slips and loss of lock. However, as long as position and attitude values of the required accuracy are available at each epoch, regardless of the processing method, the georeferencing requirements of this application can be satisfied. Therefore, whilst GASP may be considered the primary positioning tool in this application, a successful SKI solution will meet the objectives equally well.

An overview of GASP's processing strategy, and details of previous performance assessments are given in section 5.3.2. A brief description of the processing algorithms used by SKI are given in section 5.3.3.

5.3.2 GASP

5.3.2.1 Methodology

The GPS Ambiguity Searching Program (GASP) was originally developed at the University of Newcastle upon Tyne by Dr. Simon Corbett (Corbett, 1994), and has since been enhanced by the original author, Dr. Yahia Al-Haifi (Al-Haifa, 1996), and Mr. Joel Barnes. It uses a technique based on the Ambiguity Function Method, AFM, first considered by Counselman and Gourevitch (1981), to resolve integer ambiguities, and hence compute position, from a single epoch of dual-frequency GPS data. The single epoch method was adopted so that positions can be successfully computed in applications when continuous lock to satellites cannot be guaranteed. Traditional multiple epoch solutions require data from a period of uninterrupted tracking to resolve ambiguities. After a loss of lock or cycle slip, it may take a few minutes to correctly resolve ambiguities. With a single epoch solution this 'warm-up' period is not required. As soon as sufficient satellite information is available at an epoch, ambiguities can be resolved and a position can be computed with centimetre level accuracy.

Figure 5.2 illustrates the main steps involved in GASP position processing. Many of these procedures are also common to GRAPE and are described in more detail in section 5.4. To provide a general overview, these steps can be divided into three main operations:

- 1) An approximate position for the unknown receiver is computed using a double differenced code solution. Depending on the quality of the GPS data and the baseline length, the code solution typically yields approximate co-ordinates to within 0.6m (2σ) (Al-Haifi *et al.* 1998). A search volume is formed, centred on this initial approximation, containing a set of trial positions, each producing a unique set of ambiguity values.

- 2) To identify the correct ambiguity for each double differenced observation, an ambiguity function value (AFV) is computed from L1, and weighted L2 double differenced residuals, for all the candidate positions within the search space. All the candidate positions producing an AFV within a specified range of the maximum AFV are then selected for further testing. This process is described further in section 5.4.1.6, in the context of GRAPE.

- 3) A least squares adjustment is carried out for the selected ambiguity candidates, producing an estimated position with associated quality indicators. A statistical test is then conducted to determine whether the ambiguity candidate producing the lowest sum of the squares of the weighted residuals ($\hat{v}^T W \hat{v}$) is significantly more likely to be correct than any alternative candidates.

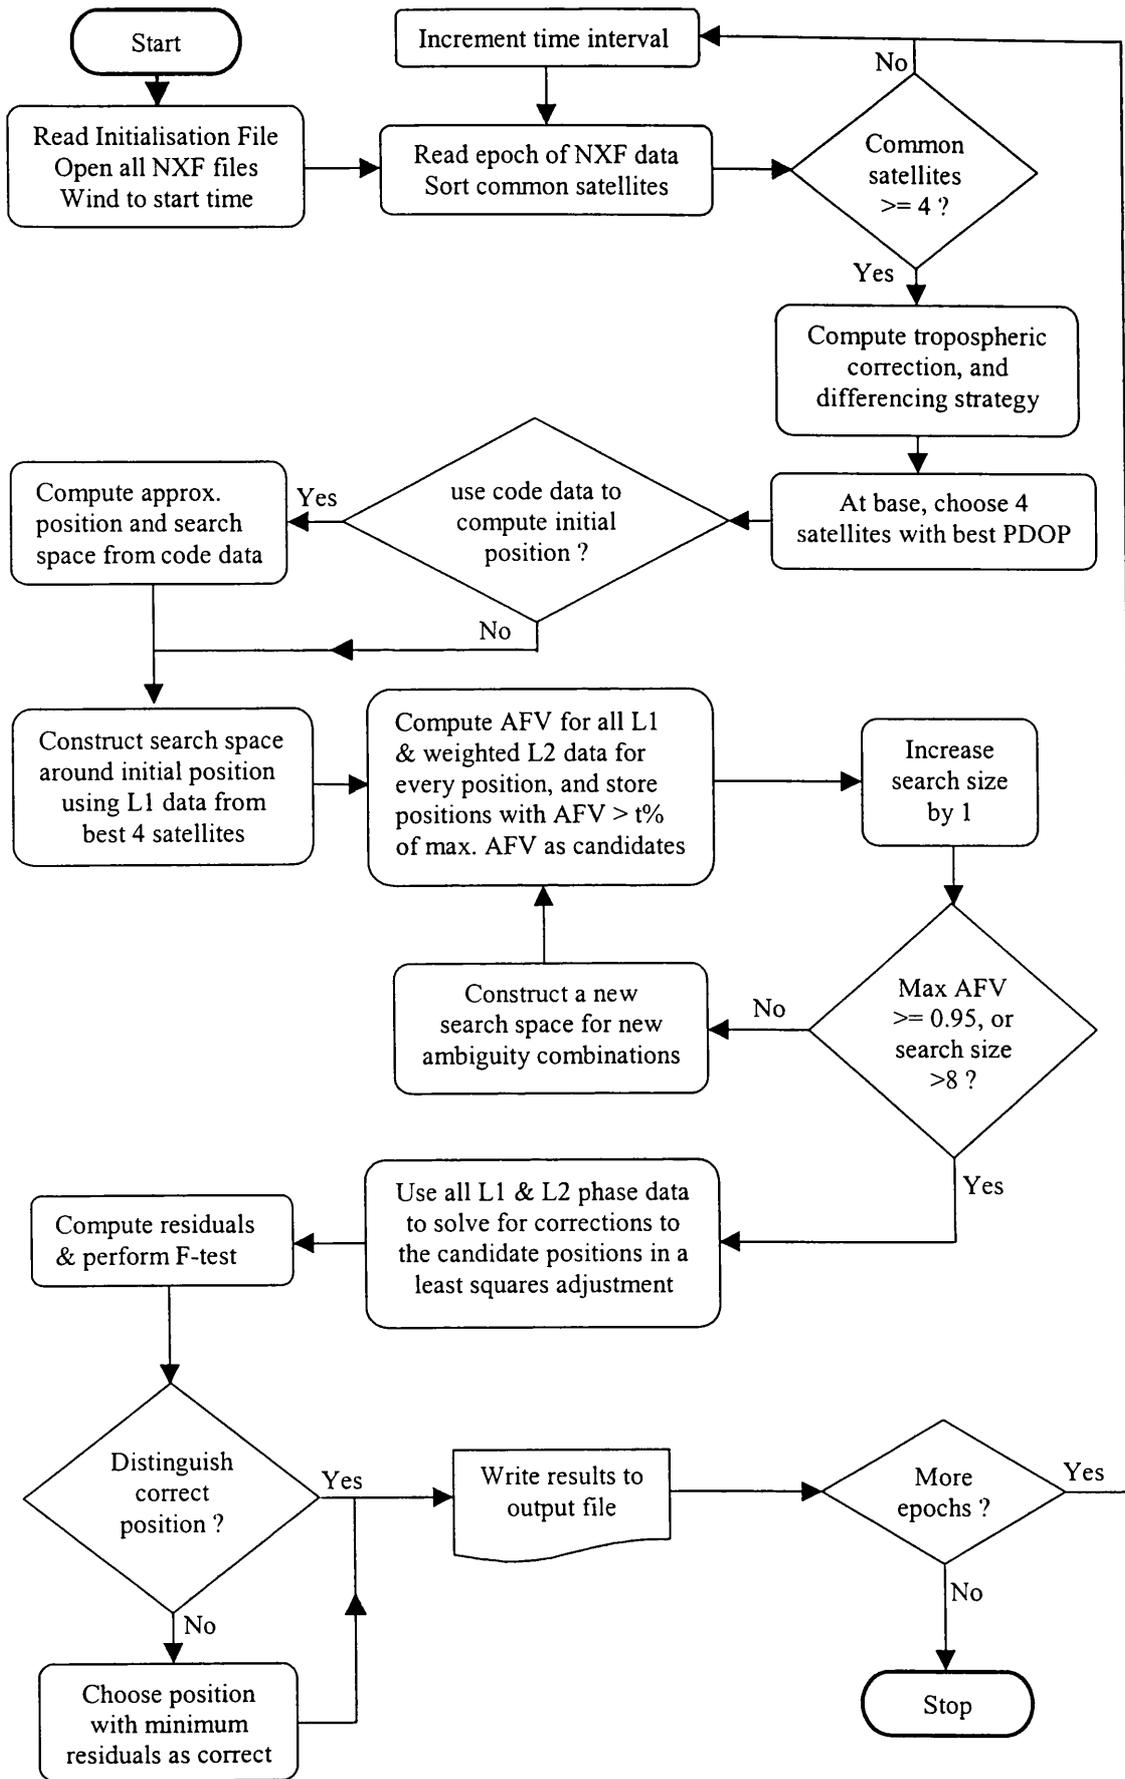


Figure 5.2: GASP processing (adapted from Al-Haifi *et al.*, 1998)

GASP can be used to position a single unknown roving receiver relative to one or more reference stations. Both the reference and roving receivers may be static or moving. The ability to determine baseline vectors when both receivers are moving make this a suitable algorithm for attitude determination using a baseline-by-baseline method, as described in section 3.4.2.1¹.

Further details of these algorithms and a full description of the software can be found in Corbett (1994), Corbett and Cross (1995) and Al-Haifi (1996).

5.3.2.2 Previous GASP Performance

In previous research projects, GASP's robustness and accuracy has been assessed in a range of operating conditions. In static tests, in low multipath environments ambiguities can be correctly resolved 100% of the time. It has also proved successful in kinematic tests over short ranges. Al-Haifi *et al.* (1998) describe a walking test in which the mobile receiver was carried in a backpack with the antenna fixed to a ranging pole, and the reference receiver was at a fixed point within 500m. A comparison of GASP co-ordinates with a solution from Ashtech's PNAV software showed agreement to within 3mm in the horizontal component and 5mm vertical. This level of agreement indicates that ambiguities have been correctly resolved at every epoch and that centimetre level accuracies are achievable.

As phase noise increases however, the performance of the system deteriorates significantly. When the software was tested with airborne data in 1994, correct ambiguity resolution occurred at less than 7% of the total epochs (Corbett, 1994). A number of refinements to the original algorithms have significantly improved the performance over the years². Over the operating distances in the current application, multipath has been identified as the limiting factor on the success of the single ambiguity resolution approach.

¹ Once vectors are determined a further step is clearly required to compute attitude. In the case of redundant measurement information, this will generally be a least squares estimation.

² Most of these refinements were implemented by Dr. Al-Haifi after a performance assessment using the original routines (Al-Haifi, 1996).

Simulation studies to quantify the extents of allowable noise on the phase measurements have shown that the algorithm will fail if the L1 phase data is in error by more than 1.0cm for a five satellite constellation, and 2.5cm for a seven satellite constellation (Cross *et al.*, 1997). Previous tests using GPS data collected from the same survey aircraft used in the current research, show the phase noise levels may be up to 2.0cm (Corbett, 1994) which will clearly influence the success of the algorithm, particularly when fewer than seven satellites are observed. Further details of GASP performance assessments can be found in Corbett (1994), Corbett and Cross (1995), Al-Haifi (1996) and Al-Haifi *et al.* (1998).

5.3.3 SKI

SKI (Static KInematic software) is Leica's commercially available GPS post-processing package. It has been designed primarily for high accuracy carrier-phase positioning over limited distances (generally for antenna separations below 15km). SKI is widely used in the GPS surveying community and has been shown to deliver centimetre level positioning accuracy in reasonable observation conditions. The version used in this research was SKI 2.2., an updated version of this software, called 'SKI Pro', is now available and claims to offer improved performance, but none of the flight data have been processed with this newer version.

SKI 2.2 resolves ambiguities using multiple epochs of data using the Fast Ambiguity Resolution Approach, FARA (Frei and Beutler, 1990). The steps involved in FARA have been summarised by Erickson (1992) as follows:-

(1) *Compute Float Carrier Phase Solution*

- estimate real values for each double difference ambiguity
- compute residuals and *a posteriori* variance factor (or unit variance)
- scale covariance matrix of parameters with *a posteriori* variance factor

(2) *Choose Ambiguity Sets to be Tested*

- choose individual integer ambiguities which fall within the confidence range of the real estimates to form sets of potential integer solutions
- reject ambiguity sets which have ambiguity pairs which exceed the confidence range of the differences of the real ambiguity estimates

(3) *Compute Fixed Solution for Each Ambiguity Set*

- compute fixed solutions, variance factors and parameter covariance matrices

(4) *Statistically Test the Fixed Solution with the Smallest Variance Factor*

- test if fixed and float solutions are compatible
- test variance factor for normal distribution
(χ^2 test on the variance factor)
- compare the smallest variance factor with the second smallest

In SKI, a condition to enter FARA is that in the case of dual frequency measurements a wide laning solution has been possible. If all the statistical tests described are successfully passed, the ambiguity resolution is accepted as successful.

SKI has been used in this research as an alternative to GASP for computing aircraft positions. A comparative solution of this kind, that uses the same raw data but applies very different processing methods, helps to validate the single epoch solution and to determine whether any problems are due to poor quality raw data or the processing methods used.

5.4 Static GPS Attitude Determination

The attitude solution used in this research has been determined by processing GPS data from the four independent receivers on the aircraft, with the GRAPE software program. GRAPE was developed by Dr. M.R. Mahmud to apply the single ambiguity resolution techniques of GASP to the problem of attitude determination (Mahmud, 1999). It implements the direct attitude algorithm described in section 3.5, i.e. attitude components are determined from a single epoch of dual-frequency GPS data using a modified double differencing algorithm.

The following section describes the design and testing of this software. Results from static tests have shown that implementing the direct attitude approach, with the increased redundancy it offers, has led to performance improvements relative to a baseline-by-baseline system, as described in section 3.4.2.1. A summary of the system performance is given, and the performance benefits it offers are described. During the current research project the existing algorithms and software have been adapted and extended to allow the successful processing of airborne data (section 5.5). To demonstrate these extensions however, it is first necessary to describe the methodology of the original software and to summarise some of the key performance results.

5.4.1 GRAPE Methodology

The following is a summary of the full description of the GRAPE processing strategy presented by Mahmud (1999). The processing involved in determining attitude components from raw GPS data using GRAPE can be divided into seven stages:

- (1) The determination of the body frame co-ordinates of the antennas.
- (2) The determination of approximate attitude.
- (3) The generation and reading of NXF files.
- (4) The formation of a search space.
- (5) The computation of modified double differenced phase values.
- (6) The computation of the ambiguity function values.
- (7) A least squares computation of the attitude values and associated quality measures with statistical tests to identify outliers and to confirm if the correct ambiguities have been selected.

The first three can be considered as preliminary steps, i.e. they are all carried out prior to actually running the GRAPE software. Details of how these steps were carried out in the course of the current research have been included where appropriate. The remaining four stages are all performed within the program. Figure 5.3 is an overview of the GRAPE processing algorithms, and clearly shows how the attitude determination software has evolved from the positioning software, GASP (Figure 5.2). Figures 5.4 and 5.5 provide further details of the ambiguity resolution and least squares adjustment stages respectively. The preliminary stages will be dealt with first, then a more detailed description of the processing algorithms within GRAPE will be given.

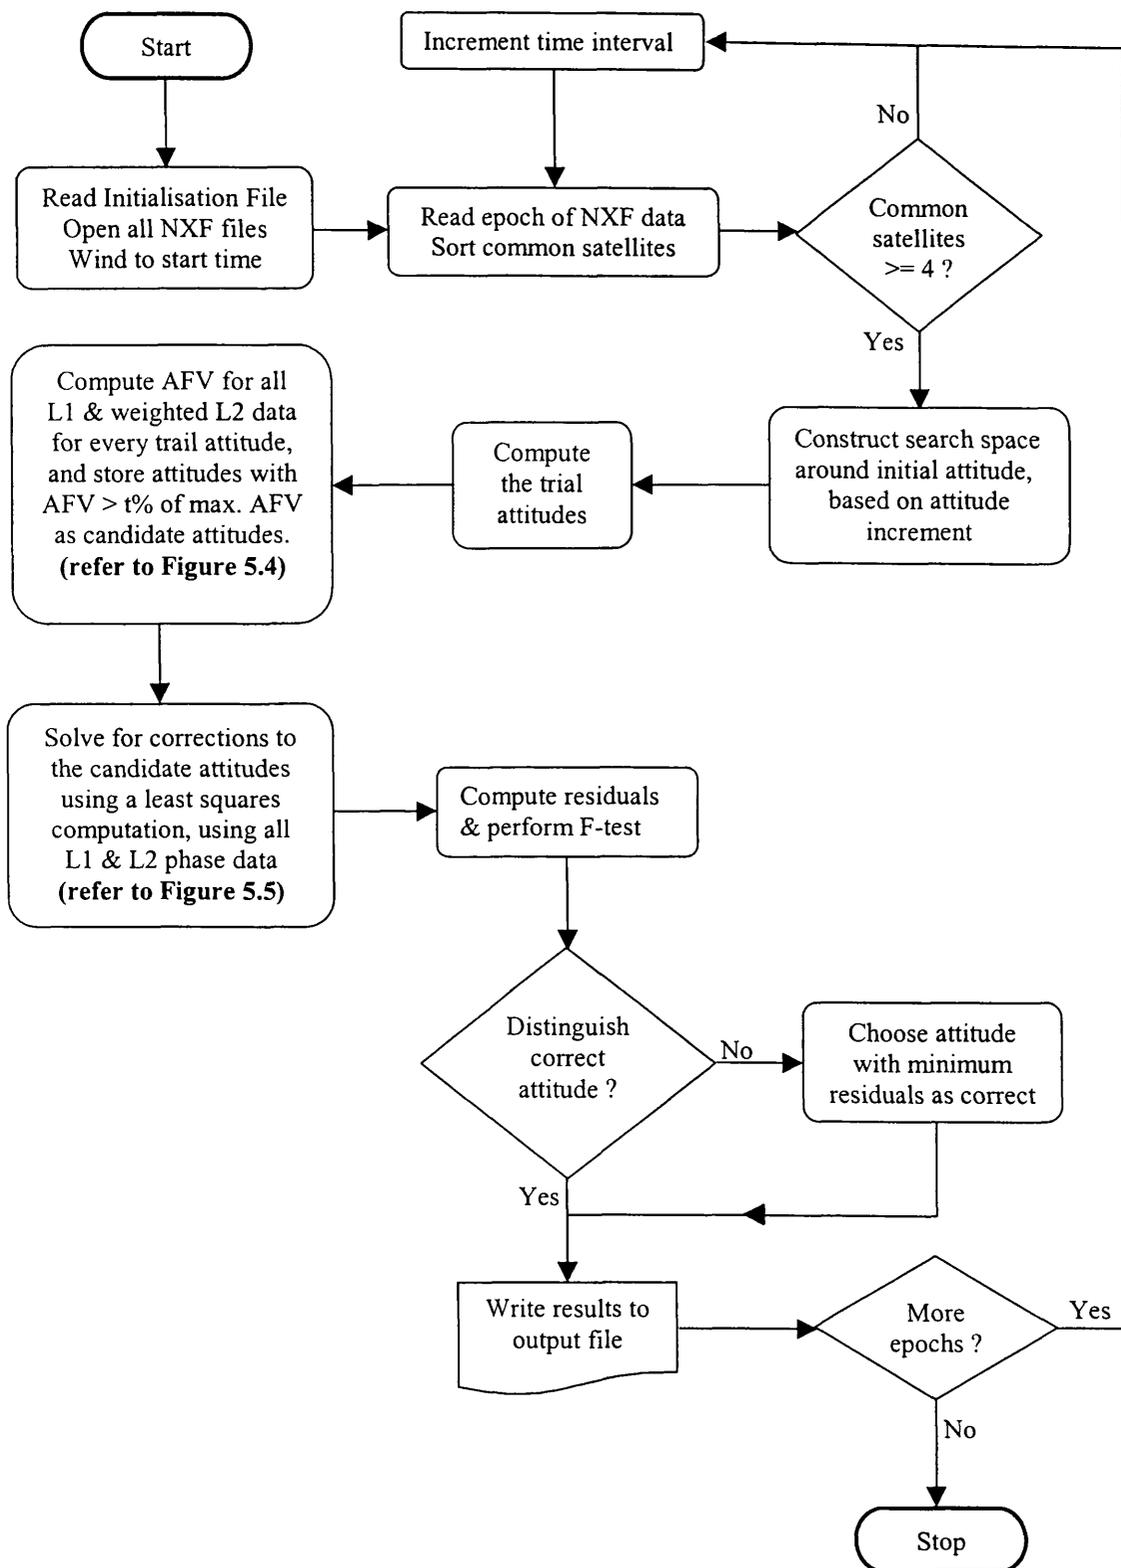


Figure 5.3: GRAPE processing (adapted from Mahmud, 1999)

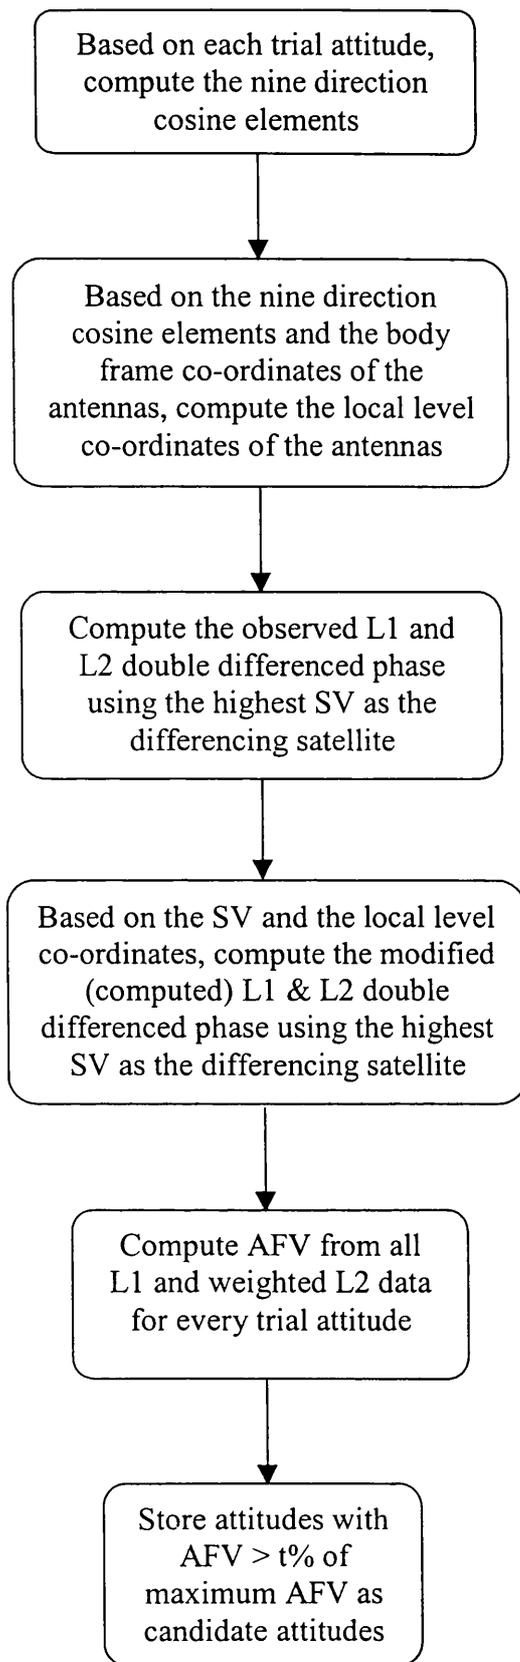


Figure 5.4: Ambiguity resolution in GRAPE (Mahmud, 1999)

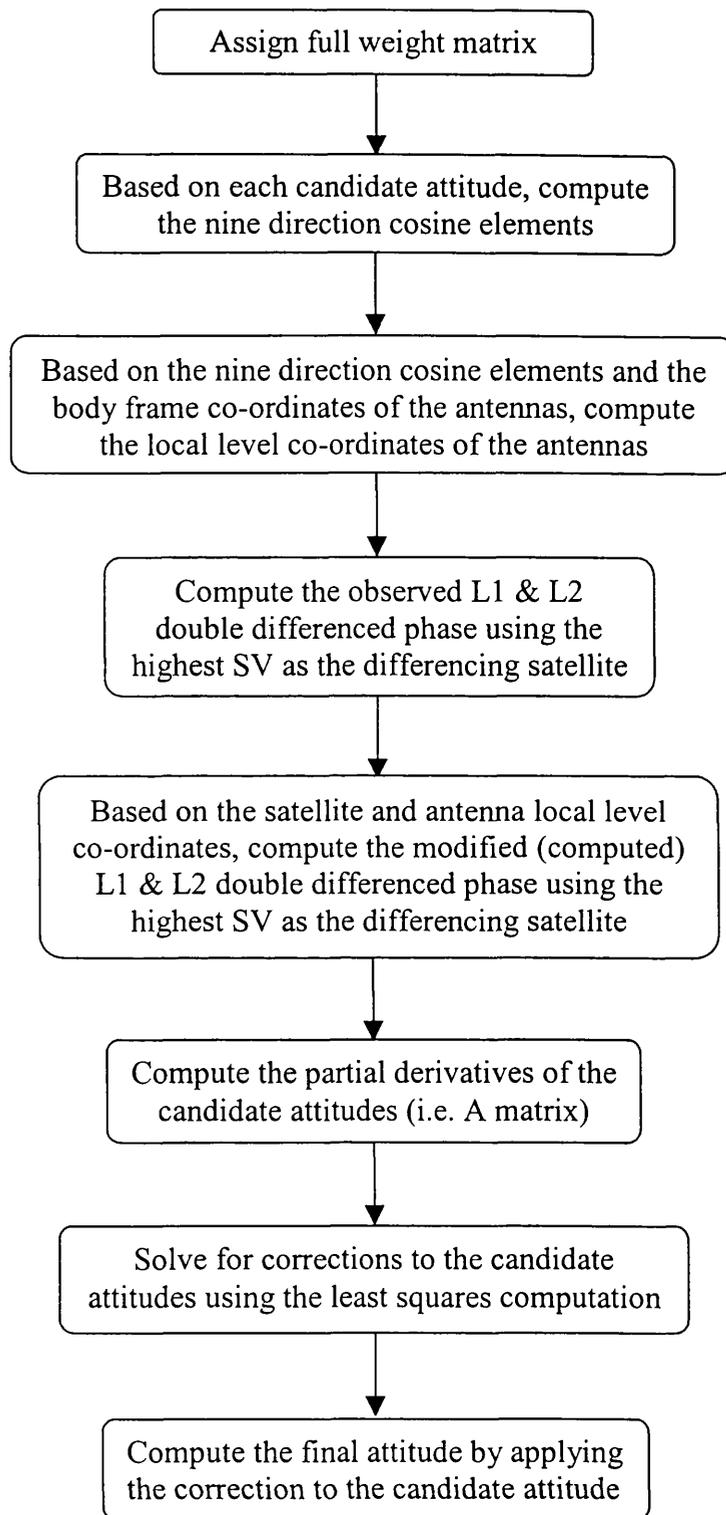


Figure 5.5: Least squares routine in GRAPE (Mahmud, 1999)

5.4.1.1 Determination of Antenna Body Frame Co-ordinates

In GPS attitude determination, attitude is defined using the antenna positions in a fixed body system and an external reference system, as described in section 3.2. Clearly then, the antenna positions in the body frame system must be established before attitude can be determined.

In the current application the co-ordinates of the aircraft antennas in the body frame system must be determined together with the co-ordinates of the scanning instruments so that GPS measurements can be related to the imagery. The antenna survey can be carried out using a static GPS survey, or by conventional survey methods, i.e. measuring angles and distances. A static GPS survey allows the vectors between antenna phase centres to be directly measured with a high accuracy, but it is not possible to position the imaging devices or inertial systems relative to the antennas by this method as they are inside the aircraft. A complete survey of all the antennas and sensors could therefore be achieved either, by using conventional survey methods throughout, or by a hybrid survey, in which antennas are positioned using GPS, and the other instruments are positioned relative to the antennas using conventional methods

The NERC ARSF aircraft has been surveyed in the past using conventional methods to position the four GPS antennas and all the imaging devices. The body frame system has been defined so that the y-axis lies on the centre-line of the fuselage, as described in section 3.2. As no sensor imagery has been available for this research, the offsets between the antennas and sensors are not applied, hence only antenna co-ordinates have been considered. These are presented in table 5.1.

Antenna Number	Location	Body Frame System Co-ordinates (m)		
		x	y	z
1	Rear fuselage	0.000	0.000	0.000
2	Front fuselage	-0.074	3.556	0.030
3	Port winglet	-0.602	2.042	0.001
4	Starboard winglet	0.591	2.046	0.001

Table 5.1: BFS co-ordinates from NERC ARSF antenna survey

Due to delays in receiving the antenna survey results from the NERC ARSF, antenna co-ordinates were also determined using static GPS methods. It was not the intention to perform a dedicated antenna survey so no data were collected specifically for this purpose. Instead, the data used was taken from a period of less than five minutes in which the aircraft remained stationary on the apron area with the engines running. To perform a dedicated static antenna survey a period of 20 to 30 minutes of data should have been recorded, preferably with the engines off and with the aircraft in an open area away from buildings. Nevertheless, in this case an antenna survey was carried out using a small segment of GPS data which appears to have relatively low levels of noise. The antenna positions from this survey, determined with SKI software, are presented in table 5.2. The figures in brackets show the difference in antenna co-ordinates from the two surveys.

Antenna Number	Location	Body Frame System Co-ordinates (m)		
		x	y	z
1	Rear fuselage	0.000 (NA)	0.000 (NA)	0.000 (NA)
2	Front fuselage	-0.074 (0.000)	3.550 (0.006)	0.005 (0.025)
3	Port winglet	-0.603 (0.001)	2.042 (0.000)	-0.006 (0.007)
4	Starboard winglet	0.588 (0.003)	2.042 (0.004)	0.026 (0.025)

Table 5.2: BFS co-ordinates from static GPS antenna survey

Whilst the agreement between the two surveys is generally very good, at the level of a few millimetres, there are some significant differences in the z co-ordinate of antennas 2 and 4, 25mm in each case. This may be due to differences in the horizontal planes defined in each survey, differences in the actual point on each antenna that was measured, biases in the GPS solution caused by multipath, or more likely, a combination of all of these. Without details of the NERC ARSF survey, the reasons for these discrepancies cannot be identified with any certainty. These co-ordinate differences will lead to a constant offset in attitude components which could be removed if there was an external reference source for calibration.

In this research, the GPS derived antenna co-ordinates have been used in the processing as they should produce a body frame system that will be consistent with the GPS data collected during the flights.

5.4.1.2 Determination of Approximate Attitude

The GRAPE algorithm defines a search space containing a set of trial attitudes from which ambiguities can be resolved. To centre this search space a starting approximation of the attitude must be available. Three antennas forming two independent baselines are the minimum requirement for three-dimensional attitude determination. By positioning three antennas, with known body frame co-ordinates, in a local level system, the composite rotation matrix $R_{\psi\theta\gamma}$ (3.9) can be formed, by rearranging equation (3.1) to

$$R_{\psi\theta\gamma} = X^{LLS} [X^{BFS}]^{-1} \quad (5.1)$$

Attitude components can be extracted from the rotation matrix as shown in equations 3.10 to 3.12. If more antenna positions are available a better solution can be achieved using a least squares estimation, although for an approximate solution this is not strictly necessary.

If the antennas are arranged on the body frame axes (or parallel to them), single attitude components, or combinations of two components, can be determined by computing the LLS co-ordinates of two antennas. In the current application for example, two antennas are located on the BFS y-axis³, both with near zero z co-ordinates (BFS), as shown in figure 5.6. The pitch (ψ) can therefore be computed using the height difference (δh) between the two antennas and the difference in northing (δn).

³ Antenna 2 is actually offset from the y axis by 7cm, but for a starting approximation all the height difference between antennas one and two can be attributed to pitch.

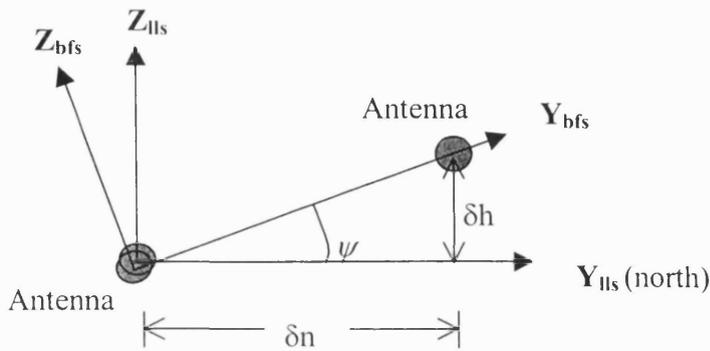


Figure 5.6: Pitch determination from a single appropriate baseline.

$$\psi = \arctan \frac{\delta h}{\delta n} \quad (5.2)$$

Using similar trigonometric relationships, and choosing appropriate baselines, all three attitude components can be approximated⁴.

In previous GRAPE tests the GPS data were gathered in a static environment allowing a single approximate attitude to be used for the computations at all epochs. For the airborne situation considered in the current research, the attitude clearly changes over time, so a single starting approximation cannot be used for any significant period. The generation of suitable starting approximations for kinematic applications is considered in section 5.5.2 and later in chapter 7, when the possibility of using an external data source to provide approximate attitude values is discussed.

⁴ BFS z co-ordinates need not both be zero or the same value. The height difference due to pitch can be isolated from any previously measured height difference in the BFS.

5.4.1.3 Generating and Reading NXF Data

GRAPE requires NXF files from at least three airborne antennas to compute attitude. The generation of NXF files from the original binary data has been discussed in section 5.2.

5.4.1.4 Determination of a Search Space

The ambiguity resolution technique used by GRAPE searches for the correct set of double difference integer ambiguities by performing computations based on a set of trial attitudes. The set of ambiguities associated with each trial attitude is made up of one integer ambiguity from each double differenced observation. For each trial attitude there are $F(L-1)(K-1)$ associated integer ambiguities, where F is the number of frequencies used, in this case two, L is the number of antennas, and K is the number of satellites. The trial attitudes lie inside a search space that is centred on the starting approximation (Figure 5.7).

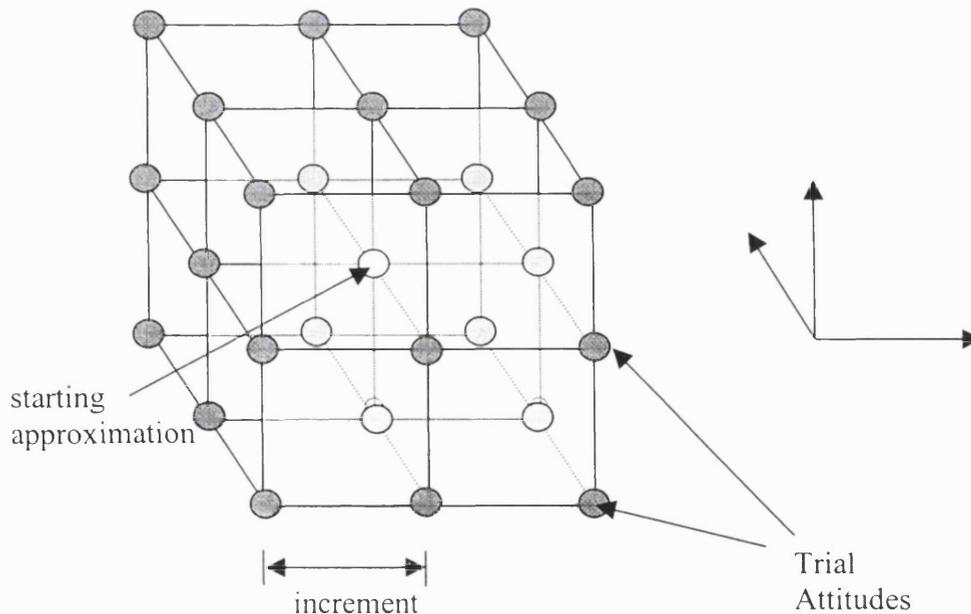


Figure 5.7: GRAPE search space

In the original version of GRAPE the search increment is fixed at 5mrad and the search size is 50mrad, for all attitude components. A search size of 50mrad means that the search space extends for 50mrad in each direction from this starting approximation, forming a search cube with three dimensions of 100 mrad, as shown in figure 5.8.

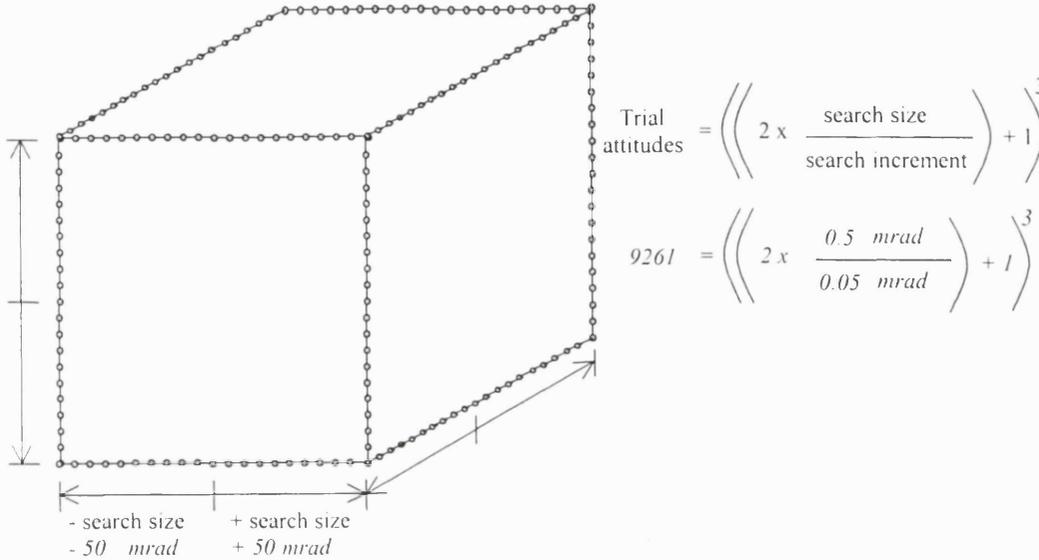


Figure 5.8: Search size and trial attitudes

A number of changes have been made to the way in which the trial attitudes are derived to make it appropriate for the current research, and future work, these are described in section 5.5.3.

5.4.1.5 Computation of the Modified Double Differenced Phase Values

The derivation of the modified double difference phase observation equation was described in section 3.5. The use of this equation allows attitude parameters to be estimated directly from GPS phase observables. This is required when computing the ambiguity function value (AFV) for each trial attitude (section 5.4.1.6), and later when computing the corrections to the candidate attitude in the least squares procedure (section 5.4.1.7). The following procedure is carried out to determine the modified double difference phase values :

- (a) Compute the composite rotation matrix $R_{\psi\theta\gamma}$, containing the nine direction cosine elements, using the trial attitude components (equation 3.9).
- (b) Compute the local level co-ordinates of all available satellites at all antennas (equation 3.5).
- (c) Compute local level co-ordinates for each antenna relative to the origin antenna.
- (d) Determine the modified double difference phase based on equation 3.30, using the local level co-ordinates of the antennas and the local level co-ordinates of all available satellites at all antennas

5.4.1.6 Computation of the Ambiguity Function Values

The ambiguity function value (AFV) for the double differenced phase residual at every trial attitude is calculated from the ‘observed minus computed’ double differenced phase (Figure 5.4). Double differences are formed for all the available independent antenna and satellite combinations using both L1 and L2 phase observations. The weighting factor applied to L2 residuals is set to 1.3 (approximately the ratio of L2 to L1 wavelength), based on the results of static tests carried out by Al-Haifi *et al.* (1998). The weighting factor increases the L2 double differenced phase residuals and subsequently reduces their contribution to the AFV.

The weighted ambiguity function algorithm may be written as:

$$AFV(\psi_i, \theta_i, \gamma_i) = \frac{1}{2(L-1)(K-1)} \sum_{l=1}^{L-1} \sum_{k=1}^{K-1} \cos 2\pi [(\nabla\Delta\Phi_{L_1})_i + (\nabla\Delta\Phi_{L_2} \times S)_i] \quad (5.3)$$

where

- $AFV(\psi_i, \theta_i, \gamma_i)$ is the ambiguity function value at trial attitude i ;
- L is the number of antennas;
- K is the number of satellites;
- $\nabla\Delta\Phi$ is the double differenced phase residual; and
- S is the weighting factor applied to L2 double differenced phase residual.

Once an AFV is computed for each trial attitude in the designated search space the maximum value is selected. All other trial attitudes that have an AFV within a specified range of this maximum are retained for the next processing step, the remainder are discarded. This range was set to be 90% of the maximum AFV for the static testing of GRAPE, based largely on experience from GASP processing. When noisier data are used, as in the current airborne example, a 90% threshold can lead to trial attitudes with the correct ambiguity combinations being discarded. The setting of this threshold for the current application will be discussed in section 5.5.

Provided that the search volume contains the true attitude, and the phase data is not too noisy, the maximum AFV should be very close to unity. In static trials the maximum AFV was consistently close to unity and only this trial attitude was retained. With noisier airborne data this is not the case, as will be discussed in section 5.5.

5.4.1.7 Least Squares Computation and Statistical Testing

A least squares adjustment⁵ is carried out for each candidate attitude that has been retained following the previous tests (section 5.4.1.6). Using this as a provisional value, the least squares routine computes corrections and residuals (Figure 5.5). The computation uses all the available L1 and L2 phase data, with a weight matrix based on the double difference mathematical correlations between the observables. Weighted least squares residuals are then computed and the solution with the minimum sum of weighted residuals is selected as the most likely solution. The Fisher test (F-test) is then used to ascertain if the solution with the minimum sum of weighted residuals can be distinguished from the other candidates. The test is carried out as follows :

The null hypothesis to be tested is

$$H_0 : \hat{\sigma}_{0(\min)}^2 = \hat{\sigma}_{0(i)}^2 \quad (5.4)$$

with the alternative hypothesis

$$H_A : \hat{\sigma}_{0(\min)}^2 < \hat{\sigma}_{0(i)}^2 \quad (5.5)$$

where

$\hat{\sigma}_{0(\min)}^2$ is the *a posteriori* unit variance for the candidate attitude with the minimum sum of weighted residuals; and
 $\hat{\sigma}_{0(i)}^2$ is the *a posteriori* unit variance for the *i*th candidate attitude, i.e. any candidate other than the one with the minimum sum of weighted residuals.

⁵ The partial derivatives needed to construct the design matrix A are given in appendix A.

The unit variance is determined as follows (Cross, 1983)

$$\hat{\sigma}_0^2 = \frac{\hat{v}^T W \hat{v}}{n - 3} \quad (5.6)$$

where

- \hat{v} is the vector of residuals;
- W is the weight matrix of a set of observations; and
- n is the number of double difference observation equations.

The test is therefore

$$F_s = \frac{\hat{\sigma}_{0(\min)}^2}{\hat{\sigma}_{0(i)}^2} < F_{v,v,1-\alpha} \quad (5.7)$$

where

- $F_{v,v,1-\alpha}$ is the critical value from the F distribution table;
- v is the number of degrees of freedom (n-3); and
- α is the level of significance (α is set to 5%).

If the null hypothesis is accepted, it can be concluded that the solution with the minimum sum of the squares of weighted residuals cannot be distinguished statistically from the other possible solutions. The attitude with the minimum sum of weighted residuals is accepted if the test fails.

Once the least squares estimation has been carried out, and a solution selected, an outlier detection routine is implemented. The statistical test procedure is the tau or τ -test, as given by Pope (1976). A τ -statistic is computed for each observation and a critical value of the τ distribution is computed. If the τ -statistic for one or more observations exceeds the critical value, the observation with the largest τ -statistic is removed from the system, and the whole least squares estimation is recomputed with the remaining data. This test is applied until no further outliers are identified, i.e. it has the ability to detect and remove multiple outliers. Full details of this procedure, including the theory and its implementation in the GRAPE software, are given in Mahmud (1999).

Some changes that have been made to the least squares adjustment are discussed in section 5.5. A weight matrix based on satellite elevations is introduced, the least squares solution is iterated, and tests for correct ambiguities based on the *a priori* baseline lengths rather than the F-test are described. The limitations of only searching for outliers in the data once ambiguities are resolved are also discussed.

5.4.2 GRAPE Performance Assessment: Static Testing

The performance of GRAPE has been tested with a static data set (Mahmud, 1999). The results from these tests have been used to assess the suitability of the algorithms, and to confirm that these algorithms have been correctly implemented within the GRAPE software. In this section, these tests are described and the systems performance is discussed in terms of ambiguity resolution success, or robustness, and accuracy. The robustness of GRAPE, using the direct attitude algorithm, is compared to that of a single baseline positioning method using GASP.

The GPS data were collected in a static mode using four dual-frequency Ashtech Z-12 receivers, each connected to a geodetic antenna. The antennas were arranged to form baselines with lengths between 20 and 40 metres. During the observation period of two hours, data from between five and nine satellites above an elevation mask of 15° were recorded at one second intervals. Using these data, a static survey, was carried out with the Trimble GPSurvey package to determine antenna phase centre positions in WGS84 co-ordinates. These values have been used to define a body frame system, and to determine the antenna positions in this fixed reference frame. The body frame was defined so that the 'true' attitude components of the array are zero in pitch and roll, with a fixed non-zero heading. The same GPS data was then processed using the GRAPE software, using the 'true' attitude values as starting approximations to centre the search space. The deviation of the computed attitude from the true attitude at each epoch is considered an error in the measurement system. The degree of variation from the fixed attitude values over the whole observation period is a measure of the accuracy of the GRAPE solution.

Ambiguity resolution was found to be 100% successful with this data set. Ambiguities were checked using the F-test and by looking for spikes in the final attitude values. Table 5.3 shows the variation of the GRAPE solution over the observation period.

	Pitch	Roll	Yaw
Maximum	2'28"	2'42"	43"
Minimum	-2.14"	-2'31"	-29"
Range	4'42"	5'13"	1'12"
Mean	6"	13"	7"
RMS	43"	49"	10"

Table 5.3: GRAPE accuracy from static testing (' minutes and " seconds of arc)

Due to the arrangement of the antenna array and the nature of the direct attitude algorithm it is not possible to attribute the accuracy of an attitude component to a given baseline. As discussed previously, the yaw component is more accurate as it is not directly affected by the height component of GPS, which remains the least well known. The mean attitude components calculated over the whole observation period show a bias of 6 to 13 arc-seconds. No precision measures from the least squares estimation have been quoted for these static tests, therefore it is not possible to assess how well GRAPE measures its own performance.

The performance of GRAPE is clearly related to the quality of the raw GPS data. During this static test, the antennas were arranged in a low multipath environment, free from obstruction, and the observation period was selected to ensure good conditions in terms of satellite availability and geometry. In such conditions, the errors in the phase observations should be minimal. Table 5.4, adapted from Mahmud (1999), gives details of the double differenced L1 and L2 phase residuals generated by GRAPE over a selected baseline.

Satellite (SV)	L1 Phase Residuals			L2 Phase Residuals		
	Max(m)	Min(m)	σ (m)	Max (m)	Min (m)	σ (m)
1	0.0097	-0.0097	0.0028	0.0193	-0.0131	0.0052
4	0.0082	-0.0127	0.0038	0.0155	-0.0089	0.0045
5	0.0102	-0.0099	0.0024	0.0112	-0.0109	0.0034
6	0.0066	-0.0071	0.0019	0.0092	-0.0104	0.0031
9	0.0083	-0.0071	0.0026	0.0093	-0.0128	0.0038
24	0.0114	-0.0082	0.0030	0.0171	-0.0134	0.0053
25	0.0110	-0.0081	0.0023	0.0156	-0.0143	0.0033
29	0.0089	-0.0076	0.0033	0.0147	-0.0137	0.0055

Table 5.4: L1 and L2 double differenced phase residuals from GRAPE

These figures indicate that the noise levels were very low, as expected. The standard deviation of the phase residual is between 1.6 and 3.8mm for L1, and between 2.4 and 5.5mm for L2. Residuals up to 19mm were recorded but in general the residuals are well below 10mm, and are usually less than 5mm. Hofmann-Wellenhof *et al.* (1994) indicate that the maximum error in raw phase data due to multipathing effects is of the order of a quarter of a wavelength (or 50mm on L1). Clearly then, this data has been recorded in a low multipath environment.

One of the reasons for developing the direct attitude algorithm was to increase the redundancy in the processing procedures. A single baseline system, such as GASP, determines three unknown position components per baseline, using double difference observations across that baseline. In the direct attitude system there are three unknown attitude components for the complete array, but observations from three independent baselines can be used in the calculation. In tests of the single epoch positioning algorithm, additional satellite observations have significantly improved the success of GASP ambiguity resolution (Corbett, 1994; Corbett and Cross, 1995; Al-Haifi, 1996; and Al-Haifi *et al.*, 1998). Therefore, it was expected that the greater redundancy in the GRAPE direct attitude algorithm would lead to significantly better performance, in terms of ambiguity resolution, than the single baseline positioning algorithm used by GASP.

To test this hypothesis, GPS data from the same test were processed using GASP to determine a set of baselines and therefore establish relative antenna positions. Each baseline was processed separately, i.e. one unknown receiver position was determined relative to one reference station. Results showed that GASP could successfully resolve ambiguities across each baseline for all epochs, just as GRAPE had done. It is possible therefore to derive three dimensional attitude by processing these baselines together with the *a priori* body frame co-ordinates. To discover whether or not GRAPE ambiguity resolution was demonstrably more robust than GASP using this data set, additional errors were added to the original phase observations. The performance of the direct attitude method, GRAPE, and single baseline method, GASP, were then compared again.

Initially, a random error with a standard deviation of 5mm was added to the L1 phase observations from a single satellite at a single antenna. The data was re-processed with the direct attitude (GRAPE) and the single baseline (GASP) software. The additional error was then increased to 10mm (σ) and again the data was re-processed. The direct attitude system was still able to correctly resolve ambiguities at all epochs, with an additional 5 or 10mm phase error, i.e. 100% success. The single baseline system failed to resolve ambiguities at 64 of the total 7,200 epochs with the 5mm error, and 175 epochs with the 10mm error, this represents a success rate of 99.1% and 97.6% respectively. This demonstrates that the additional observations in the direct attitude approach produce a more robust solution than a single baseline system.

It was also predicted that the additional redundancy of the direct attitude algorithm would allow more successful outlier detection, i.e. smaller errors could be detected more of the time. To test this hypothesis, the outlier rejection routine developed for GRAPE (section 5.4.1.7) was also implemented in GASP. A comparative test was conducted in a similar way to the ambiguity resolution test described previously, except that the additional phase error was not added until the ambiguities had been successfully resolved. A 5mm error was added to the observed phase from one satellite (SV24) recorded at one antenna. The data was processed and the number of epochs at which the biased observations were

rejected was recorded. The additional error was increased incrementally by 5mm until either the direct attitude system or the single baseline system managed to reject the observations in which the constant error was included. This process was then repeated for SV25 to see if the results differed significantly. Table 5.5 shows the number of epochs at which observations containing outliers were successfully detected and rejected by the two approaches as the introduced error was increased. Figures in brackets express the number of epochs rejected as a percentage of the total number of epochs tested.

Added Error (m)	Direct Attitude System		Single Baseline System	
	SV24	SV25	SV24	SV25
0.005	4802 (66%)	6718 (92%)	4327 (59%)	5044 (69%)
0.010	6824 (93%)	7267 (99.8%)	6425 (88%)	6290 (86%)
0.015	7209 (98%)	7281 (100%)	7002 (96%)	6837 (94%)
0.020	7318 (99.96%)	-	7235 (99%)	-
0.025	7321 (100%)	-	7167 (98%)	-

Table 5.5: Outlier rejection comparison

The results show that the direct attitude system is consistently able to detect outliers at more epochs than the single baseline system, this is due to the additional redundancy, which leads to more accurate and reliable attitude solutions.

These tests have proved that the direct attitude system is more robust than a single baseline system in terms of its ability to resolve ambiguities, and that it is more reliable in terms of its ability to reject outliers following the least squares computation. Further details on the performance of GRAPE are given in Mahmud (1999). These results have shown the potential of the direct attitude algorithm, particularly as the levels of noise in the phase measurements increase.

5.5 Kinematic GPS Attitude Determination: Extensions to GRAPE

5.5.1 Overview

The previous section described the original GRAPE software in terms of the algorithms it employs and its performance using static data. In this section some modifications to the original algorithms are described, including extensions that have been developed specifically to allow the processing of the airborne GPS data collected for the current research.

The following discussion is limited to general changes in the program design and algorithms, but it should be noted that many other changes of a more technical nature have been made to the original program. For example, numerous changes have been made to the processing parameters that the user specifies, and the format of the output files, primarily so that GRAPE output can be used as an input source for the Kalman filter described in chapter 7. Changes have also been made to the RINtoNXF, GASP and GRAPE programs to allow data recorded at a rate higher than 1Hz to be processed.

Of the more general changes described in this section, the first two relate to the construction of the ambiguity search space (sections 5.5.2 and 5.5.3). The third change (section 5.5.4), iterating the least squares estimation, affects the ambiguity resolution and the accuracy of the final solution. The fourth, the introduction of satellite elevation dependent weighting, should improve the solution from the least squares estimation by assigning more appropriate weights to the observations (section 5.5.5). Finally a moving local level system origin (section 5.5.6) is introduced to keep the co-ordinate system origins coincident in a dynamic environment. In section 5.5.7 a number of changes that have been considered, but not implemented, are described.

5.5.2 Attitude Approximation to Centre Search Volume

In section 5.4 the need to derive an approximate starting attitude on which to centre the search volume, was discussed. In a static environment a single value can be used throughout the processing period. In a dynamic environment however, this starting approximation must be updated for each epoch as the attitude of the vehicle changes. A number of methods, with varying degrees of complexity, could be used to compute a new value for each epoch. Some examples are:

- 1) Use the final attitude computed at the previous epoch to centre the search space at the current epoch. This is sufficient provided the attitude is not changing significantly between epochs.
- 2) Predict a starting approximation based on previous measurements and the predicted behaviour of the vehicle. This can be achieved by calculating a rate of change of attitude from the previous two epochs, and multiplying this by the recording interval to compute the attitude change from the previous epoch to the current epoch. The formula for this is simply:

$$x_i = x_{i-1} + \dot{x} * \delta t \quad (5.8)$$

where,

- x_i is the value of any attitude component at epoch i ,
- x_{i-1} is the value of the component at the previous epoch $i-1$,
- \dot{x} is the rate of change of the attitude component, from $\dot{x} = x_{i-1} - x_{i-2}$, and,
- δt is the recording interval.

The accuracy of the predicted value can be improved by implementing a Kalman filter so that each new starting approximation is based on all the measurements to date, and a knowledge of the vehicle's behaviour over time.

- 3) An external source of attitude data can be used to provide a starting approximation, provided any spatial and temporal offsets between instruments are determined. In chapter 7, results from basing the starting value on the measurements from a gyro are described.
- 4) A nominal 'zero' value can be used at every epoch, with a search volume designed so that it includes all possible attitude ranges that the vehicle could experience. In the case of an aircraft this may mean the search volume needs to extend from -60° to 60° in pitch and roll, and from 0 to 360° in heading, this would introduce a huge computational load.

The choice of method will depend on the frequency and magnitude of the measured attitude variations and the recording rate of the GPS receiver, as these factors dictate how much the attitude changes between epochs. For the current airborne data set, the attitude does not change significantly between recording epochs (0.1 seconds) during flight lines, therefore using the computed attitude from the previous epoch, option 1, has proved successful.

Any value that is to be used as a starting approximation must undergo some form of quality check before it is accepted. A poor starting value may lead to a search space being created that does not include trial attitudes leading to the correct ambiguity combination, except in the case of an 'all encompassing' search volume, as described in option 4. If this occurs at any one epoch, then it is likely to occur at all the following epochs. The quality check can be based on some form of reliability or accuracy measure associated with the proposed value, e.g. a standard deviation. If no such quality measure is available then rejection criteria based on the maximum attitude changes that can take place between epochs should be implemented. Currently, the AFV of the solution at the previous epoch is compared to a threshold value before it is used as starting approximation.

A simple approach, using the computed attitude at the previous epoch, has proved to be successful for the current application. An option to use an external data source will be discussed in chapter 7. If, in a future application, the GRAPE software was used as the sole means of attitude determination and the data collection conditions changed, the procedures may need to be altered.

5.5.3 Search Increment and Size

Once the search volume is centred, the search size and increment must be defined to create the set of trial attitudes (see section 5.4.1.4). In the GASP positioning software, searching is carried out in position and ambiguity space to ensure that every trial candidate represents a unique set of ambiguity values. With GRAPE this is not the case, here the search takes place only in attitude space, so each trial attitude does not necessarily represent a unique set of integer ambiguities. A single increment in one attitude component may not lead to a single change in the set of ambiguities. If none of the ambiguity values change from one trial attitude to the next the computation is inefficient, if however more than one ambiguity changes between adjacent trial attitudes the correct set of ambiguities may be missed. For the processing to be successful, a trial attitude that produces the correct double differenced ambiguity values must be included in the search volume. To ensure such a trial attitude is included, the starting approximation, the search increment, and the search size must all be appropriate.

To improve the processing efficiency, the search size should be related to the quality of the starting approximation. If the starting approximation is good then a relatively small search space can be created. In GASP, a double differenced code solution is used to provide a starting position and the standard deviation associated with this position is then used to determine the search size. The original GRAPE program used a fixed search size, which was chosen based on experience of processing the specific static data set. In the current application, it is difficult to find a good indicator of the quality of the starting approximation so again the search size has been fixed, based on experience. For the current research, variables have been introduced so that the search size and increment can

be chosen separately for each attitude component. The dimensions of the search space can therefore be varied to reflect the potential accuracy of the starting values. For example, the heading component tends to change in a smoother manner than pitch and roll. As the heading changes less between epochs, the starting approximation is likely to be better, therefore this dimension of the search space can be reduced. Any reduction in the number of trial attitudes to be considered decreases the processing time.

The search increment for GRAPE was originally fixed at 5 mrad based on tests with the static data set. This figure was chosen as it produced only one trial attitude with the correct set of integer ambiguities. However, it was later discovered that this increment leads to many changes in ambiguities between adjacent trial attitudes in the search space and therefore misses many possible ambiguity combinations. The correct ambiguities were found at each epoch because the starting approximation on which the search space was centred led directly to the correct ambiguity combination. In the current research a number of tests were carried out to find an appropriate search increment.

The baselines on the aircraft are much smaller than those formed in the static tests, 1.2 to 3.5m rather than 20 to 40m. Therefore a larger change in attitude is required before any integer ambiguities will change. In other words, a larger increment size can be used without missing a set of ambiguities. Although variables have been introduced so that different increments can be used for different attitude components, for convenience, the same increment was applied for all components in these initial tests. Tests were carried out with varying increments (0.005 to 0.1 radians), search sizes⁶ (0.02 to 0.1 radians) and observation conditions (different numbers of available satellites and antennas). The aim of this investigation was to determine the largest increment that could be used without missing possible ambiguity sets. A larger increment produces fewer trial attitudes for a fixed search size, or if the total number of trial attitudes tested is held fixed, a larger increment allows a greater range of attitude values and

⁶ A search size of 0.02 radians produces a search cube with dimensions of 0.04 radians, as described in section 5.4.1.4.

different ambiguity combinations to be tested. It also avoids repetition in the processing, which occurs when multiple trial attitudes produce the same set of ambiguities. However, it is more important that the correct set of ambiguities is not missed, even though this leads to repetition of processing.

In this investigation, the complete set of double differenced ambiguities derived from each trial attitude at a single epoch were stored. Once the search was complete, the ambiguity sets were compared to establish how many different sets were considered. Figures 5.9 and 5.10 both show the percentage of the total ambiguity sets⁷ that are tested within a 0.1 radian search space, as the search increment changes. These tests were then repeated for different observation conditions and the number of double differenced observations used is given for each line. Figure 5.9 shows the trend over a larger range of increments and therefore includes conditions in which very few of the total ambiguity sets are tested. Figure 5.10 uses the same data but focuses on the increment size required to test all, or nearly all, of the possible ambiguity combinations.

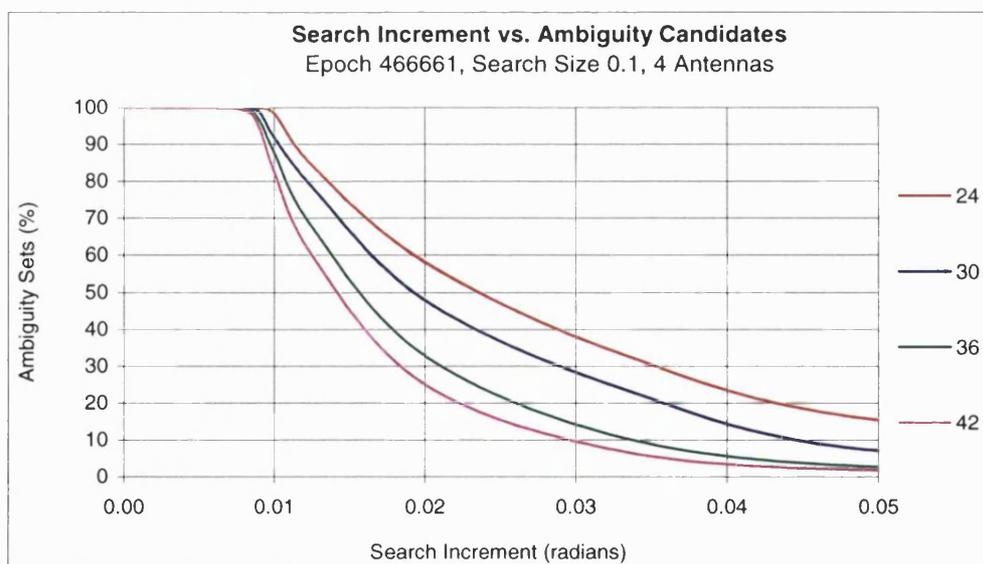


Figure 5.9: GRAPE search increment

⁷ The total number of ambiguity sets has been determined by reducing the increment until no new combinations are found.

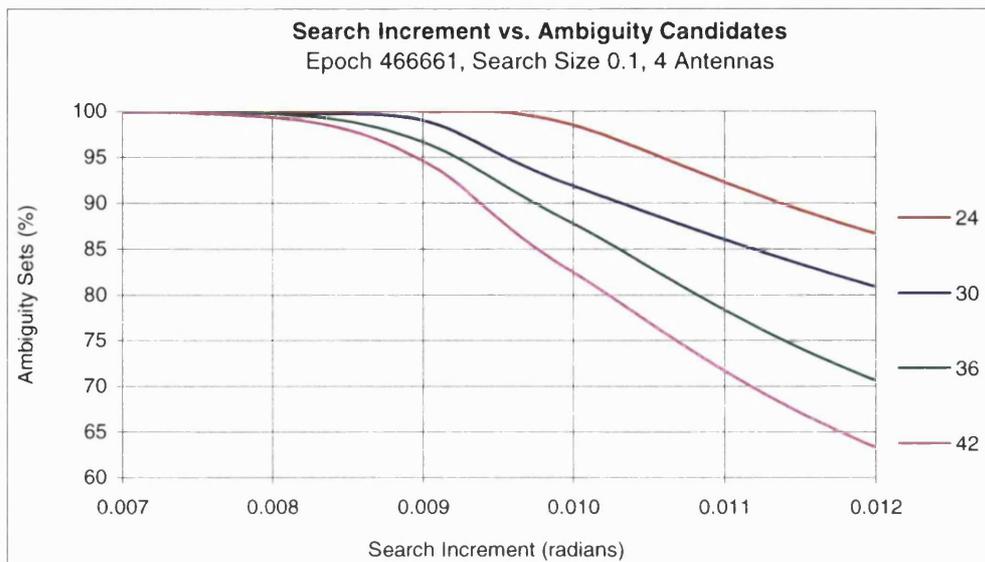


Figure 5.10: GRAPE search increment (selected values)

With fewer observations used in the solution, the total number of ambiguity sets considered in the search space is smaller and a greater change in attitude is needed to induce a change in the ambiguities, hence a larger increment can be used. With a larger set of observations, and hence unknown ambiguities, there are a greater number of different ambiguity combinations in the same search volume and a smaller change in attitude will change the ambiguities, hence the increment size must be smaller. With 42 double differenced observations there are 1075 different ambiguities within the search space at the selected epoch. With only 24 observations, 205 different ambiguity sets are tested in total.

To be sure that no possible ambiguity sets are missed in the search, an increment no larger than 0.007 radians should be chosen. This is adequate for eight satellites observed at four antennas (42 double differences, using L1 and L2). If any more than eight satellites are tracked, or more than four antennas are used, this increment must be reduced accordingly. The search increment can be scaled automatically within the program based on the number of observations used.

For the current application, a search increment of 0.01 radians has been used when data is only available from three antennas, and increments from 0.005 to 0.008 radians have been tested with four antennas. To reduce the computational load, an increment that is sufficient to test most but not all ambiguity combinations could be used. It is possible to check whether the correct ambiguities have been resolved using *a priori* knowledge of the antenna separations (see section 6.4.3). If the correct ambiguity set has been found, then this increment can be retained.

The processing time of the GRAPE program is a function of the number of trial attitudes that have to be tested, which in turn is determined by the search size and increment (see section 5.4.1.4). Using a 0.008 radian increment over a 0.1 radian search space produces 19,683 trial attitudes. These trial attitudes only lead to 205 and 1075 different ambiguity combinations using 24 and 42 observations respectively. If the search increment can be increased to 0.01 radians for example, the number of trial attitudes reduces to 9,261.

The routines within GRAPE to isolate the correct ambiguity set rely on comparing the AFV, and later the sum of the squares of the residuals, produced by different trial attitudes. Clearly, if the different trial attitudes are producing the same set of ambiguities these procedures are not actually comparing alternative ambiguity sets. In the original GRAPE software, all trial attitudes producing an AFV within 90% of the maximum AFV were retained for the least squares estimation. This can lead to a large number of trial attitudes that produce the same ambiguity combination being retained. In the course of this research, a ‘uniqueness’ criterion has been introduced to select only one trial attitude associated with any ambiguity set. This removes a great deal of redundant processing in the least squares estimation. Any tests to isolate the correct ambiguity sets are then being carried out on genuinely different sets of ambiguities.

5.5.4 Iterating the Least Squares Estimation

It was discovered that when multiple trial attitudes leading to a common set of integer ambiguities were used as provisional values in the least squares estimation, they did not produce the same final attitude, as should be the case. Once the routine to isolate and retain only a single trial attitude for any ambiguity combination was introduced, a different answer could be produced depending on the method for choosing the one representative trial attitude per ambiguity set, i.e. again, the same ambiguities are leading to different final solutions. Investigations showed that this was due to the absence of any iteration in the least squares estimation. No iteration was used in the original design because previous GASP position tests had shown that a single run of the adjustment was sufficient. In any further iterations the changes to the corrections became negligible. With the current airborne data, corrections computed after the first loop can be as large as 5 to 10 arc-minutes. An iterative procedure has now been introduced so that corrections are computed until the change between iterations is less than 1", this does not of course imply that the accuracy of the final solution is at this level. This condition is usually satisfied after 4 or 5 iterations.

Another advantage of iterating the least squares solution, is that it can aid the selection of the correct ambiguity set. If trial attitudes which initially lead to an incorrect ambiguity set are used as provisional values in the least squares estimation, the addition of iterative corrections can change the attitude until its associated ambiguity set becomes the correct set. At the end of the iteration the ambiguities from different trial attitudes may have become identical and the final solutions will then converge to the same value. If incorrect ambiguities are retained throughout the iterations, the final solution from this trial attitude will generally have a significantly larger sum of the squares of the weighted residuals ($\hat{v}^T W \hat{v}$), allowing this alternative solution to be rejected, using the F-test described in section 5.4.1.7. This could allow the previous routine to select possible trial attitudes to be less rigorous, as any trial attitudes producing incorrect ambiguity combinations will either be corrected in the least squares estimation until the correct ambiguity set is found, or rejected on the basis of the associated residuals.

It might be possible to extend this idea further by replacing the ambiguity function method with the least squares ambiguity search technique, LSAST (Hatch, 1989 & 1990) or to introduce some hybrid of the two. The advantage of using LSAST, in preference to any of the other commonly used 'on-the-fly' ambiguity resolution techniques, is that the computation time decreases as the number of observations increases. As the algorithms implemented in GRAPE have been developed to maximise the number of available observations, LSAST could be an appropriate technique. Hatch (1990) indicates that instantaneous ambiguity resolution is possible with LSAST using dual frequency data from seven or more satellites over baselines of less than 5km. However, if only five satellites are available, ambiguity resolution may take up to five minutes, especially under conditions of poor geometry. Descriptions and comparisons of commonly used ambiguity resolution algorithms are given in Erickson (1992), Abidin (1993), Corbett (1994) and Al-Haifi *et al.* (1998).

This investigation of a hybrid approach was not pursued further because results have shown that GRAPE can successfully resolve ambiguities during flight lines using the current algorithms. In this particular application it was also decided that any developments to improve the success and efficiency of ambiguity resolution should concentrate on the integration of gyro data with GPS, rather than developing stand-alone GPS solutions.

5.5.5 Satellite Elevation Dependent Weighting

The original GRAPE program employed a simple weighting strategy in the least squares estimation. Equal weights were assigned to all L1 observations, with L2 observations down-weighted by a factor of 1.3 relative to L1. It is known however, that GPS observations from low elevation satellites are subject to substantially more noise than those from higher elevation satellites. The elevation-angle dependence of a measurement's noise is due mainly to the antenna's gain pattern, with other factors such as atmospheric signal attenuation contributing to a lesser degree (Tiberius *et al.*, 1999).

To reflect this, a weighting function based on $1/\sin(\text{elevation})$ has been introduced. Figure 5.11 shows how observations from lower elevation satellites are down-weighted. From 0° to 5° equal weights are assigned rather than continuing the function, otherwise there is little point including very low elevation satellites in the solution at all as they will be given so little weight that they do not affect the final result. It is assumed, that if the elevation mask is set at 0° rather than 5° then the user has made a decision to include satellites at such low elevations. L2 observations are still down-weighted relative to L1 by 1.3.

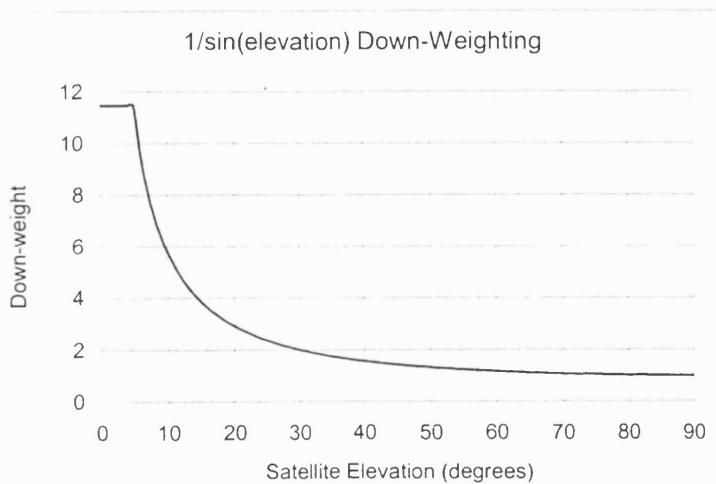


Figure 5.11: Satellite elevation dependent weighting function

5.5.6 Moving Local Level System Origin

When attitude was defined in chapter 3 as the rotations of a fixed body frame with respect to a local level system, it was noted that the origin of the two systems should be coincident. For convenience, the phase centre of antenna 1 is used to define the origin of both systems. In the original version of GRAPE, designed to process static data, a single position of antenna 1 from the NXF file header is used throughout the processing period. In the current dynamic application, an updated antenna 1 position is required for each epoch. As part of the NXF file structure an approximate antenna position is given for each observation epoch, based on a single point code solution. For dynamic data this position is used as the local level system origin.

5.5.7 Possible Future Changes

In this section some further investigations of the original GRAPE algorithms are described. The changes considered have not been implemented because GRAPE is working successfully with the airborne data without these additional refinements, and some of the suggested improvements can be achieved using alternative methods based on the integration of GPS data with an external data source. Exploiting the benefits of an integrated GPS and gyro attitude system in this way, is one of the aims of the current research project. If it was the intention to use GPS as the only source of data in future applications, and observation conditions were likely to be poorer than those experienced during the flight lines in this project, then some of these issues could be explored further.

5.5.7.1 Outlier Detection

One of the primary reasons for implementing a direct attitude determination algorithm within GRAPE was to increase redundancy in the system to allow more successful outlier detection, i.e. to improve reliability. It was predicted that the performance of the attitude determination system would be significantly improved if it was able to identify specific observations contaminated by high levels of multipath and remove them from the solution. In section 3.6 three general problems that multipath can cause were described, these were:

- (1) Ambiguities not being fixed
- (2) Incorrect ambiguities being selected
- (3) Degraded accuracy in the final solution
(with correct ambiguities)

Due to the processing methods used in GRAPE and GASP, the first situation cannot occur, a set of ambiguities will always be chosen whether they are correct or not.

The outlier detection routine implemented in GRAPE was described briefly in section 5.4.1.7 and some results of its use were given in section 5.4.4, a complete description of the algorithm and its performance is given in Mahmud, 1999. The test for outlying data is performed once the least squares estimation is complete, using the computed unit variance. If an outlier is identified, this observation is removed from the solution and the attitude is recomputed. This method can therefore address problem (3) to some extent: once ambiguities are successfully resolved it can identify and remove noisy observations that would degrade the accuracy of the final solution. This outlier detection routine does not address problem (2), the selection of incorrect ambiguities. In the investigations described, the additional errors that were introduced to the observations when testing the outlier detection routine were only introduced once ambiguities had been fixed. In practice, if a phase observation is significantly contaminated by multipath it may prevent successful ambiguity resolution. If this is the case, then

any subsequent tests that rely on successful ambiguity resolution are clearly of little use. The original outlier detection routine would only be of benefit if levels of multipath are sufficiently low to allow successful ambiguity resolution. It may then have some affect on the accuracy of the final solution. However, if noise levels are low enough to allow successful ambiguity resolution, any undetected poor data distributed amongst the full set of observations in the least squares estimation are unlikely to have a significant impact on the final result.

To improve the ambiguity resolution in the presence of multipath, outlier detection needs to take place at an earlier stage in the processing. One possibility would be to implement a check on the data during the computation of ambiguity function values. Inspection of the formula for computing the AFV for a trial attitude (equation 5.3), shows that the final value is based on a summation of the cosines of the phase residuals from each double differenced observation. In the standard Ambiguity Function Method, each individual cosine is tested against a minimum threshold value, for example Mader (1990) uses a 0.7 threshold. If the cosine from a residual falls below this value, the remaining computations are skipped and the algorithm proceeds to the next trial position. This procedure greatly improves the efficiency of the algorithm (Mader, 1992). The introduction of this step was investigated in the design of GASP, but has not been implemented. No attempt was made to introduce this in GRAPE. In both GRAPE and GASP, any threshold tests are carried out once all individual cosines have been summed to compute the final ambiguity function value for a trial attitude or position.

This check on individual cosine values could be implemented for the original purpose of improving processing efficiency, but it could also be adapted to act as an outlier detection routine. Instead of moving on to the next trial attitude when the cosine of a phase residual falls below the threshold level, it could be stored and the remaining calculations at the trial attitude could proceed. If the cosine value from an observation is significantly below the value of the cosines from other observations, this suggests there may be an outlier in that specific double difference observation. If further double difference observations using data from

the same satellite antenna combination show a similar trend, then all observations from that antenna-satellite combination could be removed from the solution. For example, in GRAPE observations from each satellite to antenna 2 are used to form double difference observations on baselines antenna 1 to antenna 2, and antenna 2 to antenna 3. By re-computing without these poor observations, the ambiguity function value at a trial attitude can become a better indicator of whether or not the correct ambiguities have been found. The amount by which the AFV produced by a trial attitude leading to the correct ambiguities exceeds the AFV of a trial attitude producing an incorrect ambiguity set, should increase.

Such a procedure needs to be implemented with caution as phase observations may not fit the calculated solution very well (i.e. they produce large residuals) simply because the trial attitude does not produce the correct ambiguities. In the GRAPE program, where there are a larger number of observations, it is possible to identify a single poor observation. For example, if measurements are recorded from seven satellites by four dual-frequency GPS receivers, there are a total of 36 independent double differenced observations. An outlier in the observations from a single satellite observed at a particular antenna would be used in only one or two of the 36 double differenced observations⁸. Some preliminary tests were carried out in which observations were rejected on the basis of the cosine of the phase residual. After re-computing without these observations, the ambiguity function values of trial attitudes producing correct ambiguities could be differentiated more easily from those associated with incorrect ambiguity sets.

A far less efficient method to try and remove a single phase observation contaminated by multipath, is to systematically remove each observation in turn and re-compute the solution. If the solution improves significantly, in terms of the ability to differentiate trial attitudes producing the correct ambiguities, when one observation is removed, it can be removed for each trial attitude in the search space. To identify outliers in more than one observation would require a large number of computations to consider each possible combination.

⁸ Observations from the differencing satellite are used in all double differences on a baseline, but as this is the highest elevation satellite it is least likely to produce very noisy phase observations.

An outlier detection routine before or at the ambiguity resolution stage has not yet been implemented as a standard step in GRAPE. This is because the current algorithms are adequate for the airborne data used in this project. Also ambiguity resolution is assisted by the use of external attitude data and the iteration of the least squares estimation.

5.5.7.2 Validating Ambiguity Combinations

To assess GRAPE's performance, it must be possible to establish with a satisfactory degree of certainty whether or not the chosen ambiguities are correct. In the standard GRAPE attitude determination program, the same statistical tests that are used in the GASP positioning program have been adopted. This procedure is described in detail in section 5.4.1.7., for convenience, it is repeated below as three simplified steps:

- 1) During the least squares estimation the sum of the squares of the weighted residuals ($\hat{v}^T W \hat{v}$) is computed for each trial attitude retained following the initial ambiguity search.
- 2) The solution with the smallest residuals ($\hat{v}^T W \hat{v}$) is selected as the most likely.
- 3) An F-test is then used to ascertain if this solution can be distinguished from the others, if it is, then it is considered to be correct, at the specified level of significance.

This procedure has been shown in the past to be a relatively good method for establishing whether or not the chosen ambiguities are correct. However, in applications where the GPS data is relatively noisy, or where there are a limited number of observations, it becomes increasingly difficult to confirm if the solution with the minimum $\hat{v}^T W \hat{v}$ is correct. In these circumstances the F-test may find that a single solution may not be statistically better than one or more alternative solutions.

In the current application, the four aircraft antennas are rigidly fixed and the baseline lengths between them are known. Possible double differenced ambiguities across any baseline used by GRAPE can therefore be tested by using these lengths when positioning one antenna relative to another. If the baselines computed after applying these ambiguities agree to within a few centimetres with the *a priori* baseline lengths, then it can be concluded that the ambiguities have been successfully resolved. The procedure can be summarised as follows:

- 1) Process the data from four antennas using GRAPE to determine possible sets of ambiguity values. For each possible set, store the ambiguities for each double difference observation, on each of the three independent baselines.
- 2) Using GASP, treat one antenna as a known reference and position a second antenna relative to it using the appropriate ambiguity values from GRAPE.
- 3) Repeat step (2) for the remaining baselines processed by GRAPE.
- 4) Compare the baseline lengths computed from GASP with the known *a priori* values. If they agree to within a few centimetres, GASP can be considered to have used the correct ambiguities. This then confirms that the GRAPE ambiguity resolution has been successful.

To allow this form of ambiguity validation to be performed in an automated way, some of the positioning algorithms from GASP need to be integrated with the GRAPE package. This would allow a 'built-in' check on the final ambiguities using the known *a priori* baselines. This would not greatly increase the processing burden as many of the procedures in the two programs are common - note the similarity of figures 5.2 and 5.3, the processing overviews. Also, the current GRAPE ambiguity search, which works in attitude space only, could be made more efficient by combining it with a GASP search initially. Defining a search space containing only trial attitudes that lead to different ambiguities would make the processing significantly more efficient.

5.6 Concluding Remarks

This chapter has described the algorithms and software used to derive position and attitude information from the GPS data collected during the two flight trials. An overview of the entire processing strategy was given initially, followed by a more detailed discussion on specific aspects of position and attitude determination.

Two alternative software packages have been used for determining positions. The first, GASP, resolves ambiguities from a single epoch of data using algorithms based on the ambiguity function method. The second, SKI, is a widely used suite of GPS post-processing software that implements the fast ambiguity resolution approach (FARA) to resolve ambiguities on a multiple epoch basis.

Attitude values are determined using GRAPE, a software package that has evolved from GASP and implements the direct attitude determination algorithm. This software and more particularly its processing algorithms have been described in some detail, and its performance, as assessed in previous research has been summarised.

Refinements and additions to the original GRAPE program, that have been undertaken as part of the current research project, were then described. These changes have been made to improve the existing ambiguity resolution procedure and to allow kinematic data to be processed.

- A routine has been introduced to centre the ambiguity search space at an appropriate starting approximation for each epoch.
- Investigations have determined the search increment and search size that must be used to ensure the correct set of ambiguities will be included amongst the candidate sets. The increment size is now computed automatically depending on the number of double differenced observations in the solution. A procedure has also been added to ensure that multiple trial attitudes producing the same ambiguities are not passed from the ambiguity search to the least squares estimation.

- Introducing iteration to the least squares adjustment leads to a more accurate final solution, and can be used to identify the correct set of ambiguities.
- A satellite elevation dependent stochastic model has been introduced.
- To keep the origins of the body frame system and the local level system coincident, a new procedure has been implemented so that the origin antenna position is updated at each epoch.
- Some additional processing parameters and output options have been introduced.

A number of other issues, mainly relating to ambiguity resolution, have been investigated, but the program has not been changed as a result. The performance assessment in the following chapter explains why any further refinements were not considered necessary for the current research.

Chapter 6 Position and Attitude Results from GPS Flight Data

6.1 Introduction

This chapter presents position and attitude results that have been determined using the methods described in chapter 5, from GPS data recorded by the four independent airborne receivers and the primary reference station. To put the following results in context, some details of the observation conditions during the test flights are given initially, and the availability of GPS data from each of the independent receivers is discussed (section 6.2). Position and attitude results are then presented and analysed to make an assessment of the system's performance in terms of robustness (i.e. the ability to correctly resolve ambiguities) and accuracy.

Initially it was intended that the complete georeferencing solution would be assessed in terms of the final on-the-ground positioning accuracy. The ground co-ordinates of a number of targets in the survey area were determined prior to the flight, as described in section 4.3.3. The accuracy of the system can be assessed by comparing the co-ordinates of these targets derived from directly georeferenced scanner imagery, with these reference values. The resulting differences can be attributed to the accuracy with which the position and attitude of the scanner was determined, and the image stability. As none of the scanner imagery recorded on the two test flights is available for analysis this complete calibration is not possible.

The attitude solution could also have been calibrated using a reference solution from the dedicated GPS attitude determination unit, the ADU2, which was operated in parallel to the four independent receivers during the flights. Although a comparison of this kind cannot calibrate the accuracy of the independent receiver solution beyond the accuracy level of the ADU2, this degree of calibration would have been adequate to assess whether or not the attitude

accuracy requirements were being met. As none of the ADU2 data recorded during the flight has been made available, this calibration is also not possible.

As neither of these methods of calibration can be used, some alternative less comprehensive methods have been employed. The details of these methods for assessing the system's robustness and accuracy in terms of position and attitude are described, and the results from applying these methods are discussed.

Section 6.3 presents an assessment of the positioning solutions determined by GASP and SKI. In section 6.4 a similar assessment is made of the attitude results using data from the independent GPS receivers. This data has been processed using GRAPE software, which implements the direct attitude algorithm. Section 6.5 discusses the issue of recording rates in the context of position and attitude. Some concluding remarks are made in section 6.6, principally to summarise the key findings of the performance assessment.

6.2 Observation Conditions

This section describes how the data that were eventually processed and analysed, were selected from the complete data-sets gathered during the two test flights. It also describes the observation conditions in which the data were collected in terms of satellite availability and geometry, and the completeness of the recorded information.

The first test flight on day 288 (15 October 1998) lasted 70 minutes and recorded 10 flight lines of two to three minutes each. The second test flight, carried out on the following day, lasted 60 minutes and recorded 7 flight lines of similar length. A number of criteria have been applied to select appropriate sections of data, these are as follows:

- 1) Imagery is only recorded on flight lines, hence it is only necessary to provide exterior orientation parameters during these periods to fulfil the aims of this project.

- 2) The chosen flight lines should represent typical observation conditions.
- 3) Flight lines in which surveyed ground targets are imaged should allow an assessment of on-the-ground positioning accuracy. These lines should therefore be processed in preference to other flight lines on which no surveyed features are available for calibration¹.
- 4) GPS data must be available from a least three of the four GPS antennas to achieve a full attitude solution.
- 5) AHRS data, which will provide reference attitude values, is only recorded during flight lines. A further reason for limiting any analysis to flight lines.

As a result, six flight lines (four from the first flight and two from the second) during which sufficient data are available to allow a full attitude solution, have been selected for processing. Details of these flight lines are given in table 6.1. Each flight line is described in terms of its duration, the number of receivers on which GPS data were recorded, an approximate heading, and satellite observation conditions (number of available satellites, SVs, and geometry, PDOP, using two elevation masks). Two points to note are:

- 1) The durations of the flight lines are shorter than was originally intended, 50 to 80 seconds instead of 120 to 180 seconds. This is because GPS data has generally only been processed for periods when AHRS data is available and this does not cover full flight lines.
- 2) On day 288, there is no airborne data available from antenna 3 (the port winglet antenna), due to logging problems with the Ashtech Z-Sensor and PC configuration. As a result, a full attitude solution is not available at 10Hz as data from the 5Hz Ashtech Z-12 receiver, connected to antenna 4, must be used in the solution.

¹ As no scanner imagery has been available for calibration this consideration was ultimately not relevant, however the selection of data was made whilst the image data was still expected.

Line Number	Day	Start Time GPS seconds of week <i>hh:mm:ss</i>	End Time GPS seconds of week <i>hh:mm:ss</i>	GPS Data Availability	Heading	Observation Conditions			
						10°		20°	
						SVs	PDOP	SVs	PDOP
L001	288	388045 <i>11:47:25</i>	388115 <i>11:48:35</i>	3 antennas	West	9	2.1	8	3.0
L003	288	388680 <i>11:58:00</i>	388730 <i>11:58:50</i>	3 antennas	West	8	2.9	8	2.9
L006	288	389470 <i>12:11:10</i>	389550 <i>12:12:30</i>	3 antennas	North	8	2.7	7	3.5
L008	288	389760 <i>12:16:00</i>	389820 <i>12:17:00</i>	3 antennas	North	8	2.7	7	3.2
L013	289	467770 <i>9:56:10</i>	467830 <i>9:57:10</i>	4 antennas	West	8	1.8	5	4.3
L015	289	468530 <i>10:08:50</i>	468600 <i>10:10:00</i>	4 antennas	West	7	2.0	6	2.2

Table 6.1: Processed flight lines

Appendix B gives a more detailed description of the GPS data processed in this project. The quality and completeness of the recorded data is a function both of the observation conditions and the logging problems experienced during the flight, as noted in section 4.3.4. The amount of available GPS data varies from receiver to receiver. The Z-Surveyor connected to antenna 2, recorded GPS data to an internal memory card, this data set has no gaps. All the files from the three remaining receivers have some data gaps. The completeness of the data recorded by each aircraft receiver clearly dictates the number of epochs at which position and attitude solutions can be computed.

Gaps in these data sets are due entirely to logging problems and not to complete losses of lock on all tracked satellites during aircraft manoeuvres. Where any losses of lock do occur due to manoeuvres, they are generally restricted to one or two satellites, not complete outages on all satellites. To maintain normal operating conditions the pilot made no specific attempts to limit the degree of banking during turns to minimise any possible discontinuities.

Data gaps vary in nature from short frequent gaps of one or two epochs, to longer gaps of thirty or forty seconds. If longer gaps occur during a flight line on more than one receiver then it is generally not possible to produce a usable solution for that entire flight line. One advantage of a single epoch solution is that as soon as sufficient information is available again at an epoch after a gap, the ambiguities can be resolved and position and attitude values can be computed. A multiple epoch solution would require a period of consistent observations to recover ambiguities after a long data gap. Filtering and smoothing could be used to provide an approximate solution during a relatively long data gap. With these data frequent short gaps do not prevent a solution, they simply reduce the frequency of the available solution from a potential 5Hz to 4 or 3Hz on day 288, and from 10Hz to 8 or 6Hz on day 289.

6.3 Positioning Results

6.3.1 Introduction

During the test flights, the four independent receivers on the aircraft and a local reference receiver recorded GPS data. Details of the observation conditions and the completeness of each data-set are described in section 6.2 and appendix B. Full dual-frequency code and carrier phase observables were recorded, and the aircraft stayed within 6km of the reference station. These conditions allow a high accuracy, fixed ambiguity carrier-phase solution to be derived. A solution of this type will generally be accurate to a few centimetres, well within the 0.25 to 1m requirement described in section 2.3.2.

The aircraft antenna positions have been determined from the GPS data using two distinct approaches. The first uses a single epoch ambiguity resolution technique based on the ambiguity function method (AFM), implemented in the GASP software package and the second uses a commercial post-processing package, SKI, which uses the fast ambiguity resolution approach (FARA) using multiple epochs of data for ambiguity resolution. These software packages and the algorithms they employ have been described in section 5.3.

To derive all six exterior orientation parameters using only a single epoch of data, would require GASP to be used for positioning whilst GRAPE is used to derive attitude. A benefit of this combined solution is that it is insensitive to cycle slips and loss of lock. An assessment has therefore been made of GASP's ability to resolve ambiguities during the flight lines and hence produce accurate positions. It should be noted however, that whilst a single epoch position solution may offer advantages for the current and future applications, if a multiple epoch solution is able to produce usable positions, then this approach satisfies the direct positioning requirements of this direct georeferencing project equally well. Additionally, if the positions of at least three aircraft antennas can be determined using either a single or multiple epoch positioning algorithm, then this information can also be used to determine all three attitude components (see section 3.4.2.1).

The ability of a carrier-phase solution to deliver centimetre level positioning accuracy depends on successful ambiguity resolution. Even if the data used in the subsequent least squares estimation is relatively noisy, the final solution should be accurate to around 2 to 3cm in plan and 3 to 5cm in height, provided the correct ambiguities have been found. The main part of this assessment will therefore concentrate on the success rate of ambiguity resolution (section 6.3.2). In section 6.3.3, an assessment of final positioning accuracy is made, noting the difficulty of quantifying accuracy in the absence of a reference solution.

6.3.2. Ambiguity Resolution

Successful ambiguity resolution relies on accurate fractional phase measurements. As the level of noise in the phase observations increases, the accuracy of these fractional phase measurements deteriorates and the ambiguity resolution becomes less successful. The primary source of noise in this application is multipath, which is therefore the limiting factor on the success of ambiguity resolution and hence the overall positioning performance. As described in section 5.3.2, previous tests have shown that GASP fails to resolve ambiguities if L1 phase data is in error by more than 25mm for a seven satellite constellation². By default SKI will not attempt to resolve ambiguities if the phase noise RMS is greater than 10mm. This limit can be increased when noisy data is processed so that a fixed ambiguity solution is determined, but the confidence with which the ambiguities are resolved decreases accordingly.

² For the flight lines considered in this research there are always at least seven satellites available above a ten degrees elevation angle (section 6.2).

6.3.2.1 Validation Methods

In the current investigations, the following four methods have been used to assess whether the correct ambiguities have been determined by the two positioning software packages:

1) Internal Quality Indicators from SKI and GASP

SKI and GASP both produce values that indicate either directly or indirectly whether ambiguities have been correctly resolved. These values can be broadly divided into two categories; those that relate specifically to ambiguity resolution, and those that assess the accuracy of the final solution, and can therefore be used to infer ambiguity success. These two categories are described below for each positioning package.

(i) Specific Ambiguity Resolution Indicators

SKI

For each epoch SKI produces one of three ambiguity indicators

- **Y** means ambiguities have been resolved
- **Y*** means ambiguities have been resolved but the result should be treated with caution
- **N** means ambiguities have not been resolved

The frequencies for which ambiguities have been resolved will also be indicated: either L1, L2, L3 (ionosphere free linear combination) or L5 wide lane. If the ambiguities for L1 and L5 are known then the ambiguities of L2 can be determined.

FARA statistics listing the test criteria and results can also be produced. The values listed are:

- *rms a priori* limit for phase noise rms,
- *rms float* actual rms value before fixing ambiguities,

- *rms fix 1* rms value after fixing ambiguities for the best suitable ambiguity set,
- *rms fix 2* rms value for the second best suitable ambiguity set,
- *error prob.alpha* defines the significance level of 100%- alpha on which the hypothesis test is carried out.

An example of such a summary is given below.

```

Statistical Summary
-----
rms a priori      :    10.0 [mm]
rms float        :     4.8 [mm]
Test (rms float not significantly bigger than rms a priori)
:passed
Error prob. alpha :    5.000 [%]

rms fix 1        :     5.2 [mm]
Test (rms fix 1 not significantly bigger than rms a priori)
:passed
Error prob. alpha :    5.000 [%]

Test (rms fix 1 is significantly smaller than rms fix 2)
:passed
Error prob. alpha :    0.001 [%]

```

```

Search Statistic
-----
set#    rms      ratio
  1     0.0052   1.0000
  2     0.0091   1.7491

```

GASP

Using GASP there are no indicators of ambiguity success that are as straightforward as SKI's 'yes/no' system. Al-Haifi (1996) and Mahmud (1999) both used the following three values to judge ambiguity success at an epoch:

- *MAFV* the maximum ambiguity function value at an epoch,
- *pA* the number of trial positions leading to AFVs within a fixed range of the maximum value (typically $0.9 \cdot \text{MAFV}$),
- *pF* the total number of candidates passing the F-test.

The methods for computing the above values are described in sections 5.3.2.1, 5.4.1.6, and 5.4.1.7.

If p_A and p_F are unity and MAFV is greater than 0.95 then it is accepted that ambiguities have been resolved. These test criteria have proved appropriate using static data with low measurement noise. In such an environment, when the single epoch ambiguity resolution method is successful, these parameters will generally have values that satisfy the test criteria. At an epoch, or epochs, where ambiguities are not resolved the derived values will usually fail these tests, hence the criteria provide a useful relative test. However, in situations with greater measurement noise, typical of a dynamic environment, these test criteria are often inappropriate. If the criteria are met it is almost certain that the correct ambiguities have been found, if however the criteria are not met it does not follow that incorrect ambiguities have been selected. In the current application, trial positions producing AFVs as low as 0.73 have led to correct ambiguities. Ambiguity function values will also vary depending on the number of satellites observed, which again makes them a better relative guide than an absolute test.

With noisier data it becomes more difficult to isolate the correct set of ambiguities, leading to p_A and p_F values greater than unity. This means that the F-test has failed to identify the solution with the lowest sum of the squares of the weighted residuals ($\hat{v}^T W \hat{v}$) as significantly better than any other possible solution, however, the chosen solution does generally produce the correct ambiguities. Noisy data is also more likely to produce a trial position leading to an AFV significantly lower than the maximum AFV, which is in fact the trial position producing the correct ambiguities. It may therefore be necessary to carry more trial positions forward to the least squares estimation by retaining candidates which produce an AFV within 80% or 70% of the maximum value, rather than the standard 90%. Even though it produces a relatively low AFV, the correct trial position can generally be accepted on the basis of the summed residuals following the least squares adjustment.

The limitations of judging ambiguity resolution success from these parameters, particularly considering their relative nature, means that these indicators cannot be used as the sole means of assessing ambiguity resolution performance.

(ii) *Final Solution Quality Measures*

It is often possible to infer whether or not ambiguities have been correctly resolved by inspecting the precision measures associated with the calculated positions. SKI and GASP both produce values from which the standard deviation of each component can be determined. SKI produces the rms *a posteriori* and a co-factor matrix (which is the inverse of the normal equation matrix) from which a full covariance matrix can be formed, and hence the standard deviation can then be derived. GASP produces a standard deviation for each position component which can be scaled using the unit variance to provide an internal accuracy indicator.

When SKI is able to produce a fixed ambiguity solution the estimated standard deviation for any position component is generally well below 20mm, typically between 5 and 10mm. When ambiguities have not been resolved the standard deviations are around 1m in plan and 1.5m in height.

GASP produces more conservative estimates of accuracy, generally around 10mm in plan and 15 to 20mm in height when ambiguities have been correctly resolved. Standard deviations are produced from an elevation dependent stochastic model and do not change to reflect the success of ambiguity resolution. The unit variance reflects the level of agreement between the complete set of observations and the computed solution. It does increase if ambiguities are not resolved correctly, but it is not sufficiently sensitive to be used directly as an indicator of ambiguity resolution success.

2) **Consistency of the Calculated Solution**

By simply inspecting the final positions from GASP it is possible to make a qualitative assessment of the ambiguity resolution. For example, plotting the height component against time will show whether a consistent solution has been produced. Figure 6.1 shows the height determined by GASP for antenna 4, from the reference station. Clearly no consistent correct set of ambiguities has been

found. Such a test cannot identify when the solution is correct and when it is not, but it clearly shows if a problem exists and demonstrates that further processing is required to correctly resolve ambiguities.

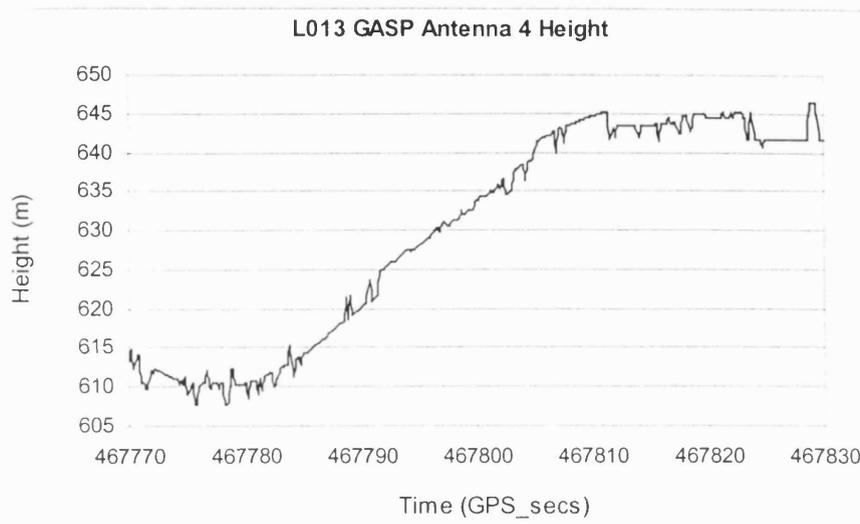


Figure 6.1: GASP derived height of antenna 4 during flight line L013.

A similar check cannot be applied to SKI as the algorithm will produce a smoothed trajectory regardless of whether ambiguities have been correctly resolved. If the ambiguities are not correctly resolved there will be a constant bias in the solution but there are unlikely to be obvious steps in the solution as there are with GASP where ambiguities are resolved independently at each epoch.

3) Comparisons of SKI and GASP solutions

SKI and GASP use different processing strategies and algorithms, as described in section 5.3. If the same data are processed using both approaches and the resultant positions for the unknown airborne receiver agree to within a few centimetres, it is almost certain that both packages have correctly resolved the ambiguities. It is highly unlikely that this level of agreement is the result of both packages consistently using the same wrong set of ambiguities. If this were the case, the fact that the ambiguities were incorrect should be indicated to some extent by internal quality measures, particularly those from SKI, as described above. If either approach produces a solution with proven ambiguity success, this can then be used as a reference solution for the other approach.

4) Vector Checks Between Aircraft Antennas

As there are four antennas on the aircraft, each of which can be positioned relative to the same local reference station, a check can be carried out using the baseline lengths between aircraft antennas. Two or more airborne antennas are processed independently from the same reference station (clearly the observations at the reference station are common). From the positions of these two antennas the baseline between them is derived. If this baseline length agrees with the *a priori* value from the ground survey to within a few centimetres it is inferred that the ambiguities have been successfully resolved for each antenna.

There is a possibility that a consistent vector length could result from the consistent use of an incorrect ambiguity set, but this is unlikely and can be ruled out using a combination of the other tests described. This test may be the most trustworthy as it uses independent data from each antenna and therefore applies a more external test than an internal ‘goodness of fit’ test.

One drawback with this approach is that it requires two or more antennas to be positioned successfully before confirming successful ambiguity resolution. There are situations in which the phase observations recorded at two or three of the aircraft antennas are consistently too noisy to resolve ambiguities with the remaining antenna not being affected in the same way. Data from this antenna can be used to produce a successful fixed ambiguity solution, but this success cannot be confirmed unless a second equally successful solution is available.

6.3.2.2 Robustness Results

Through combining the methods described in section 6.3.2.1, it is possible to assess, with a reasonable degree of certainty, whether or not ambiguities have been successfully resolved using each processing approach, i.e. GASP and SKI. Following some preliminary tests, a strategy was devised to assess the ambiguity resolution success, or robustness, for each approach using these validation methods. This processing strategy is outlined in figure 6.2.

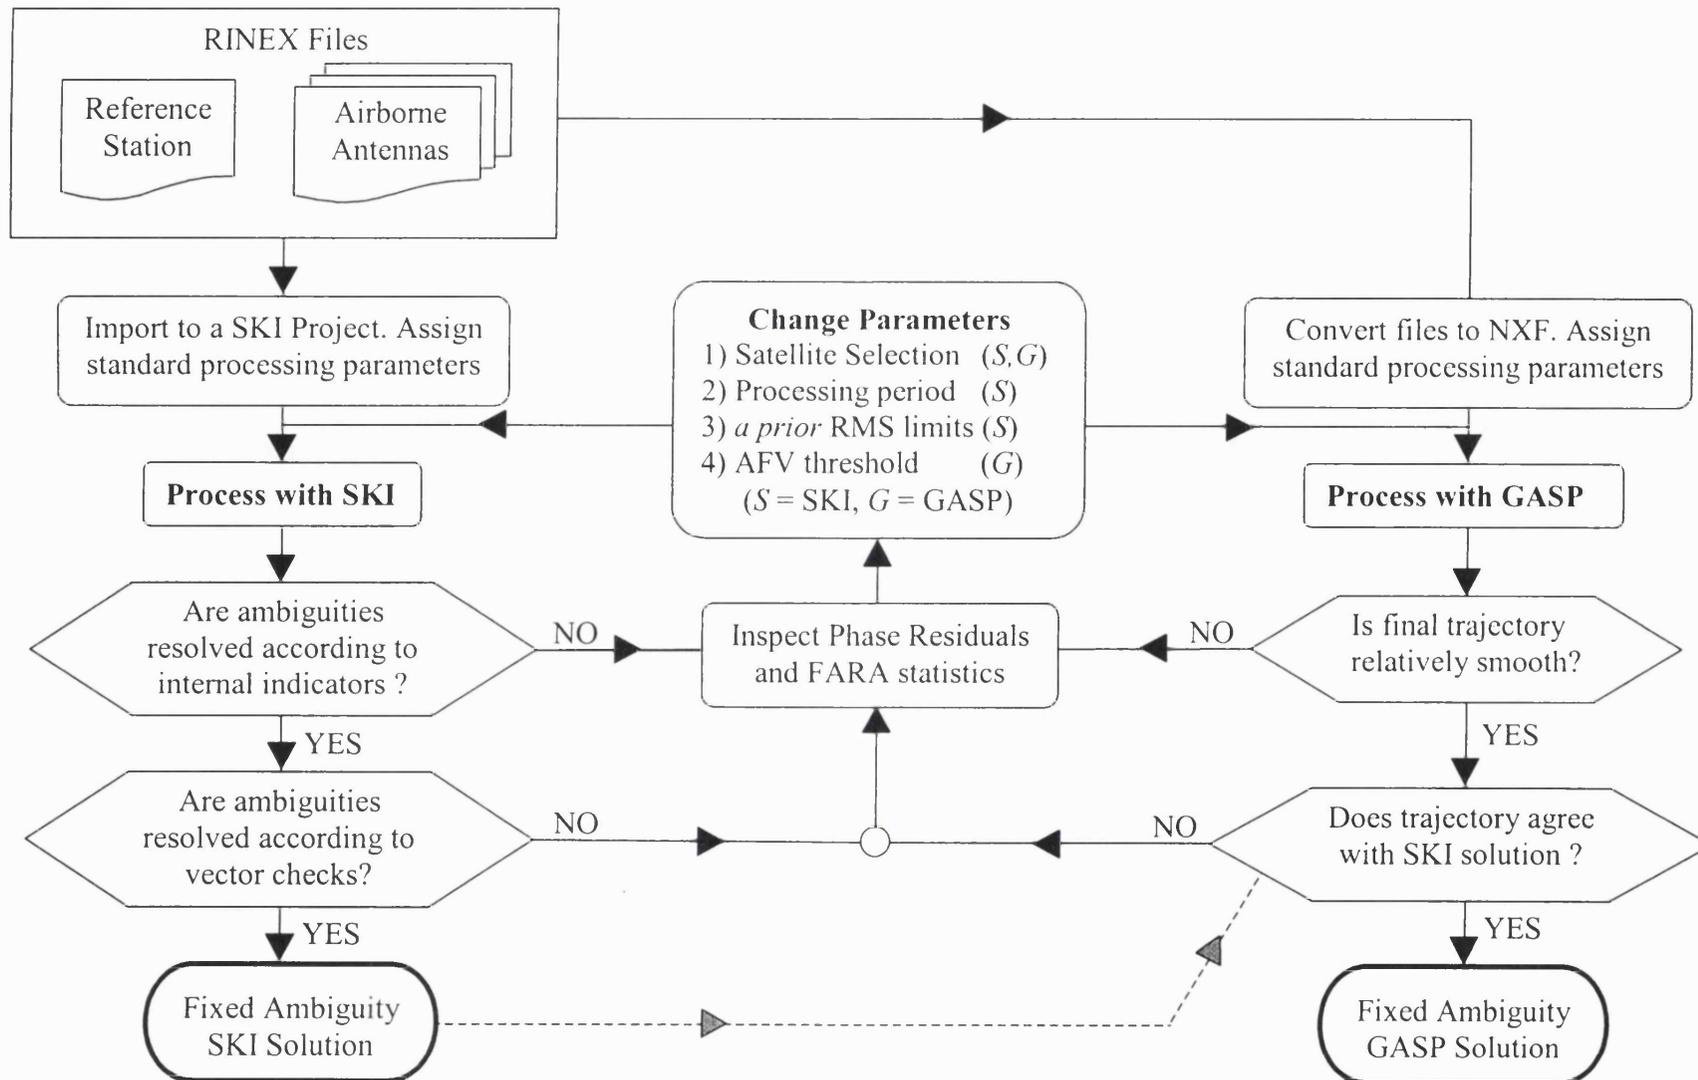


Figure 6.2: Ambiguity validation

The figure shows two distinct paths for SKI and GASP, both using the same GPS data. When using SKI, the RINEX files are imported into a project where they are stored in a standard Leica format. To process the data with GASP it is first converted to the NXF format, as described in section 5.2. Initially the data are processed using standard processing parameters. For SKI this includes a 20° elevation mask, and a limit on the phase measurement rms of 10mm. For GASP the elevation mask is also set at 20° initially with an AFV threshold of 90%, i.e. only trial positions producing an AFV within 90% of the maximum AFV are tested in the least squares estimation.

SKI Processing

Once SKI has processed the data, the internal ambiguity resolution indicators are inspected. If they indicate that the ambiguities have been successfully resolved, the computed antenna positions are then used to check the aircraft antenna baseline lengths. If all these criteria are satisfied the solution is considered to be a trustworthy fixed ambiguity solution. This solution is used later as a reference for the GASP results.

If on the other hand, any of these tests indicate that there are problems with the ambiguity resolution, the single-differenced phase residuals for each satellite-baseline combination are inspected, together with the FARA statistics. Based on this information a number of parameters can be changed and the data re-processed. The performance can be improved either by identifying and removing noisy data from the solution, or by changing thresholds to allow processing to continue even in the presence of these noisy measurements.

To remove noisy data, the processing period can be altered, or particular satellites can be removed from the solution. SKI requires a period of continuous good observations, i.e. no significant measurement noise, to resolve ambiguities. Ideally periods of at least ten minutes should be used, but in this application such a period will often include measurements with significant noise, coinciding with

an aircraft manoeuvre. Processing sessions must be carefully chosen to maximise the amount of good data used whilst avoiding periods of noisy data. If the phase residuals from individual satellites indicate that some satellite-antenna combinations are more contaminated by multipath than others, observations from these satellites can be removed.

If it is not possible to identify and remove noisy data, or if removal leaves too few observations for a successful solution, the limit on allowable phase noise can be increased. Increasing this limit means that SKI will attempt to resolve ambiguities when otherwise it would simply stop processing. Increasing this limit and allowing a set of ambiguities to be chosen despite the noisy phase data, increases the probability that an incorrect set of ambiguities will be accepted and is generally not recommended, particularly for short observation periods. In this application however, the resultant solution can then be validated using alternative methods.

By iteratively processing with adjusted parameters, SKI was able to successfully resolve ambiguities on baselines between the reference station and all aircraft antennas³ during every flight line, with one exception. During the flight line L015 on day 289, SKI is unable to resolve ambiguities between the reference station and antenna 3. This is due to insufficient measurements from the antenna 3 data-set. Of the 700 epochs on this flight line, data was only recorded from antenna 3 at 62 epochs and a similar pattern occurs in the periods immediately before and after the line was flown.

GASP Processing

A similar iterative process of inspecting the output and then changing parameters has been used to determine the best possible GASP solution. The first test of the GASP solution is the smoothness of the trajectory. If different ambiguity sets are selected at different epochs the final trajectory shows the kind of steps and spikes apparent in figure 6.1. If an inconsistent pattern is apparent the data are

³ On day 288 no data were recorded by the receiver connected to antenna 3.

reprocessed with different parameters. If the trajectory shows a smooth trend, indicating that a consistent set of ambiguities has been determined, this solution is then checked against the SKI reference solution. If the agreement is at the level of a few centimetres per epoch, the GASP solution is considered to have correctly resolved ambiguities at that epoch. The comparison is made at every epoch during the flight line to give an overall assessment of GASP ambiguity resolution.

If GASP fails to resolve ambiguities the phase residuals produced by SKI, using similar processing parameters, are inspected and the parameters are adjusted to try and improve the solution. It was found that with particularly noisy measurements the trial position leading to the correct ambiguity resolution did not produce the maximum AFV, or an AFV within 90% of the maximum value. In these cases the threshold was lowered to 70 or 80% and for many epochs this then led to the correct ambiguities being selected based on the minimum $\hat{v}^T W \hat{v}$ criteria.

Significant improvements in the success of ambiguity resolution were achieved by identifying noisy phase observations from specific satellites to specific antennas and removing them from the solution. During flight lines L006, L008 and to a lesser extent L003, successful ambiguity resolution was rarely achieved on baselines involving antenna 2 and antenna 4. For the same period and processing parameters, ambiguities were correctly resolved at all epochs for the baseline between the reference station and antenna 1. The ambiguity success rates using observations from all satellites above the 20° elevation mask, processed with standard parameters are expressed for these periods as a percentage in table 6.2.

Flight Line	Reference Station to Antenna #		
	1	2	4
L003	100%	100%	62.5%
L006	100%	0.5 %	37.2%
L008	100%	1.2%	0%

Table 6.2: GASP ambiguity success rate with standard parameters.

Inspection of the phase residuals produced by SKI using similar observations showed that satellite PRN 4, produced significantly larger phase residuals at antennas 2 and 4 than any other satellites. Tables summarising some important characteristics of these phase residuals are presented in appendix C. The residuals from PRN 4 reached 38mm in the case of L2 at antenna 4 during line L008, and the mean residual was 21mm. The mean residual from the other satellites varies between 2 and 9mm with maximum values between 6 and 15mm. This trend of phase residuals being significantly larger for PRN 4 than the other satellites is apparent for L1 and L2 observations recorded at antennas 2 and 4 during flight lines L006 and L008. During these flight lines, PRN 4 is the lowest elevation satellite, around 25° and has an azimuth of 196° (appendix C). These two flight lines head approximately north, hence the satellite signal appears to approach the antennas from the rear of the aircraft. It seems likely that the increased measurement noise is caused by multipath due to signal reflectance from the tail plane and fuselage. Removing PRN 4 from the solution improves the success rate of ambiguity resolution to 100% and 91.5% for antennas 2 and 4 respectively on line L006, and to 99.6% and 97.7% for antennas 2 and 4 respectively on line L008.

During flight line L003 there is an additional low elevation satellite, PRN 16. The phase residuals from this line show that observations from PRN 4 and 16 are noisier than those from the remaining six satellites. The magnitudes of the residuals are significantly less than that observed during L006 and L008. With all available satellites above the 20° elevation mask used in the solution, ambiguities are successfully resolved at all epochs between the reference station and antennas 1 and 2, but the success rate is only 62.5% for antenna 4. Removing satellites 4 and 16 improves the ambiguity success rate between the reference station and antenna 4 to 100%. The solution is clearly more successful when PRN 4 observations are used during L003, than when they are used during L006 and L008. This decrease in the ambiguity resolution success over time can be related to the elevation of PRN 4 which is decreasing, and the presence of additional, although relatively noisy, observations during L003 from PRN 16.

During these periods, SKI was able to resolve ambiguities on baselines between the reference station and antennas 1 and 2 using standard parameters without removing any satellites. To resolve ambiguities between the reference and antenna 4 it was necessary to remove PRN 4 during L006, or to increase the limit on allowable phase noise from 10 to 30mm.

The final GASP ambiguity resolution performance after iteratively changing parameters and re-processing is summarised in table 6.3. For each flight line the percentage of epochs at which ambiguities were successfully resolved is shown for each antenna (processed from the reference station).

Flight Line	Reference Station to Antenna #			
	1	2	3	4
L001	100%	100%	NA	78.6%
L003	100%	100%	NA	100%
L006	100%	100 %	NA	91.5%
L008	100%	99.6%	NA	97.7%
L013	52.7%	99.8%	89.6%	35.2%
L015	100%	100%	37.1%*	76.2%

* Only 62 epochs of data are available in total from a possible 700

Table 6.3: Final GASP ambiguity resolution performance

The following conclusions can be drawn from these results:

- Ambiguities for baselines between the reference station and the fuselage antennas, 1 and 2 are resolved more consistently than for the winglet antennas 3 and 4. This suggests that the locations of the winglet antennas may make them more prone to multipathing affects.

- For every flight line it has been possible to correctly resolve ambiguities and hence determine accurate positions for at least one antenna all the time⁴.
- For the majority of the flight lines it is possible to position three antennas at any epoch, this is important to note as it means that all three attitude components could be derived using antenna positions

The robustness of GASP may be improved further by more consistently identifying and removing satellites that lead to large phase residuals. To date there is no automated procedure in GASP or SKI for removing individual satellites based on a data rejection criteria. To improve the GASP results on lines L006 and L008, observations from PRN 4 were removed for the entire flight line. Results may also be improved by removing the observations from individual satellites, or combinations of satellites, in turn at each epoch. During lines L013 and L015 there are three satellites with elevations between 10° and 25° all of which produce relatively large phase residuals at certain times. Tests to remove one, two, or three of these satellites from the solution at various epochs have produced some improvements in the ambiguity resolution success, but to experiment with every possible combination of satellites at every epoch over these flight lines was considered unnecessary, particularly as the position requirements for this application can be adequately met without achieving 100% ambiguity resolution success from GASP for all antennas at all times.

To summarise, a fixed ambiguity solution can be achieved for all the airborne receivers during all the selected flight lines using the multiple epoch ambiguity resolution method implemented in SKI. Using the GASP single epoch algorithm, ambiguities can be correctly resolved on baselines between the reference station and all of the aircraft antennas, most of the time. In general, GASP is able to produce a fixed ambiguity solution at two or three aircraft antennas during all the flight lines.

⁴ On L013 ambiguities are not resolved on antenna 2 at one epoch, but they are resolved at antenna 3 at that epoch.

6.3.3 Accuracy Assessment

The only way in which the accuracy of any measurement system can be assessed with complete certainty is to compare the estimates obtained from the system to known reference values. The difference between the estimated value and the known value then represents the true error of the measurement system. To assess the accuracy of GPS carrier-phase positioning, numerous tests have been undertaken in a static environment where a reference solution can be established relatively easily, see for example Remondi (1985) and Corbett (1994). The following procedure is usually adopted:

- 1) Set up two GPS receivers over points whose co-ordinates are known.
- 2) Treat one of the receivers as a known reference station and the other as unknown.
- 3) Determine the unknown station position relative to the reference station.
- 4) Compare the co-ordinates computed for the unknown station to the existing 'truth' set to assess the system's true accuracy and to assess how well the system has predicted the actual errors.

In a dynamic environment, there is generally no convenient reference of this kind. Experiments have been designed on land, in which a moving vehicle retraces an exact route a number of times, the positions on the route determined from different runs can then be compared. This approach uses the repeatability of the solution to make inferences about the accuracy. It is also possible to precisely survey the fixed route using conventional survey techniques, and to then compare the GPS derived route against the reference. Results derived from this 'test-bed' approach are described in Edwards *et al.* (1999) and Corbett and Cross (1995). These tests not only attempt to measure the accuracy of the system, they also assess how well the internal accuracy measures produced by the system are describing the actual performance.

It is generally accepted, based on numerous calibration tests of this kind, that if the integer ambiguities have been correctly resolved in a GPS carrier-phase solution, the final accuracy of the position will be of the order of 1 to 2cm in plan and 3 to 5cm in height.

In this application therefore, once it has been established that ambiguities have been correctly resolved, it is possible to make some assumptions regarding the overall accuracy of the solution. The validity of these assumptions has been tested for the current data set to allow a more comprehensive assessment of the accuracy to be made. As no reference solution is available in this application, similar approaches to those used in the previous section have been implemented. In this case, three methods have been used to make an assessment of the positioning accuracy. These methods are described in turn, together with some results from their implementation.

6.3.3.1 Precision Measures from SKI and GASP

As described in section 6.3.3.2, SKI and GASP produce standard deviations for each position component derived from the covariance matrix of the estimated parameters and the unit variance, both of which are determined during the least squares estimation. Quality measures from the many GPS processing software packages have been shown to be overly optimistic about the accuracy of the position results. This is generally due to the implementation of relatively simple stochastic models that make some incorrect assumptions regarding the nature of the errors in the system and particularly their correlation (Tiberius, *et al.*, 1999).

Barnes *et al.* (1998) showed that the accuracy estimates from an elevation dependent stochastic model, as used in GASP and SKI, are largely geometry dependent and do not follow many of the trends apparent in the 'true' position errors. The effects of multipath, which may significantly increase the observation noise level, have not been incorporated into the stochastic or functional models of GASP or SKI. Despite the limitations of the stochastic models, the standard deviations do still provide a convenient guide to the accuracy of the final solution.

Figures 6.3 and 6.4 show the standard deviations derived from SKI and GASP respectively for the position of antenna 2 over flight line L013.

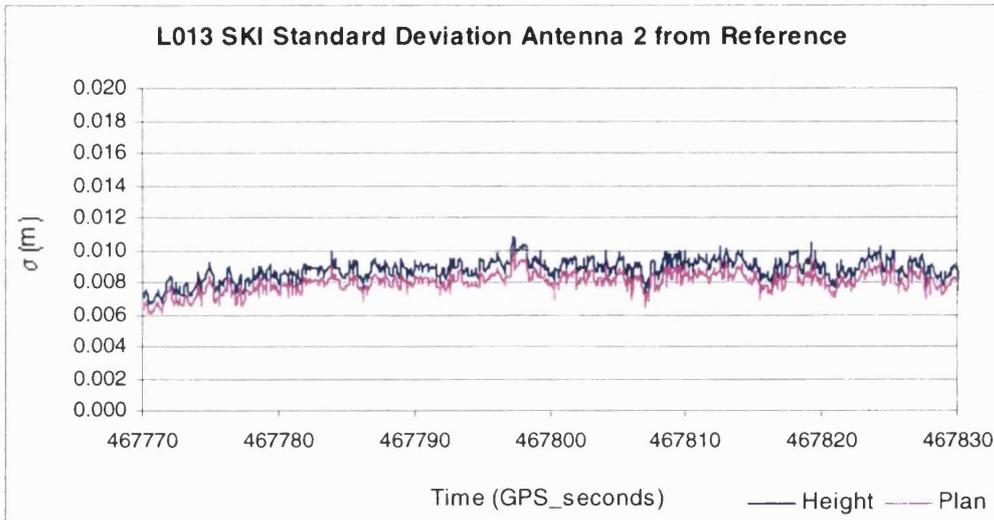


Figure 6.3: Standard deviations from SKI processing. L013, antenna 2.

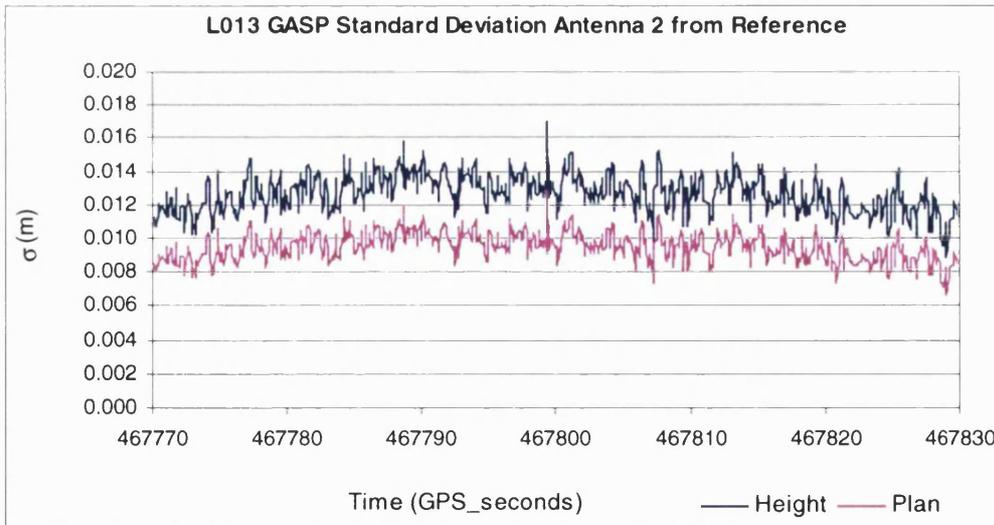


Figure 6.4: Standard deviations from GASP processing. L013, antenna 2.

The standard deviations in plan and height determined from this particular data set by SKI are typical of those recorded on all flight lines at all antennas. Typically the standard deviation in plan is between 5 and 15mm, and increases to between 8 and 20mm in the height component.

GASP estimates of accuracy are more conservative, typically 8 to 20mm in plan and between 10 and 30 mm in height. The standard deviation in height reaches 35mm to 40mm for antenna 4 on a number of flight lines, due to the higher noise levels noted previously. These accuracy measures can only be relied upon once the correct ambiguities have been validated.

6.3.3.2 Comparisons of SKI and GASP Solutions

The level of agreement between SKI and GASP solutions can also be used to make some inferences about the accuracy of each system. Neither solution can be considered significantly more accurate than the other, hence one solution cannot be used as a reference for the other. Instead, the difference between the two solutions is measured and it is assumed that the magnitude of this difference can be equated to some extent with the accuracy of each individual solution.

Figures 6.5 and 6.6 show the difference in plan and height between the SKI and GASP solutions for antenna 2 on flight line L001 and antenna 1 on L008, respectively.

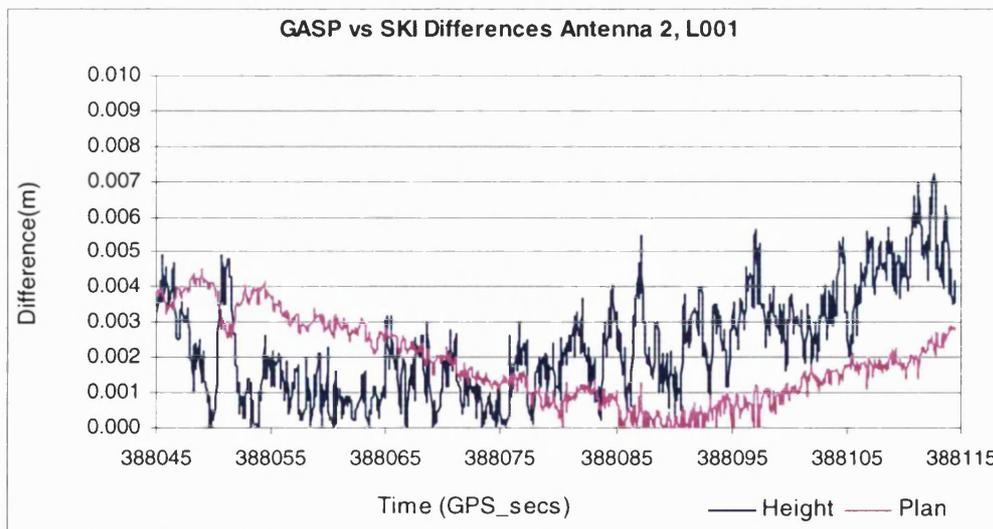


Figure 6.5: GASP and SKI agreement for antenna 2 position on L001.

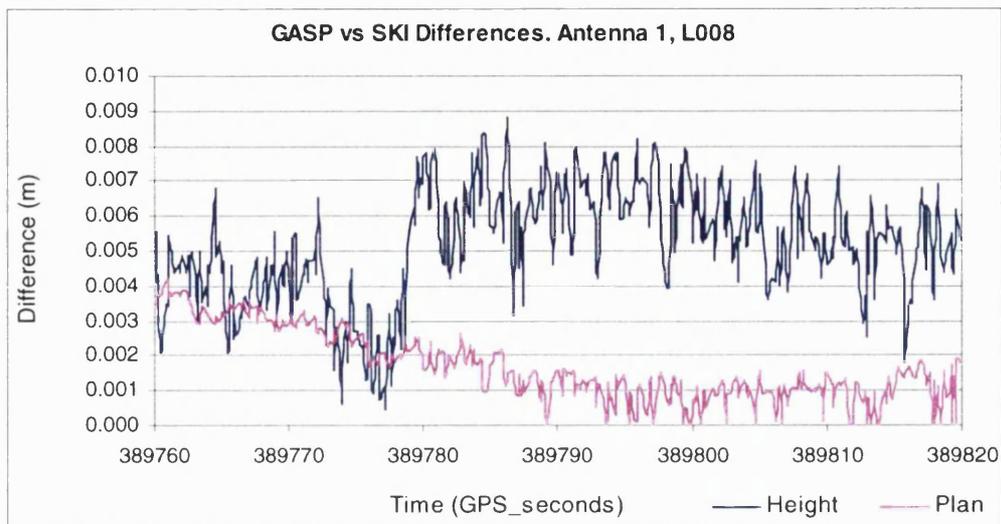


Figure 6.6: GASP and SKI agreement for antenna 1 position on L008.

These figures are typical of the level of agreement between the two solutions for all antenna positions during each flight line, provided ambiguities are resolved successfully. For each pair of solutions (i.e. SKI and GASP positions for each antenna during each flight line) the size of the differences were computed. These values range between 3 and 12mm in plan, 5 to 22mm in height, 95% of the time. In order to produce this level of agreement, each individual solution must be of a similar accuracy.

These results compare favourably with those from tests conducted jointly by Delft University of Technology, the Netherlands National Aerospace Laboratory, and the Applied Physics Research Centre (Husti and Sluiter, 1999). In their tests, airborne GPS data were post-processed with three alternative software packages. The results generally agreed to within 20mm horizontally and 50mm in height. It should be noted however, that in these tests the maximum distance between the aircraft and the reference was 35km. During the flight lines considered in the current project, the maximum distance is 3.8km.

6.3.3.3 Vector Checks between Aircraft Antennas

The baseline lengths between the aircraft antennas, established in a ground survey, can also be used as an accuracy reference. If two aircraft antennas are positioned from a common reference station, the co-ordinate differences between these antennas can be used to compute a baseline length. The level of agreement between this derived distance and the *a priori* distance provides some measure of the accuracy with which each antenna has been positioned.

Figures 6.7 to 6.10 show the difference between the derived distance and the *a priori* value for a selection of antenna baselines. The first three plots (figures 6.7 to 6.9) are vectors derived from SKI positions on line L006, the fourth (figure 6.10) is the vector between antennas 1 and 2 derived from GASP positions on line L001.

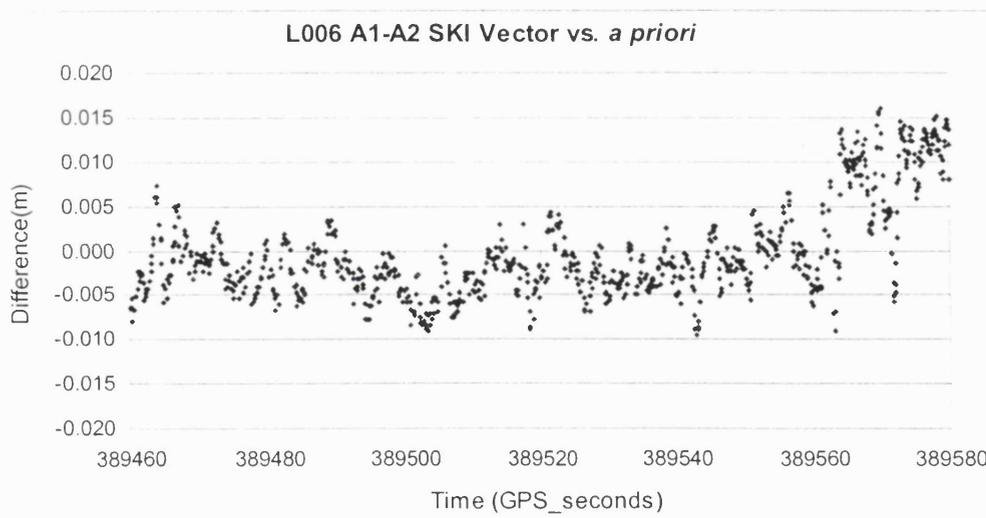


Figure 6.7: A1-A2 vector from SKI positions, L006.

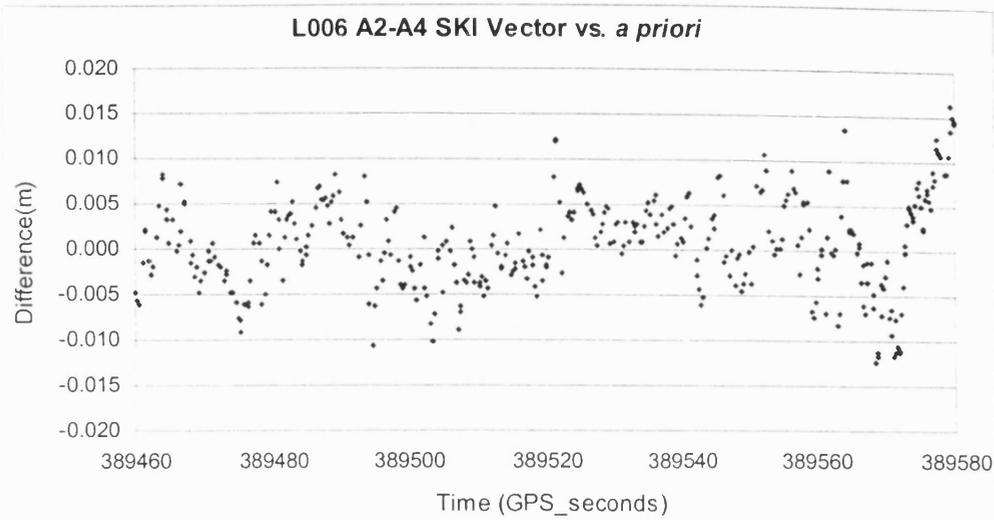


Figure 6.8: A2-A4 vector from SKI positions, L006.

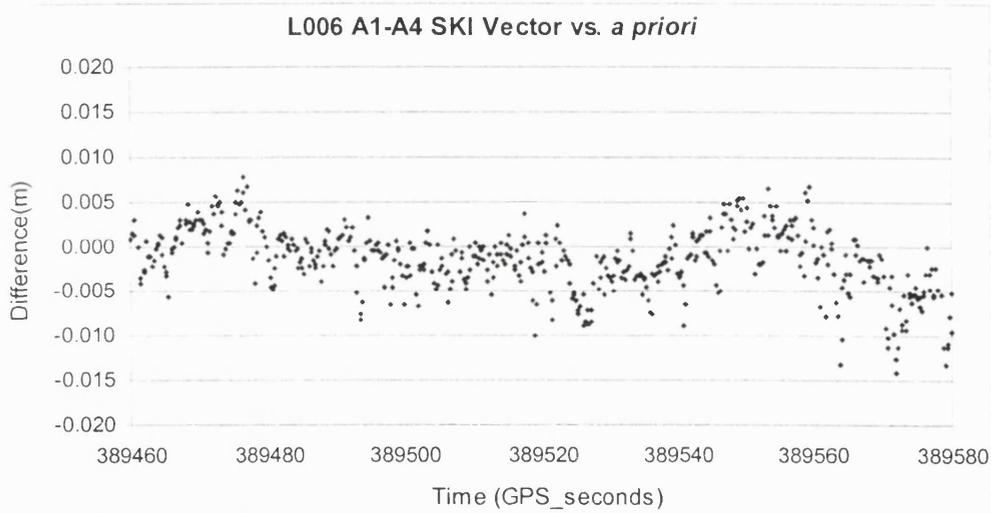


Figure 6.9: A1-A4 vector from SKI positions, L006.

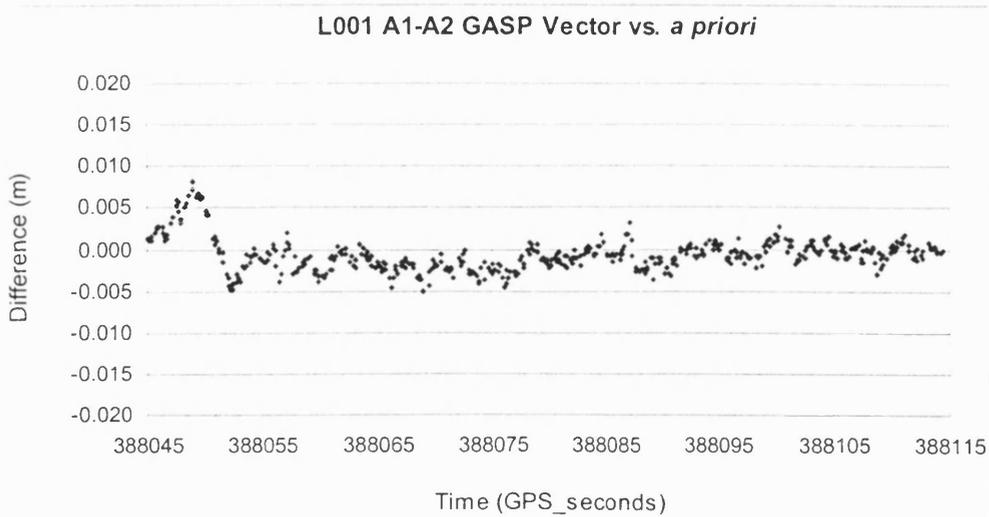


Figure 6.10: A1-A2 vector from GASP positions, L001.

These plots have been chosen as they are representative of vector results from both SKI and GASP derived antenna positions, on all flight lines. They show that for the majority of the data set analysed, the distances derived from positioning the aircraft antennas differ from the *a priori* distances by between 5 and 10mm, with the maximum difference around 20mm. This suggests that each antenna is positioned from the reference station to a similar level of accuracy.

6.3.3.4 Positioning Accuracy Summary

Combining the results from these different approaches to assess accuracy, suggests that the aircraft antennas are generally positioned to between 5 and 10mm in plan and 10 to 30mm in height, although these values may increase to around 20mm and 40mm respectively when noisier data are used in the solution. The accuracy with which the perspective centres of the scanners are positioned depends on the accuracy of the antenna positions *and* the accuracy with which these values can be related to the scanner. Providing that the level of accuracy with which the antennas are positioned can be successfully transferred to the scanners, the positioning accuracy is well within the 0.25 to 1m general specification for airborne scanners described in section 2.3.2. These centimetre level positioning errors will make a very small, perhaps negligible, contribution to the overall error budget of the current direct georeferencing system.

Standard fixed ambiguity carrier-phase solutions of the type used in this research are usually restricted to ranges of less than 10 or 15km. In general operation, it may not be possible to fly the survey aircraft within this range of a local reference station at all times. In these circumstances an alternative position solution would need to be used. Various techniques have been developed that increase the distance over which ambiguities can be resolved, or which produce accurate solutions without fixing ambiguities to integer values. Colombo *et al.* (1995) for example, used long-range carrier phase DGPS to position a moving vehicle from a reference station more than 1000km away, with an accuracy of 10cm. No investigation of long range positioning accuracy has been undertaken in the course of the current research.

6.4 Attitude Determination Results

6.4.1 Introduction

This section presents the attitude results produced from processing GPS data recorded by the four independent aircraft receivers with GRAPE. As discussed in section 5.4.1, GRAPE employs the direct attitude algorithm, outlined in section 3.5, to estimate three dimensional attitude components directly from GPS phase observations. This method was adopted for the current research project after its benefits in terms of robustness and reliability were demonstrated by Mahmud (1999). The original software was developed and tested using static data. The changes necessary to allow these algorithms to be tested in a dynamic environment were described in section 5.5.

In section 6.4.2, the attitude values determined by GRAPE are compared with alternative attitude solutions to allow a preliminary assessment of the system's performance. The first of these alternative solutions comes from the gyro based AHRS, the second is an attitude solution derived from GPS positions. The degree to which attitude values from GRAPE reflect trends in the aircraft position is also demonstrated. Essentially these tests assess whether or not GRAPE is producing 'reasonable' attitude values.

Performance is then described in section 6.4.3 in terms of the ability to correctly resolve ambiguities, and a comparison is made between the robustness of the direct attitude approach implemented in GRAPE, and a baseline-by-baseline approach, of the kind described in section 3.4.2.1. In section 6.4.4 an assessment is made of the accuracy of the GRAPE attitude solution.

It was the intention to use the attitude solution from the ADU2 to calibrate the independent receiver approach to the accuracy level of the ADU2. A further calibration would have used the total 'on-the-ground' positioning error. By quantifying the errors due to positioning the scanners from GPS measurements, and any further errors due to the scanners themselves, the remaining error could then be attributed to the combined affects of the attitude errors in each component.

To do this however, requires the scanner imagery, which has not been made available. As no reference solution of sufficient accuracy is available to properly calibrate the attitude accuracy, alternative methods have been developed and applied to assess the GRAPE solution.

6.4.2 Preliminary Assessment of Attitude from GRAPE

To provide an initial assessment of the attitude results from GRAPE, three different methods have been used. The first compares the values from GRAPE with those from the AHRS, the gyro based instrument described in section 4.2.4. The second method uses attitude values derived from antenna positions as a comparative solution. Finally, a simple test checks whether or not the manoeuvres that can be seen from plots of the aircraft positions are reflected in the attitude solution.

In the current application, where the AHRS comparative solution is available, the remaining checks are largely redundant. They do however provide some further corroborative evidence, and they may be of use in future applications where the GPS solution may be the sole means of attitude determination.

6.4.2.1 GRAPE and AHRS Comparisons

In this application, the most straightforward way to test if GRAPE is producing reasonable answers, is to compare the attitude solution from any flight line with the corresponding solution from AHRS. The comparison between two different forms of measurement, in this case GPS and gyros, does not explicitly yield accuracy information beyond the specification of the reference instrument. The absolute accuracy of the AHRS solution, 1° in pitch and roll, and 2° in heading (95%), is not good enough to ‘truth’ the GRAPE solution, however, the resolution of AHRS in relative mode is 36”. It therefore records *trends* in attitude relatively well and it is easy to see whether or not GRAPE has recorded the same trends.

With the data available, it is not possible to calibrate the axes misalignments that are known to exist between the reference frame defined by the GPS antenna array and the internal reference frame of the AHRS. These misalignments lead to offsets in each attitude component. The largest offset occurs in the heading component as the AHRS unit defines an arbitrary 'north' value as a heading reference each time the unit is switched on. As it is only operated during flight lines, the AHRS is switched on at the start of each line and a new 'north' value is defined. To allow convenient comparisons of the trends in each solution any offsets have been removed.

Figures 6.11 to 6.14 show the pitch, roll and heading produced from GRAPE and AHRS on flight line L008⁵. In figure 6.13 the two heading solutions have been deliberately offset by one degree to allow them to be differentiated. Figure 6.14 shows a 15 second section of the heading component in greater detail, with the offset removed. In each plot, the blue diamonds represent individual attitude values from GRAPE processing. The magenta line is actually made up of discrete values from the AHRS solution, but due to the frequency of the solution (64Hz) these discrete points appear to form a continuous line. During this period, no GPS data were recorded from antenna 3, hence the solution is based on data from antennas 1,2 and 4 only. The inclusion of antenna 4 means that the maximum potential measurement rate for a three dimensional attitude solution is 5Hz.

⁵ Results from the other five flight lines show similar features. They have not been shown to avoid repetition that does not add any new information.

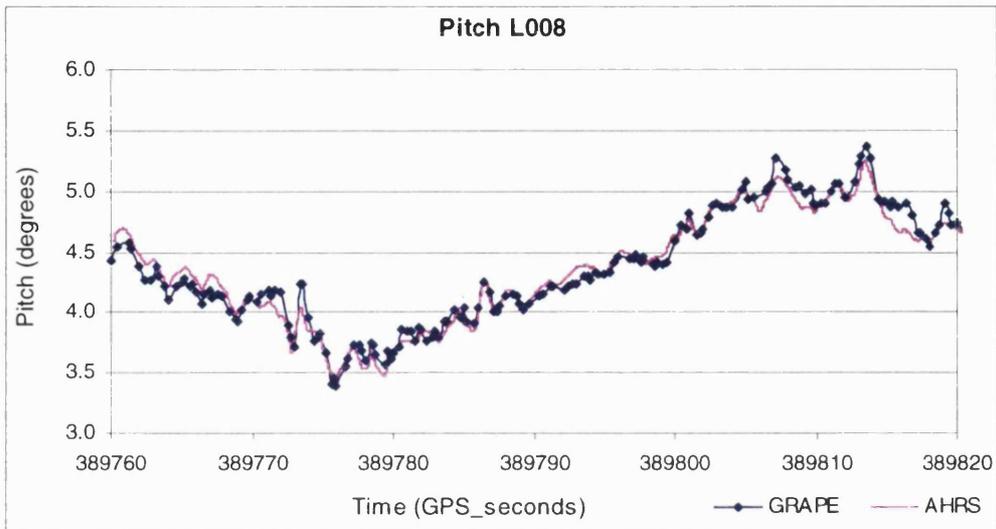


Figure 6.11: GRAPE and AHRS pitch. L008

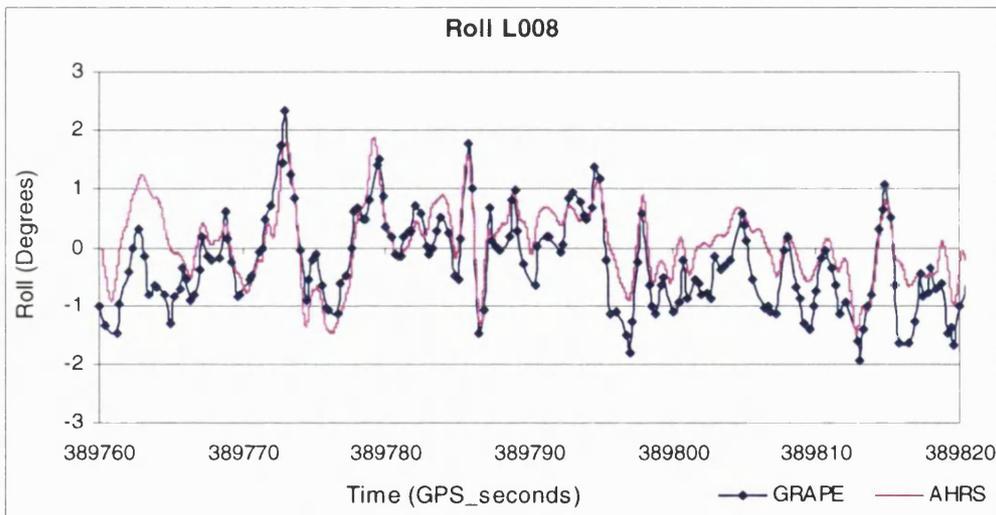


Figure 6.12: GRAPE and AHRS roll. L008

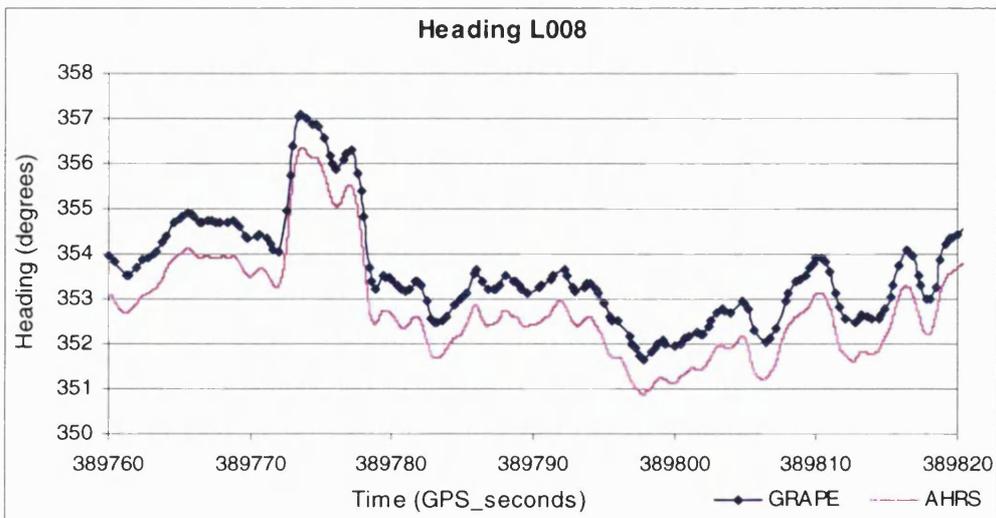


Figure 6.13: GRAPE and AHRS heading. L008

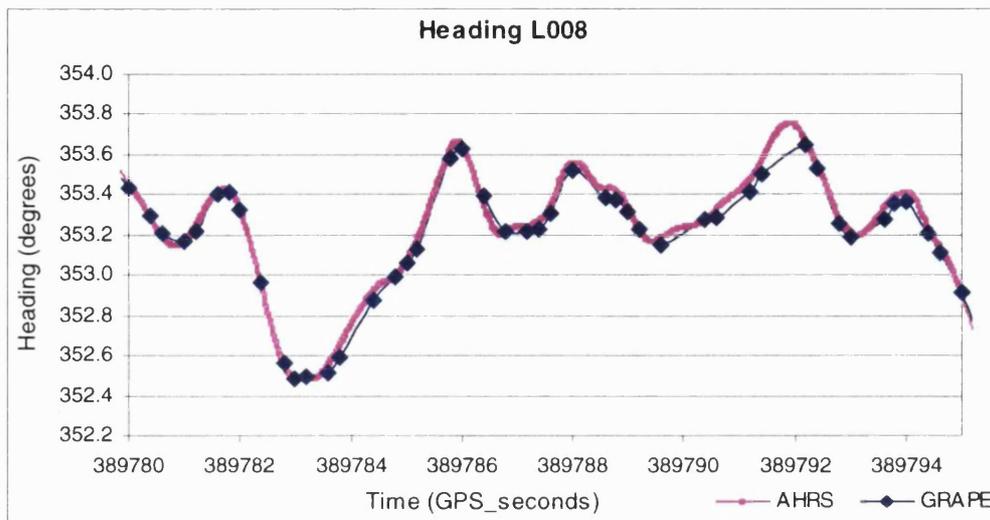


Figure 6.14: GRAPE and AHRS heading. L008 – 15 second period.

These plots illustrate some important points.

- The general agreement between the two solutions is good. The GPS solution from GRAPE is recording the same attitude trends as the AHRS instrument.
 - In pitch, the solutions generally agree to within 5 minutes of arc.
 - In roll, the solutions generally agree to within 20 minutes of arc.
 - In heading, the solutions generally agree to within 2 minutes of arc.

- This level of agreement suggests that GRAPE is consistently selecting the same ambiguities for the complete set of double differenced phase observations at each epoch. If different ambiguities were chosen at different epochs there would be clear spikes or steps in the solution. Although a consistent set of ambiguities is being determined, it is possible that it is not the correct set, consistent but incorrect ambiguities would lead to an offset in the solution. All of the offsets between the GRAPE and AHRS solutions have been attributed to axes misalignments but the possibility that some of this may be due to a bias in the GRAPE solution, caused by incorrect ambiguities cannot be ruled out. Ambiguity resolution is discussed in more detail in section 6.4.3.

- The frequency of the GPS attitude solution from GRAPE has clearly been affected by the data gaps discussed in section 6.2. As a result, the potential 5Hz solution has not been achieved, generally only three values have been computed per second. This is most apparent in figure 6.14 where only 44 GRAPE values from a potential 75, are available over the 15 second period. The lower sampling rate resulting from data gaps can lead to attitude trends being missed, an example of this can be seen in the heading component at time 389792 (figure 6.14).
- Despite the lower frequency of the GPS solution it still appears to record the majority of attitude changes measured by the much higher rate AHRS. Section 6.5.1 discusses the measurement rates of attitude solutions in more detail.
- Despite the generally good level of agreement between the solutions there are also some significant differences. These are most noticeable in the roll component where differences as large as one degree are apparent. A number of factors will contribute to these discrepancies, the two main ones are:
 - The AHRS device and the GPS antennas are not measuring exactly the same physical changes. The GPS antennas effectively measure the average attitude changes experienced by the airframe. The AHRS is rigidly fixed to the ATM scanner, which employs damping devices to reduce the effects of vibration on the imagery. Therefore, the AHRS instrument is measuring a damped version of the airframe motion that the GPS antenna array records
 - As there is no GPS data available from one of the winglet antennas (antenna 3) this effectively reduces the baseline over which roll is measured from 1.2m to 0.6m. As the accuracy of GPS attitude determination is proportional to the baseline length this decreases the potential accuracy of the GPS derived roll component by a factor of two.

6.4.2.2 Comparisons with Attitude Derived from Antenna Positions

In applications where a supplementary navigation device like the AHRS is not available, alternative means are required to make a similar general assessment of GRAPE's performance. One possibility is to derive attitude from the antenna positions, either by performing a least squares adjustment using all available information, or by deriving individual attitude components using selected antenna baselines.

In the following examples, the antenna positions used to derive attitude have been determined relative to the ground reference station using GASP, as described in section 6.3. Usually when determining attitude from GPS data, the baselines between aircraft antennas are estimated, as described in section 6.4.3.2. Figures 6.15 and 6.16 show the pitch and heading values derived from the positions of antennas 1 and 2 during L001. These two antennas are situated on top of the fuselage at the front and back of the aircraft. Although antenna 2 does not lie exactly on the x-axis of the body frame system (see section 5.4.1.1), an offset can be applied to the heading component, and it is also reasonable to attribute all height differences between these two antennas solely to pitch. In both plots, the blue line represents the attitude solution derived from GASP positions and the magenta line shows the GRAPE solution.

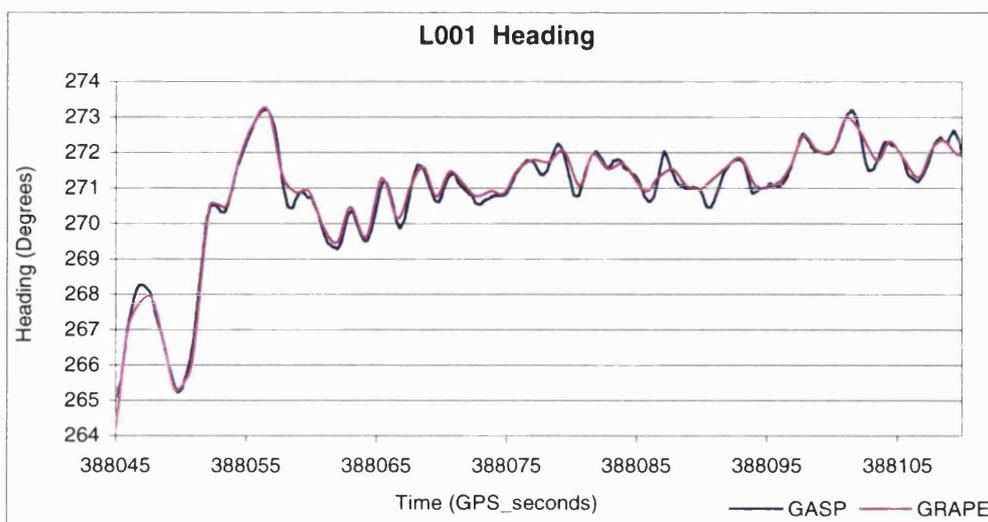


Figure 6.15: Heading derived from GASP antenna positions and GRAPE. L001.

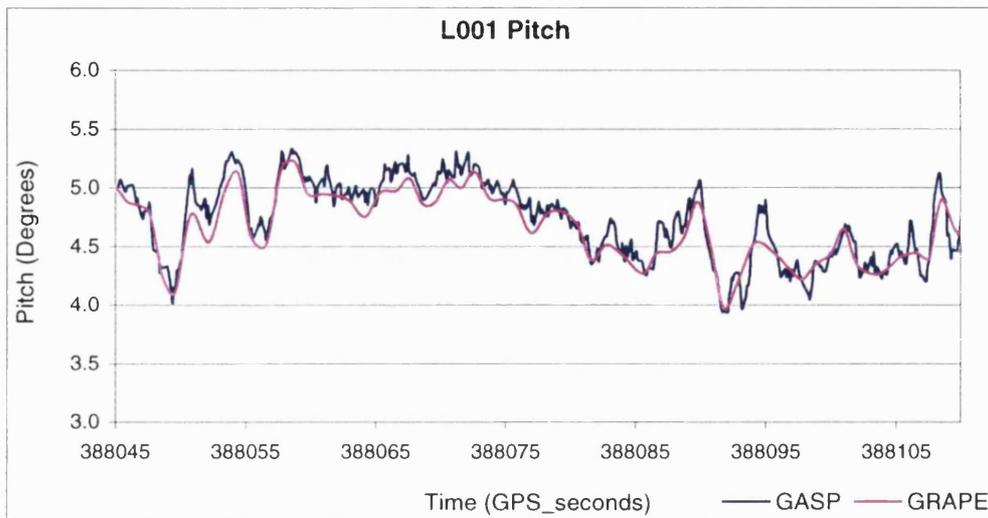


Figure 6.16: Pitch derived from GASP antenna positions and GRAPE. L001.

The agreement between the two solutions provides further evidence that GRAPE is successfully computing reasonable attitude values. These plots also show that the attitude solution derived from GASP positions is detecting high frequency changes, particularly in pitch, that the GRAPE solution does not measure. This is due to the varying availability of GPS data. The GRAPE solution is restricted to a maximum potential sampling rate of 5Hz due to the inclusion of GPS data from the Z-12 receiver connected to antenna 4. The frequency of the final solution is reduced further by data gaps in the measurements from antenna 4 and antenna 1. The two attitude components derived from GASP positions, do not rely on any data from antenna 4 and therefore have the potential to produce a 10Hz solution. In this example the final solution from GASP is actually around 8Hz due to gaps from antenna 1, however this is still significantly higher than the 3Hz GRAPE solution.

6.4.2.3 GRAPE Assessment using Position Trends

A further test of the validity of GRAPE's solution, can be made by checking if trends in position are adequately reflected in attitude. Figures 6.17 and 6.18 show the aircraft height and plan position respectively between 09:45:00 and 09:50:00 on day 289 (16/10/98). This five minute period starts as the aircraft begins to move along the runway before take off, and covers the initial climb to the nominal survey height of 600m. Figures 6.19 to 6.21 show the attitude values determined by GRAPE for this period.

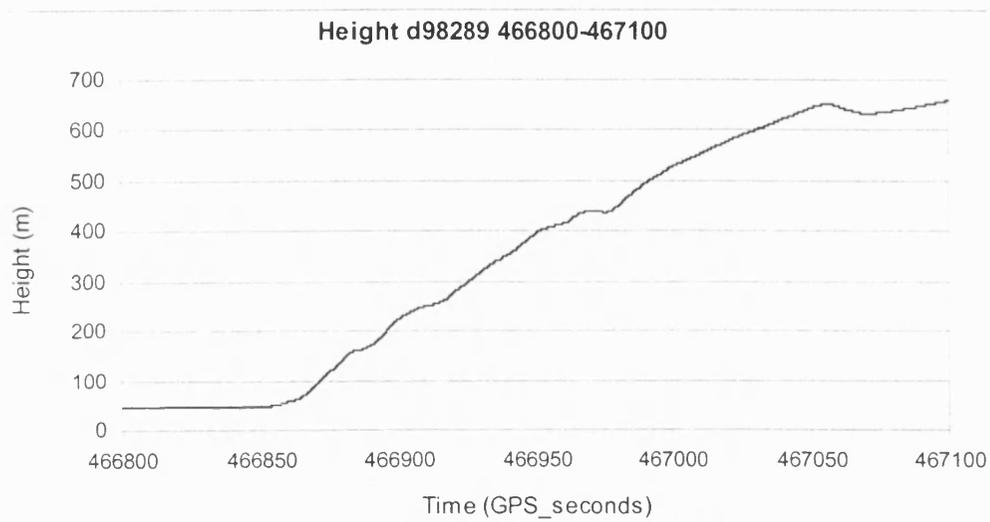


Figure 6.17: Aircraft height. 9:45-9:50, 16/10/98.

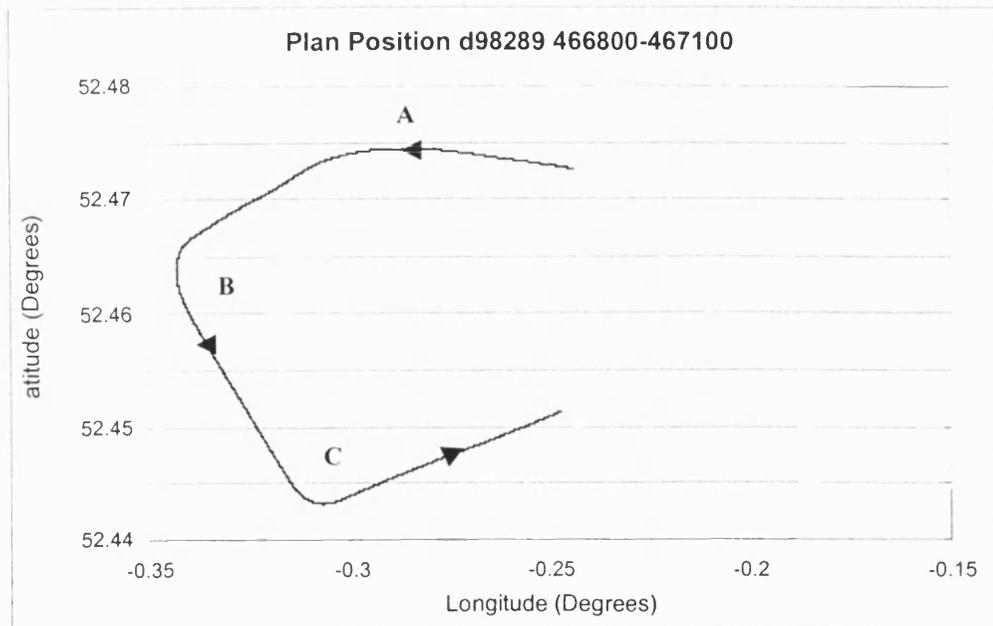


Figure 6.18: Aircraft plan position. 9:45-9:50, 16/10/98.

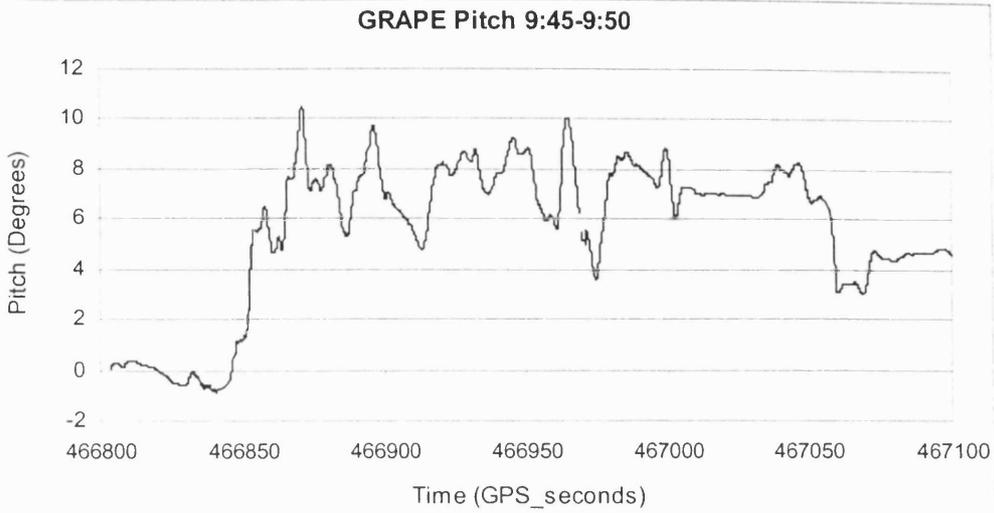


Figure 6.19: Aircraft pitch from GRAPE. 9:45-9:50, 16-10-98.

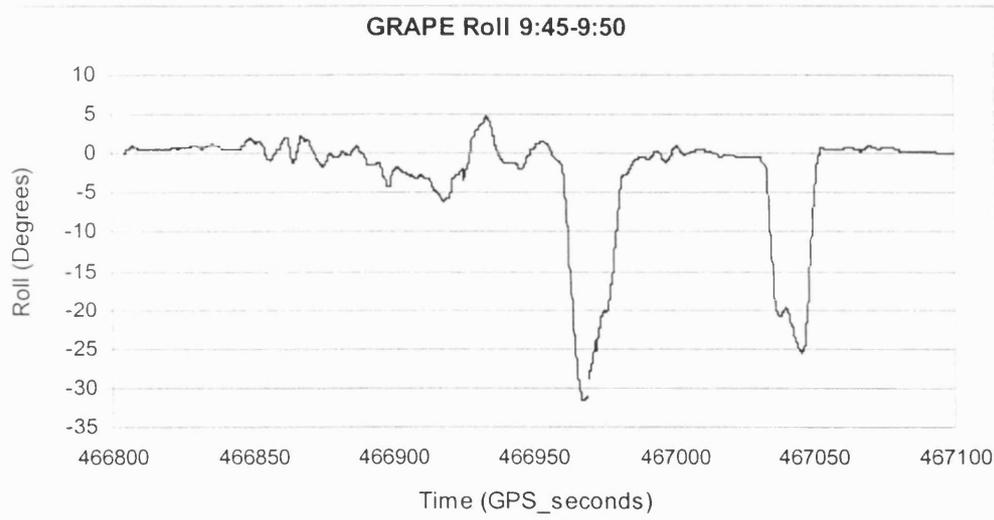


Figure 6.20: Aircraft roll from GRAPE. 9:45-9:50, 16-10-98.

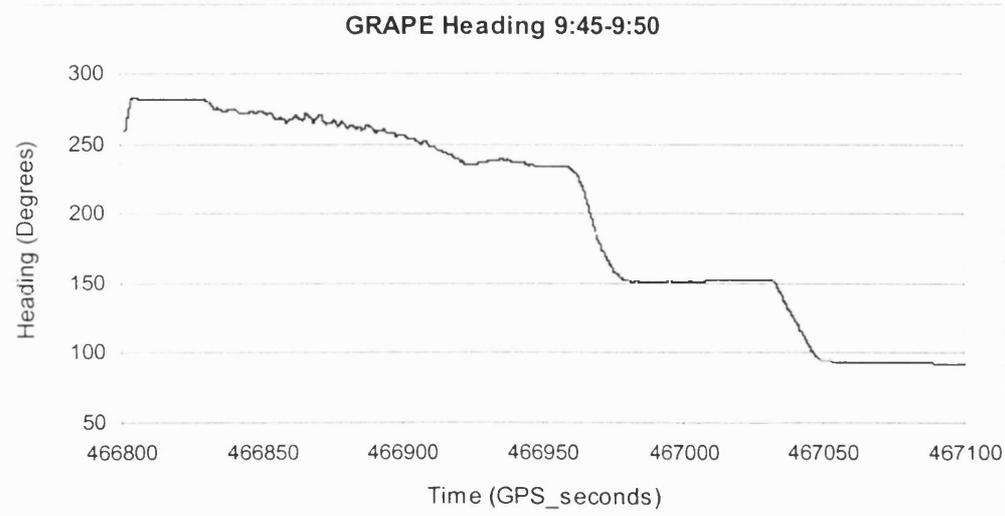


Figure 6.21: Aircraft heading from GRAPE. 9:45-9:50, 16-10-98.

Whilst the aircraft is on the runway, the height remains almost constant (figure 6.17) and the pitch of the aircraft (figure 6.19) remains close to 0° . Once the aircraft has taken off, it climbs steadily and hence the pitch is positive, corresponding to 'nose-up'. The variations in the magnitude of pitch can be directly related to short periods where the rate of climb reduces. During the final 30 seconds of this period the aircraft climbs less steeply and this is reflected in the decrease of pitch to around 5° .

The plan position (figure 6.18) shows distinct phases within the selected period. During the first section, from the start to point 'A', the aircraft is on, and then over, the runway. This is reflected by almost constant roll and heading values (figure 6.20 and 6.21). From point 'A' until turn 'B', the aircraft is gradually turning anti-clockwise, which is reflected by slowly decreasing heading and a negative roll of increasing magnitude, corresponding to 'right wing up'. The motion during turn 'B' is apparent in all components of the GRAPE solution. The roll reaches 32° at the apex of the turn, at which point the number of tracked satellites falls below four and GRAPE is unable to produce a solution for a short period. The change in heading is also clearly recorded. Between turns 'B' and 'C', the aircraft holds an almost constant heading with a steady rate of climb and only a small amount of roll. Motions during turn 'C' are again reflected in all components of the GRAPE solution.

This rather crude check shows that the attitude values from GRAPE show the correct trends. More detailed analysis is required to obtain a more rigorous assessment of the system's performance.

6.4.2.4 Summary of GRAPE Performance from Preliminary Tests

The comparisons of the GRAPE and AHRS solutions indicate that GRAPE appears to be producing reasonable results from the airborne antenna data. Further checks can be made on the general quality of the GRAPE solution by comparing it to an attitude solution derived from GPS antenna positions, or by checking if patterns in the aircraft position are adequately reflected in the attitude components. In this application, where an alternative source of attitude data is available from the AHRS, these additional tests are not strictly necessary. If however any similar work was undertaken in which GPS data were processed using GRAPE, and no external attitude solution was available, these methods would be of use.

The absolute accuracy of the AHRS solution is not sufficient to calibrate the GRAPE solution, therefore further investigations were conducted to more rigorously assess GRAPE performance. In the following section a method to validate the ambiguity resolution is described and applied. In section 6.4.3 an assessment is made of the likely accuracy of the attitude values produced by GRAPE in the current application.

6.4.3 Ambiguity Resolution

6.4.3.1 Ambiguity Validation

GRAPE must correctly resolve ambiguities in order to produce a suitably accurate attitude solution directly from GPS observations. Validating the success of ambiguity resolution is therefore the primary concern when assessing the performance of GRAPE.

The method used to validate ambiguity values from GRAPE is relatively simple in principle, it relies on establishing a set of ‘truth’ ambiguity values using GASP positioning. The procedure can be summarised as follows:

- 1) Position one aircraft antenna either relative to the local reference station, or simply from a single point code solution, at an epoch.
- 2) Use this antenna position as a fixed reference and position one or more additional aircraft antennas relative to this reference.
- 3) Compare the antenna baseline length from this relative positioning to the *a priori* vector length.
- 4) If the computed length agrees with the *a priori* value to within a certain tolerance (chosen as 20mm), then the ambiguities can be considered to be correct⁶.
- 5) These ambiguities are then used as a ‘truth’ set for GRAPE. If a ‘truth’ ambiguity value can be established for each double-differenced phase observation used in the GRAPE solution then every ambiguity value can be checked at each epoch.

⁶ By combining two or more vectors, the possibility that an antenna was inaccurately positioned using incorrect ambiguities but the vectors were still within tolerance, can be ruled out.

To build up this ‘truth’ set it is not necessary to successfully resolve ambiguities using GASP for each baseline at every epoch. As no cycle slips or losses of lock occur during the flight lines, the correct ambiguities established at one epoch can be used for a complete flight line. Comparing the attitude solution from GRAPE with alternative solutions, as described in section 6.4.2, allows the consistency of the ambiguity resolution to be assessed, but establishing the correct double differenced ambiguities and then using them as ‘truth’ values is a more rigorous approach.

By determining the correct ambiguities in this way, and comparing them with the values determined by GRAPE, it was found that ambiguities were being correctly resolved at all epochs on all six flight lines.

6.4.3.2 Single Baseline vs. Direct Attitude Ambiguity Resolution

One of the reasons for developing the direct attitude algorithm described in section 3.5, was that it should allow a more robust solution than the baseline-by-baseline approach (section 3.4.2.1), i.e. it should lead to more successful ambiguity resolution. This improved robustness results from the additional redundancy in the solution provided by directly estimating attitude components from all the available independent double-differenced phase observations, formed from all satellites observed at each antenna. The correlation between the redundancy in the solution, and the robustness of the single epoch ambiguity resolution method, has been demonstrated previously in the context of position determination (Corbett, 1994; Corbett and Cross, 1995; Al-Haifi, 1996, and Al-Haifi *et al.*, 1998). Later, Mahmud (1999) showed that the additional redundancy of the direct attitude algorithm made GRAPE more robust than a single baseline positioning algorithm, in this case GASP. The results of these tests have been summarised in section 5.4.4. A similar investigation has been carried out using the airborne data sets recorded for the current project.

In section 6.3.2, the robustness of the single epoch positioning algorithm implemented in GASP was assessed when fixing the positions of aircraft antennas relative to the local reference station. It was noted however, that in general it would not be possible to replicate these results when the aircraft is operating at greater distances from the reference station, typically beyond 10 or 15km. In these circumstances attitude could be determined using a baseline-by-baseline method by positioning one aircraft antenna at an epoch and using this ‘master’ antenna as a reference to determine the baselines to other aircraft antennas. This is the method adopted in the standard procedure for baseline-by-baseline GPS attitude determination, as described in section 3.4.2.1. and is the method used to find the correct ambiguities used in section 6.4.3.1.

The robustness of the direct attitude system compared to that of a system using single baseline estimation, was tested on all six flight lines. After a lengthy process of iteratively changing various GASP processing parameters, processing the GPS data, and then inspecting the results, it was found that this single baseline approach was able to correctly resolve ambiguities on all the aircraft antenna baselines for the duration of each flight line. In order to achieve this success however, noisy observations from particular satellites had to be removed from the solution. The direct attitude determination algorithm used by GRAPE can successfully resolve ambiguities when these noisy observations are included in the solution, hence it provides a more robust solution. It was also found that even when GASP had successfully resolved ambiguities on a single baseline it was often not possible to identify the correct solution on the basis of the F-test, described in sections 5.3.2.1 and 5.4.1.7. Some of the significant results from these investigations are now presented.

i) Robustness Comparison

In section 6.3.3 it was shown that noisy measurements recorded at antennas 2 and 4 from satellite PRN 4, had prevented successful ambiguity resolution when aircraft antennas were positioned relative to the ground reference station during flight lines L006 and L008. Similar tests were carried out using data from the same period, but this time estimating baselines between the aircraft antennas. Table 6.4 shows the percentage of epochs at which ambiguities were successfully resolved for each baseline on both flight lines, with and without PRN 4 in the solution.

Flight Line	Baseline	Ambiguity Resolution Success (%)	
		With PRN 4	Without PRN 4
L006	A1-A2	100	100
	A1-A4	0	97.3
	A2-A4	0	100
L008	A1-A2	98.2	100
	A1-A4	0	100
	A2-A4	0	100

Table 6.4: Ambiguity resolution success using a single baseline solution.

With observations from PRN 4 included in the solution, it is not possible to resolve ambiguities on the baseline between antennas 1 and 4, and the baseline between antennas 2 and 4. The inclusion of PRN 4 does not only lead to incorrect ambiguities for double difference observations including PRN 4⁷, the level of noise introduced into the solution also prevents a number of other ambiguity values being correctly determined. Once observations from PRN 4 are removed from the solution, ambiguities can be successfully resolved at almost every epoch from both flight lines.

⁷ The correct ambiguities for double differenced observation involving PRN 4 were determined by using the GRAPE values in GASP and then checking the resultant solution using vector lengths.

The phase residuals described in appendix C, and referred to in section 6.3.2.3, clearly demonstrate the larger than average noise in phase measurements from PRN 4 at antennas 2 and 4. Measurements from antenna 1 are not subject to the same noise levels, and the noise level is lower at antenna 2 than it is at antenna 4, hence, ambiguities can be resolved for the baseline between antennas 1 and 2, even when PRN 4 observations are used.

These results again show the detrimental affect on ambiguity resolution of specific noisy measurements, illustrated by both the improved performance when PRN 4 observations are removed, and also the variation in performance over different baselines when PRN 4 observations are included in the solution.

GRAPE, using the direct attitude algorithm, with its increased redundancy, is able to resolve ambiguities when the measurement noise reaches levels that prevent successful ambiguity resolution using a single baseline method. Specifically in this example, GRAPE can resolve the ambiguities for double difference observations between antennas 2 and 4 when observations from PRN 4 are included in the solution.

ii) Confirmation of Correct Ambiguities

It can also be shown that even when GASP can successfully resolve ambiguities it may be unable to confirm that the chosen solution is in fact correct. If, after the initial ambiguity search, more than one trial position is passed forward to the least squares estimation, an F-test, as described in 5.4.1.7, is used to determine whether the chosen solution is significantly more likely to be correct than an alternative solution using different ambiguities. The unit variances of the possible solutions are used as the test criteria. Even though the chosen solution with the smallest sum of the squares of the weighted residuals ($\hat{v}^T W \hat{v}$) is correct, statistically one solution may not be distinguished from another. Table 6.5 shows the percentage of epochs during each flight line at which one, two, three or more than three possible solutions pass the F-test.

Flight Line	Baseline	Number of possible solutions passing the F-test, (%)			
		1	2	3	3+
L006	A1-A4	30.0	36.2	26.2	7.6
	A1-A2	55.0	37.7	7.7	1.7
	A2-A4	60.0	27.5	12.5	0
L008	A1-A2	100	0	0	0
	A1-A4	89.5	7	3.5	0
	A2-A4	98.3	1.7	0	0

Table 6.5: Number of possible solutions following the GASP F-test.

In the worst case, that of the baseline formed by antennas 1 and 4 during L006, the correct solution can only be distinguished statistically from other candidates for 30% of the time. Previous tests of the GASP software have demonstrated that not only does the ambiguity resolution improve as the number of observations increases, so too does the ability to distinguish the correct solution (Corbett, 1994; Al-Haifi, 1996). The larger number of observations in the direct attitude system relative to the single baseline system, allows GRAPE to distinguish the correct solution as statistically more likely than any other, at all epochs, on all flight lines.

6.4.4. Accuracy Assessment

Whilst the comparisons of the attitude values from GRAPE and AHRS in section 6.4.2 showed the general agreement of the solutions, it was noted that AHRS values are not sufficiently accurate to fully calibrate the GRAPE solution. The absolute accuracy of the GRAPE solution being tested cannot be assessed beyond the 1° to 2° absolute accuracy of the AHRS reference solution. The agreement of the attitude *trends* measured by the two solutions can perhaps be used to calibrate the relative accuracy of the GRAPE solution to a higher degree, as the resolution of the AHRS solution is significantly higher than its absolute accuracy. Furthermore, as the GPS based solution is not subject to the same time dependent drifts as the AHRS, the absolute accuracy of this solution should be similar to its relative accuracy in this situation. Accepting that this relative agreement in the two solutions can be used to assess the absolute accuracy of the GRAPE solution to some degree, the resultant accuracies are approximately 5, 20 and 2 arc minutes in pitch, roll and heading respectively.

Further evidence from additional sources is required to make a more comprehensive assessment of the system's accuracy. The internal quality measures produced by GRAPE have not been used to assess its performance as they are dependent on models which are similar to those used to determine the attitude parameters themselves. Instead, methods that make use of some external form of calibration have been employed.

Section 6.4.4.1 describes an assessment of GRAPE attitude accuracy using data gathered whilst the aircraft remained stationary with the engines running. This assessment assumes that any recorded attitude changes are attributable to errors in the measurement system not actual changes in the aircraft attitude. In section 6.4.4.2, the accuracy estimates from baseline positioning, using GPS data collected during the flight lines, are used to infer the attitude accuracies that should be achieved using the same data. These values are summarised in section 6.4.4.3, and an overall estimate of the accuracy is made based on all the available information.

6.4.4.1 Static Testing

The performance of GRAPE was originally assessed by Mahmud (1999) using a static data set as described in section 5.4.4. The ‘truth’ attitude values were determined by computing the average position for each antenna in the array using the complete two hour data set, and then deriving a body frame system such that the ‘true’ pitch and roll values were zero and the heading value was a non-zero constant. The same data that were used to compute the antenna positions were then processed with GRAPE, and the solution was compared to these ‘truth’ values. Any deviation of the computed attitude value at each epoch from the ‘truth’ value was then considered to be the error in the GRAPE solution. If similar results were achieved using the current antenna array with its much shorter baseline lengths, 1.2 to 3.5m instead of 20 to 40m, the expected accuracies would be approximately 6 arc minutes in pitch, 20 arc minutes in roll, and 1.5 arc minutes in heading (σ).

A similar investigation has been conducted in this research using a three minute period of static data collected on all four aircraft antennas whilst the aircraft was stationary with the engines running. In this test the ‘truth’ solution used was the average GRAPE attitude value for each component over the observation period. Again, deviations from this truth at an epoch are assumed to be the system errors and not actual motions of the aircraft. Figure 6.22 to 6.24 show the magnitude of the pitch, roll and yaw errors respectively over this three minute static period. Table 6.6 summarises some important characteristics from these results.

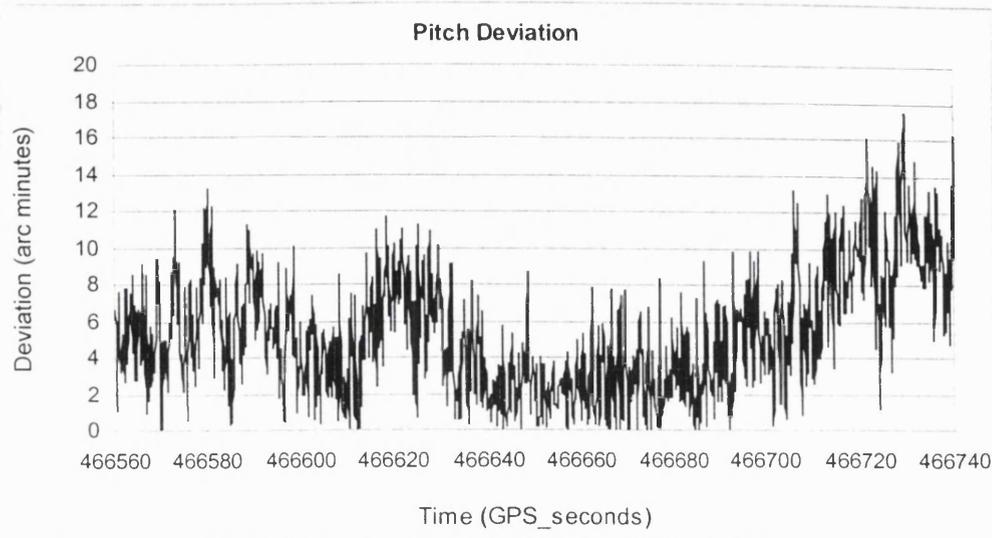


Figure 6.22: Deviation in the GRAPE pitch solution from a three minute static period.

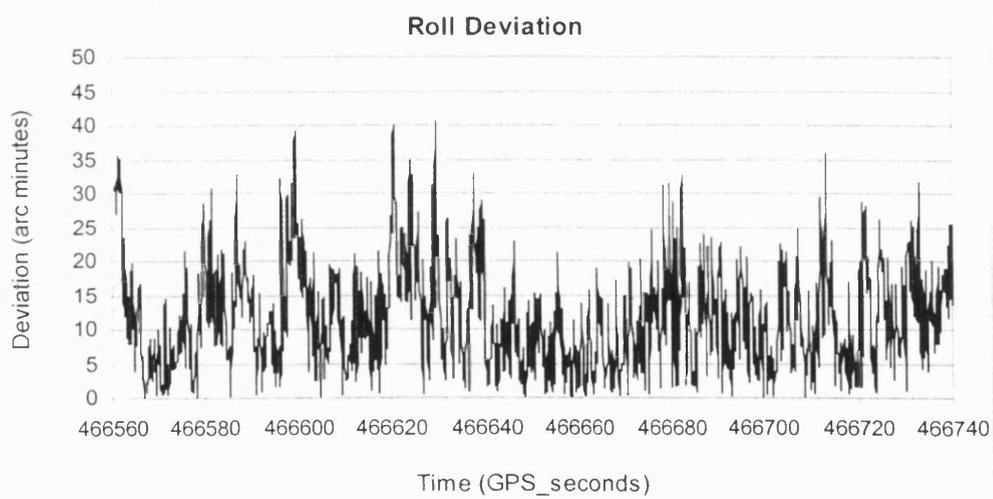


Figure 6.23: Deviation in the GRAPE roll solution from a three minute static period.

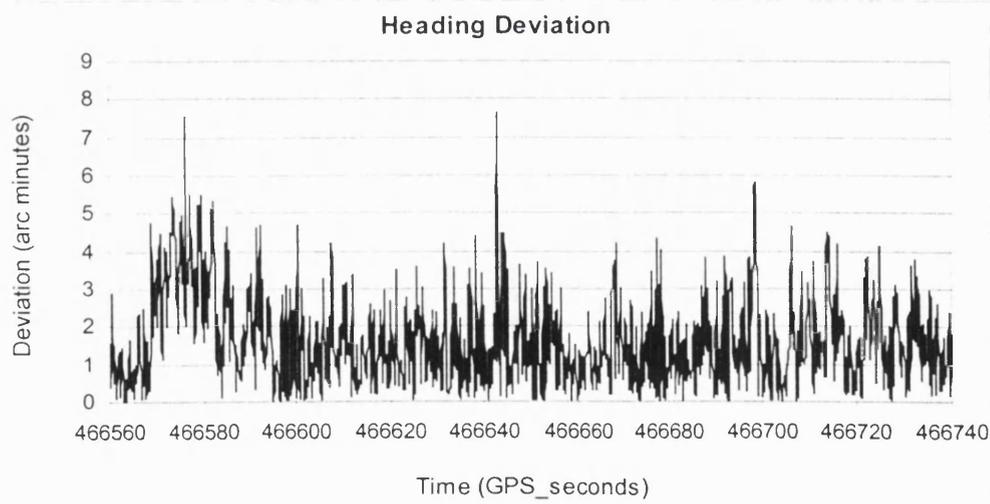


Figure 6.24: Deviation in the GRAPE heading solution from a three minute static period.

	Pitch	Roll	Heading
Mean	7.0'	50.3'	200° 57.3'
Minimum	-10.6'	14.9'	200° 49.7'
Maximum	20.2'	1° 30.8'	201° 05.0'
Range	30.8'	1° 15.9'	15.3'
Deviations			
σ (68%)	6.8'	14.6'	2.0'
95%	11.3'	25.9'	4.0'
100%	17.5'	40.5'	7.7'

Table 6.6: GRAPE attitude accuracy using static data.

The standard deviations from this test correspond closely to the predicted accuracies based on previous tests (Mahmud, 1999). It appears that any increase in noise due to increased levels of multipath in the current observation environment has not significantly altered the accuracy of the solution. Although these accuracy values refer specifically to static data, they may also be applicable for airborne periods. During flight lines where the aircraft maintains a regular speed and heading, and any tilts (pitch and roll) are minimised, the observation conditions will be fairly similar to these static tests. The reflecting surfaces around the antennas will cause similar levels of multipath whether the aircraft is stationary or moving. This static test however, cannot replicate the full range of satellite-antenna geometries, and the associated multipath, that could be experienced during flight lines.

6.4.4.2 Attitude Accuracy based on Positioning Accuracy

The potential accuracy of the GRAPE attitude solution can also be estimated based on the accuracy with which baselines can be determined from the same GPS data, using standard positioning techniques. Baseline lengths between aircraft antennas have been determined by positioning one antenna relative to a temporary reference antenna on the aircraft, using GASP. The co-ordinates of this reference antenna are calculated during the RINEX to NXF conversion using a single point, single epoch, code position, as described in section 5.2. Figures

6.25 and 6.26 show the computed baseline lengths between antennas 1 and 2, and antennas 3 and 4 respectively, during flight line L013. Results from estimating the complete set baselines between aircraft antennas for this flight line are given in table 6.7.

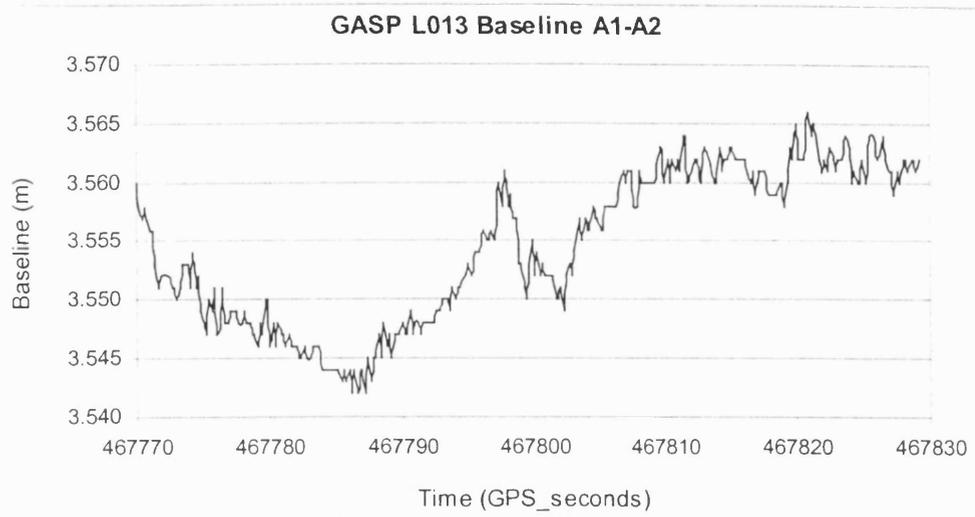


Figure 6.25: Antenna 1 to antenna 2 baseline length using GASP, L013.

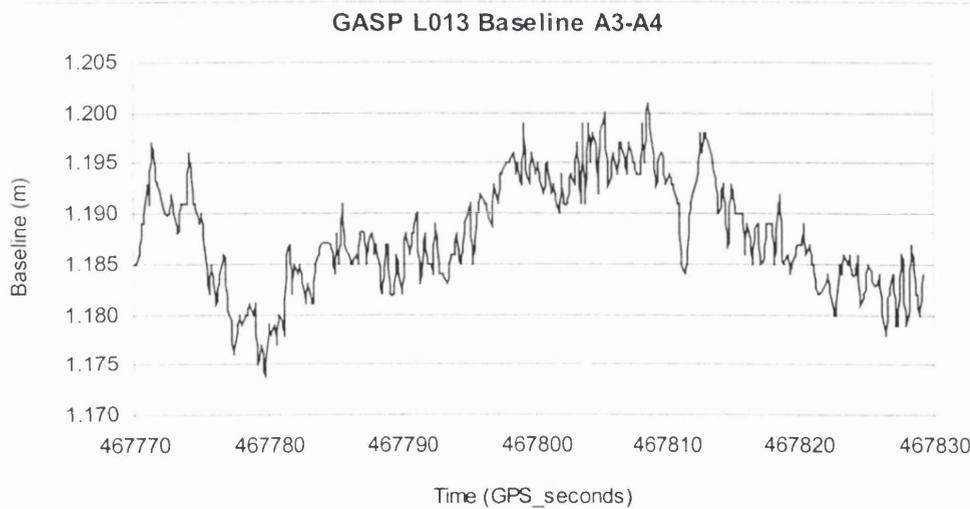


Figure 6.26: Antenna 3 to antenna 4 baseline length using GASP, L013.

Baseline	A1-A2	A1-A3	A1-A4	A2-A3	A2-A4	A3-A4
Max (m)	3.556	2.137	2.121	1.616	1.672	1.201
Min (m)	3.542	2.126	2.130	1.596	1.640	1.174
Range (m)	0.024	0.011	0.009	0.020	0.032	0.027
Mean (m)	3.555	2.134	2.125	1.606	1.656	1.188
σ (m)	0.007	0.003	0.002	0.006	0.009	0.006

Table 6.7: Summary statistics from all baseline estimations during L013.

The first point to note is the improvement in positioning performance when antennas are positioned relative to an aircraft reference antenna instead of the ground reference station. The results presented in section 6.3.3 show significant problems in resolving ambiguities when antennas 1 and 4 were positioned relative to the ground reference station. When these antennas are positioned relative to other aircraft antennas, ambiguities can be resolved at all epochs for each baseline. Using a reference antenna on the aircraft instead of the ground station reduces the baseline lengths from a few kilometres to a few metres.

The standard deviations in the baseline lengths vary between 2 and 9mm for these six baselines. It is difficult to determine exactly how height and plan errors from a single baseline will propagate through to errors in each attitude component. When the direct attitude algorithm is used with observations from an array of antennas, errors may tend to cancel out, or they may combine.

Using the following simple assumptions, the accuracy of each attitude component has been estimated.

- 1) The error in height is approximately twice the plan error.
- 2) The pitch error is largely dependent on the accuracy with which the height difference is determined over a fixed 3.55m baseline.
- 3) The roll error is largely dependent on the accuracy with which the height difference is determined over a fixed 1.2m baseline.
- 4) The heading error is largely dependent on the accuracy with which the plan difference is determined over a fixed 3.55m baseline.

Table 6.8 shows the predicted accuracies for each attitude component, based on these assumptions, using standard deviations in position of 5 and 10mm.

	Position Error σ	
	5mm	10mm
Pitch	4' 21"	8' 43"
Roll	13' 00"	25' 59"
Heading	2' 26"	4' 51"

Table 6.8: Attitude errors predicted from positioning errors.

6.4.4.3 Attitude Accuracy Summary

None of the methods described is able to provide a complete assessment of the accuracy of the GRAPE solution. The limitations of each approach have been discussed in the appropriate sections. Taken together, the figures determined using these methods, do provide a reasonable assessment of the system performance. Table 6.9 summarises the predicted accuracy for each attitude component based on these investigations. In the final column an overall figure is given based on all of the available evidence. This constitutes the best accuracy guide in the absence of a reference solution from the ADU2 or scanner imagery. All figures are given in minutes of arc.

	AHRS Comparison	Static Test	Positioning Accuracy	Combined
Pitch	6	6.8	4.5-9	7
Roll	20	14.6	13-26	20
Heading	1.5	2	2.5-5	3

Table 6.9: Attitude accuracy from all available sources.

These accuracies of 7' in pitch, 20' in roll and 3' in heading, are almost identical to those of the ADU2, the dedicated GPS attitude determination instrument installed on the survey aircraft (section 4.2.4). They also agree with the typical accuracy figures for GPS based attitude systems documented in previous research (section 2.3.3.2). More importantly perhaps, these figures indicate that the system is able to meet the attitude accuracy requirements of 5 to 10' in pitch, 10 to 20' in roll and 3' to 5' in heading stated in section 4.4. The relative accuracies of the three attitude components were as expected, with the limiting factor on the overall performance of the system being the error in roll.

6.5 Recording Rates

The GPS based direct georeferencing system developed in the course of this research has been assessed in terms of its robustness and accuracy for determining position and attitude values. The rate at which these values are measured must also be considered when assessing the system's ability to provide sufficiently accurate orientation parameters for each scan line. Some interpolation will clearly be required to relate navigation data to imagery as the ATM scanner can record up to 50 lines per second whilst the GPS solution is only available at 10Hz, at best. The recording rates are discussed separately for position and attitude as the requirements and available data vary.

6.5.1 Position Rate

Kusevic and Mrstik (1997) investigated the frequency with which GPS data needed to be recorded in order to position the perspective centre of an aerial camera within a given error budget of 5cm. They concluded that with a relatively simple interpolation procedure using only the GPS epochs immediately before and after the instant of exposure, this accuracy requirement could be met with a 1Hz solution. The accuracy requirements for the current application are considerably lower than those of a photogrammetric survey, suggesting that a 1Hz solution should be more than adequate.

Figure 6.27 shows the height of antenna 2 during flight line L001, determined using GASP. The complete 10Hz solution is shown as a series of small blue diamonds, with a decimated 1Hz solution represented by larger magenta squares. At this scale it appears that the 1Hz solution can describe the aircraft trajectory quite precisely, and that a simple interpolation procedure would probably be adequate.

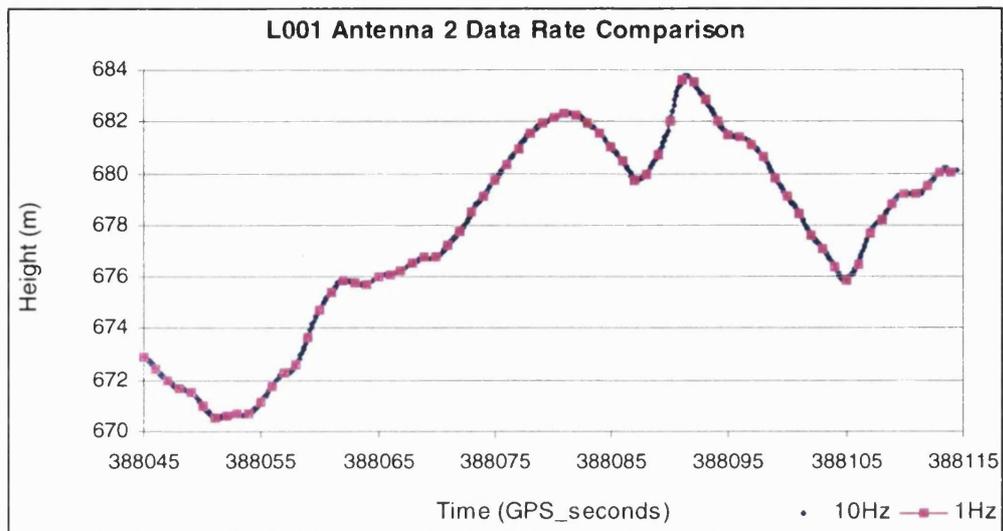


Figure 6.27: Antenna 2 height on L001, using a 10Hz and 1Hz solution.

Figure 6.28 shows a 20 second period of the data-set used in figure 6.27 at a larger scale. The height values for periods between the 1Hz GPS epochs have been determined using a simple linear interpolation between adjacent recorded heights.

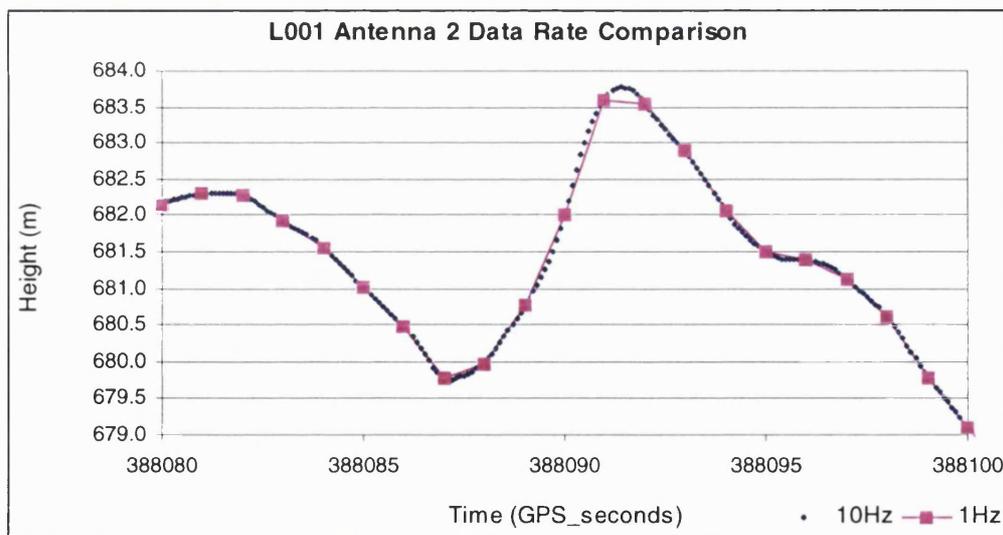


Figure 6.28: Antenna 2 height for 20 seconds of L001, using a 10Hz and 1Hz solution.

Generally the heights interpolated from a 1Hz solution agree to within a few centimetres with the measured 10Hz values. There is a period however, between epochs 388091 and 388092 where the interpolated values differ from the measured values by up to 25cm. The error introduced by the interpolation procedure can be reduced by using a more refined non-linear interpolation algorithm or by processing the data with a filter. The possibility of using a lower

data rate and applying a suitable interpolation procedure is of practical interest as it allows a smaller GPS data set to be recorded and processed.

To date, no further investigations of interpolation procedures have been conducted using the current GPS data. The first reason for this, is that as part of the original processing strategy, the positions from GASP will be filtered using a Kalman filter. This should assign usable positions for each scan line between measurement epochs. Secondly, although the aircraft antenna positions can be positioned with centimetre level accuracy, the overall error budget of the direct georeferencing system is dominated by orientation errors. The on-the-ground positioning accuracy will be at the level of a few metres (see section 6.4.3), hence any increase in the positioning accuracy of a few centimetres due to improvements in the interpolation process is likely to have a negligible effect on the total error budget.

6.5.1 Attitude Rate

Prior to the test flight it was intended that gyro data from the AHRS would be integrated with GPS data to provide a higher frequency attitude solution. In effect, the AHRS values would be used as an interpolator between GPS epochs. However, the preliminary comparisons of the two solutions, described in section 6.4.2.1, have shown that in general the AHRS solution does not appear to detect high frequency motions that a GPS solution misses due to its lower recording rate. In fact, when a 10Hz GPS attitude solution is available⁸, it appears to record motions that the AHRS solution does not.

Figure 6.29 shows the GRAPE and AHRS solutions for the roll component, during a 15 second period of L013, when a 10Hz GPS attitude solution is available. Individual GPS values are represented by blue diamonds, and the magenta squares represent AHRS values.

⁸ A 10Hz solution is only available on L013 when an almost complete data set is recorded from antenna 3.

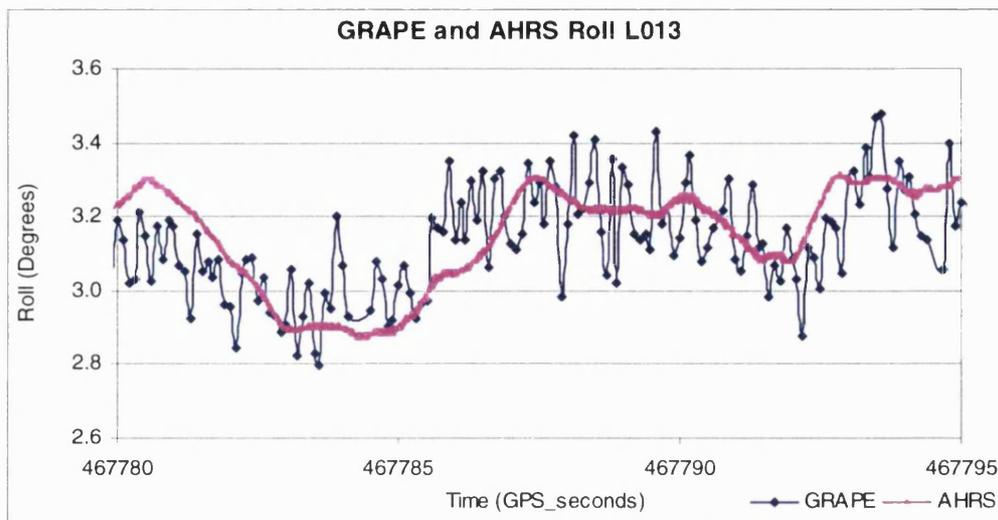


Figure 6.29: GRAPE and AHRS roll. L013.

The roll solution from GRAPE exhibits an almost regular cyclic pattern with a frequency of 5Hz and an amplitude of around 6 arc minutes. The pitch and heading components exhibit a similar pattern but with an amplitude generally below 1 arc minute. Analysis of the height difference between antennas 2 and 3 has shown a corresponding pattern, i.e. a cyclic variation of around 2mm. For all three components the AHRS does not record a similar motion, there are a number of possible reasons for this:

- 1) The high frequency motion recorded by the GPS antenna array may be the result of an airframe vibration that is not experienced by the AHRS, due to the damping devices attached to the ATM scanner on which it is mounted. If this is the case, the ATM scanner will also not be subject to this vibration. However, the CASI scanner, which is rigidly fixed to the airframe with no damping devices, may experience similar attitude changes to the antenna array, hence this high rate solution may be required to more accurately georeference this imagery.

- 2) The AHRS and ATM may experience this vibration, but due to a smoothing algorithm applied within the AHRS⁹, the recorded solution may have smoothed out a genuine motion that was measured at the full 64Hz rate. If this is the case, the ATM and CASI imagery, would be subject to these vibrations, and therefore would not be georeferenced as accurately using the AHRS solution.
- 3) This trend could possibly be a result of a systematic error in the GPS data. However, there does not appear to be an obvious GPS error source that would account for this high frequency regular pattern.

Without scanner imagery it is not possible to confirm whether this trend is experienced by either of the two scanners, but it is possible that the 10Hz GPS solution is able to detect significant high frequency changes in attitude that the AHRS cannot.

A similar test was conducted to establish how well a 2Hz GPS solution (the recording rate of the Ashtech ADU2, the standard positioning and attitude instrument on the aircraft) can record the movements experienced by the aircraft. In figures 6.30 and 6.31, a 10Hz attitude solution (represented by blue diamonds) has been compared to a 2Hz solution (represented by red circles), both from GRAPE.

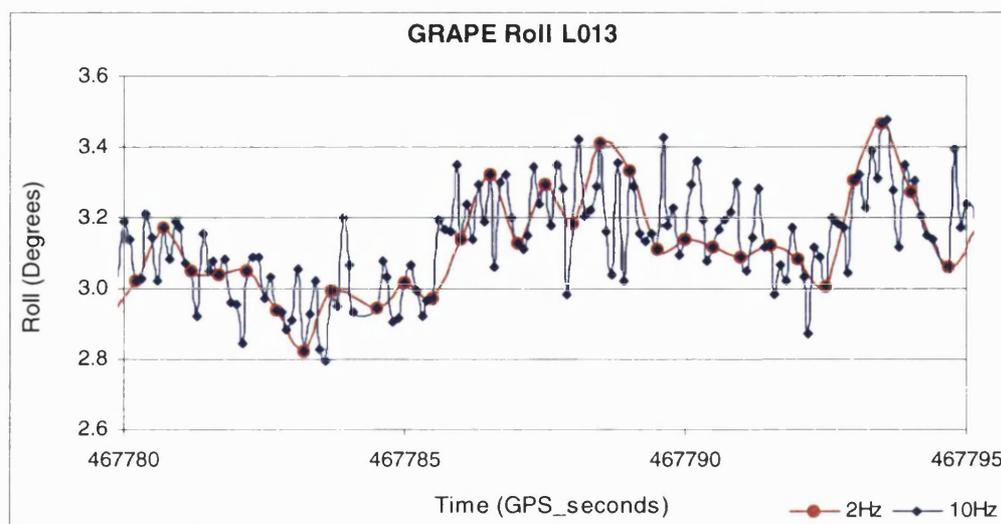


Figure 6.30: GRAPE roll. 2Hz and 10Hz solutions.

⁹ To date, no details of this filter have been made available, so it is unclear what the effective sampling rate is.

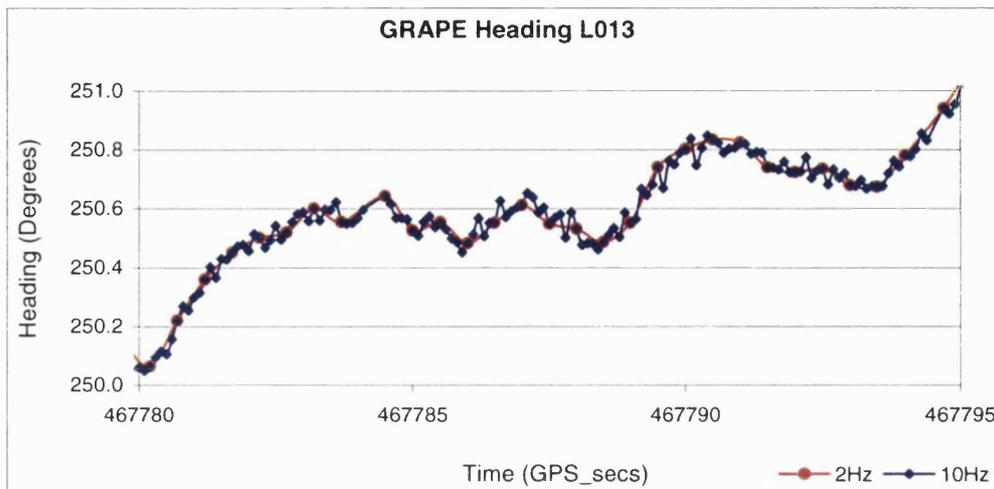


Figure 6.31: GRAPE heading. 2Hz and 10Hz solutions.

Figure 6.30 shows the same high frequency variation in the roll component noted above. Clearly the 2Hz solution is not recording these motions, nor could motions of this nature be interpolated using the attitude values available from a 2Hz solution. If these motions are experienced by the scanners, it is necessary to have attitude values available at a frequency of 10Hz or higher. The heading (figure 6.31) and pitch components are less affected by these high frequency changes, and a 2Hz solution may be adequate for these components.

6.5.3 Summary of Recording Rates

- A 1Hz GPS solution appears to be adequate to transfer centimeter level positioning accuracy from GPS measurement epochs to each scan line. Given the total error budget in this application, a lower frequency solution with a suitable interpolation procedure might also be a practical option.
- A 5Hz GPS solution is generally able to record all the changes in attitude that the AHRS measures.
- It is possible that the 10Hz GPS solution from independent receivers has recorded motions that the AHRS and ADU2 do not record. The ability of the georeferencing solution to detect this motion only becomes a benefit however, if this same motion is experienced by either of the scanning devices. To date it has not been possible to determine whether or not this is the case.

6.6 Concluding Remarks

In this chapter the performance of the direct georeferencing system that has been developed for the current research has been described. Position and attitude determination have been assessed in terms of robustness and accuracy, and the rate at which these solutions must be recorded has also been discussed. It was noted throughout, that the methods available for calibrating performance have not been as rigorous as was originally intended. Despite the limitations of the validation methods, the following conclusions can still be drawn:

- A fixed ambiguity, carrier phase solution can be used to position each aircraft antenna, relative to the local reference station, at every epoch during the selected flight lines with the multiple epoch ambiguity resolution method used in SKI.
- Using the GASP single epoch ambiguity resolution technique, ambiguities can be resolved between the local reference station and the aircraft antennas for the majority of epochs.
- The accuracy of the positions determined in this application, from either solution, is generally between 5 and 10mm in plan, and 10 to 30mm in height.
- The attitude solution produced when GPS data from the four independent aircraft receivers are processed using the direct attitude method, closely matches the attitude solution from the gyro-based AHRS, and an attitude solution derived from antenna positions. The attitude solution also reflects position trends well.
- Ambiguities can be successfully resolved at every epoch during the selected flight lines using the direct attitude method. In order for a single baseline positioning method to replicate this success, specific noisy measurements, almost certainly caused by multipath, must be removed from the solution. The direct attitude method is therefore more robust than a single baseline method.

- The accuracy of the attitude solution is approximately 7 arc minutes in the pitch component, 20 arc minutes in roll, and 3 arc minutes in heading.
- Whilst a GPS solution of 1Hz or lower may be adequate to provide positions for the scanners, a higher recording rate appears necessary for the attitude components. The exact rate required cannot be confirmed without scanner imagery, but it may be as high as 10Hz.
- The direct georeferencing system developed in this research has the potential to position imaged features in a mapping co-ordinate system with an accuracy of around 4m at a flying height of 600m. It therefore meets the objectives of this project, and would be a suitable means of providing georeferencing information for the majority of tasks undertaken by the Natural Environment Research Council's Airborne Remote Sensing Facility.

Chapter 7 Integration and Filtering

7.1 Introduction

Although the results presented in chapter 6 have shown that position and attitude solutions, derived entirely from GPS data, meet the requirements of this project, there are still benefits that can be realised by data integration and filtering. This chapter discusses ways in which attitude data from different sources can be integrated and filtered to improve robustness, accuracy and processing efficiency. The positions determined from GPS data can also be filtered, but this does not significantly improve the results using the current data-set.

Section 7.2 describes a method to improve the ambiguity resolution procedure used in GRAPE, the direct GPS attitude determination program, by using attitude values produced by the gyro-based AHRS instrument. This integrated approach has benefits in terms of processing efficiency, and could potentially make the GPS processing method more robust.

In section 7.3, the use of Kalman filters to improve the position and attitude solutions is discussed. For the position solution, Kalman filtering does not significantly change the final values, it does however, provide a convenient means of detecting outlying data and can be used as an interpolator to assign positions to scan lines. By combining data from various sources a Kalman filter does provide important benefits for the attitude solution. An integrated GPS/gyro attitude solution can remove the affects of gyro drifts and axes misalignments, and provides orientation for the otherwise arbitrary AHRS heading. The attitude filter has been designed to accept GPS solutions from the specific system developed in this research, i.e. GPS data from independent receivers, processed with GRAPE software, as well as solutions from the ADU2, dedicated GPS attitude unit. It accepts ADU2 values as an input, because it has been adapted to form part of the NERC ARSF's integrated data system (IDS), as described in 4.2.1. In general operation the ADU2 will provide the only attitude solution from GPS data for the survey aircraft. The four independent receivers were installed specifically for the two test flights undertaken in October, 1998 (see section 4.3).

7.2 GPS and Gyro Integration for Ambiguity Resolution

The ambiguity resolution method used in this research computes a set of double differenced ambiguities for each trial attitude within a search volume, centred on a starting approximation (see section 5.4.1). To ensure the search volume contains a trial attitude leading to the correct ambiguities, its size must be related to the accuracy of the starting approximation. When a good starting value is not available, a search volume can be constructed around some arbitrary value (eg. zero in all components) that is sufficiently large to include any possible attitude values that an aircraft could experience, as described in section 5.5.2. This involves computations for a huge number of trial attitudes and is therefore very inefficient. In periods when sufficient good GPS observations are available for a successful solution, final attitude values from the previous epoch can be used as starting approximations, or values can be predicted based on two or more previous epochs. This approach has proved successful in periods of relatively stable attitude, as experienced during the selected flight lines. If however, the change in attitude between two epochs becomes increasingly unpredictable due to greater aircraft dynamics, or a longer time interval between measurements, which may occur when the aircraft turns and some antennas are unable to receive satellite signals for a time, this approach may not be adequate, or may only be successful with a very large search volume. The reliance on values from previous epochs also means that this is not a genuine single epoch solution.

One way to overcome this problem, is to use values from the AHRS solution as starting approximations to resolve ambiguities. It has been found, that provided any axis misalignments and time offsets between the two systems are taken into account, ambiguities can be resolved using a single attitude value, i.e. no searching is required. Float ambiguities are determined for the starting attitude, these are then rounded to integer values, and the least squares process solves for the corrections to the initial values.

Some caution is needed if a single attitude value from AHRS is used to determine ambiguities. It must be possible to identify if inaccurate values, leading to the wrong ambiguities have been used. This situation can be identified by setting threshold levels for the ambiguity function value and the unit variance of the solution. The ambiguities can also be validated by using them in a baseline estimation and checking the resultant vector length, as described in section 5.5.7.2. The dependence of the GRAPE solution on values from AHRS means the preliminary comparison between these two solutions, as described in section 6.4.2.1, is no longer an independent test of GRAPE's performance.

Applying this integrated method for periods when ambiguities can also be resolved with a GPS only solution, which is the case on all the selected flight lines, shows that the ambiguities and final attitude values from the standard and integrated methods are identical. Table 7.1 shows how the processing time of the GRAPE solution is related to the number of trial attitudes that must be considered in the ambiguity search.

Starting Value	Range (radians)	Trial Attitudes	Processing time per epoch
Previous epoch	+/- 0.1	9261	100.0%
Previous epoch	+/- 0.05	1331	7.5%
Previous epoch	+/- 0.02	125	3.6%
AHRS	0	1	1.5%

(Note. The searching increment is fixed at 0.01 radians)

Table 7.1: Search volume and processing time

The actual computation time will depend on the processing speed of a specific computer, therefore the time taken for the largest search volume has been considered as a benchmark. The figures clearly show how the processing time increases with the search volume. Using AHRS data there is only a single trial attitude to consider, making the processing significantly more efficient.

The integrated approach also has the potential to improve the robustness of the single epoch ambiguity resolution method by allowing successful ambiguity resolution at epochs when this is not possible with a standard solution. It is not possible to show this with the current data, as ambiguities can be resolved using standard methods for all periods when the AHRS solution is available. However, in situations when the levels of measurement noise are higher, or when there are fewer available satellites, the standard ambiguity resolution method may be unsuccessful. If an AHRS solution was available at those epochs, ambiguities could be correctly resolved. The outlier rejection routine, described in section 5.4.1.7, could then identify any of the erroneous data that prevented ambiguities being resolved with the standard method, and remove them from the solution before computing the final attitude values. An integrated approach of this kind is robust and reliable, i.e. it can resolve ambiguities and identify outliers in the GPS data.

This integrated method does not rely on any values from previous epochs, therefore it is a truly independent single epoch method for attitude determination. This is consistent with the original single epoch ambiguity resolution philosophy of both GASP and GRAPE. If this integrated ambiguity resolution technique is used in GRAPE, and GASP is used to determine positions, then an independent solution for all six parameters of exterior orientation is derived at every measurement epoch.

7.3 Filtering

7.3.1 Introduction

In section 2.3.3.3 various Kalman filter designs, for integrating GPS and INS data were discussed. When designing these filters, the data available, and the level at which these data should be treated, are major considerations. Generally data are filtered at as low a level as possible, e.g. GPS double differenced pseudorange, carrier phase and phase rate observations, or the corrections to them, form the measurement vector in the filter, rather than derived positions or velocities. In the current research, the overall processing strategy has been based around two single epoch ambiguity resolution programs, GASP and GRAPE, to determine position and attitude respectively, as described in chapter 5. The measurements used in the filter are therefore derived position and attitude values, and not raw measurements, limiting the possible level of processing. When the GPS solution comes from the ADU2, rather than independent receivers, it will also be in the form of derived values and not raw observables.

The level of filtering is further restricted by the absence of any position information sources, other than GPS, i.e. there are no solutions to integrate. In a more comprehensive GPS/INS configuration, the INS comprises gyroscopes and accelerometers, allowing position, as well as attitude information to be derived. The filter then makes use of all the available position information from the GPS and the INS.

For this research, a simple 9 state (3 components of position, velocity and acceleration) filter was designed to process the positions determined by GASP or SKI. However, as the positions have been determined with centimetre level accuracy using a carrier phase solution (section 6.3), the filtered solution shows negligible changes relative to the initial position values. Furthermore, any smoothing¹ of the computed trajectory may remove small changes in position that were actually experienced by the airframe and one, or both, of the scanners. One

¹ The term 'smooth' is used here in its general sense and not to denote that the Kalman filter has necessarily applied smoothing as well as filtering.

benefit of applying the filter, is that by implementing an outlier test, based on the predicted residuals (see section 7.3.2.5), any positions computed at epochs when incorrect ambiguities were determined in GASP, can be identified and removed from the final trajectory². The filter is also a convenient method for assigning positions computed at GPS measurement epochs to scan lines, using synchronised time data, i.e. it can act as an interpolator.

As two attitude solutions, using different data sources (GPS and gyros) are available, a Kalman filter is a useful tool as it allows measurements from different sources to be integrated in an optimal way. It also allows outlier detection routines to be implemented, and can compute attitude values at scan times. For the current application, a filtering program was designed to integrate the attitude solution from the independent GPS receivers with the AHRS angular rates and attitude values. This software was extended during a contract carried out at UCL on behalf of the NERC ARSF to allow additional inputs from the ADU2, dedicated GPS attitude determination unit, and the two scanners. The design of this filter is described in the following section (7.3.2) and results from processing various data with this filter are presented in section 7.3.3.

² This procedure will identify isolated epochs with incorrect ambiguities that lead to spikes in the trajectory. It will not automatically remove a series of positions based on incorrect ambiguities.

7.3.2 Attitude Filter Design

7.3.2.1 Design Overview

Comparisons presented in section 6.4.2.1, showed that in general, the GPS attitude solution determined using the direct attitude algorithm, is able to measure all the changes in attitude that the higher rate AHRS detects. This suggests that the GPS solution for the airframe motion does not require additional input from the gyro data. However, it was also noted that despite the relatively good agreement between the solutions, there were also some significant differences, and one of the reasons for this is that the GPS antenna array and the AHRS are not measuring exactly the same physical changes.

Due to the equipment configuration on the NERC ARSF survey aircraft, it is likely that the AHRS measurements provide a better estimate of the attitude changes experienced by the ATM scanner than the GPS solution does. This is because the AHRS and ATM are rigidly connected, and are less sensitive to the airframe motion due to damping devices. Integration allows the short-term accuracy of the AHRS gyro solution to be used, whilst removing any longer term drifts using GPS attitude data. It is particularly important to have GPS derived heading values to combine with AHRS data as the AHRS heading is aligned to an arbitrary north value for each flight line. A Kalman filter provides a convenient means of combining these two data sources with assumptions regarding the aircraft's dynamics, in order to produce optimal estimates of the exterior orientation parameters for each scan line. In the following discussion, the GPS attitude measurements can either come from the direct attitude algorithm using data from the independent receivers, or from the ADU2 dedicated GPS attitude determination unit. Both solutions use data from the same antenna array, and are of similar accuracy, from the point of view of filtering, they differ only in the recording rate.

As discussed in section 2.3.3.3, a Kalman filter describes the state of a system at an epoch by combining measurements that allow estimations of those states, with predictions of the same states, based on previous information, and knowledge of the system's behaviour over time. In order to do this, some assumptions need to be made regarding the behaviour of the system. In reality these assumptions will not always be correct but any uncertainty can be incorporated into the filter to some degree by applying appropriate weighting. For the current application the following assumptions have been made:

- 1) The airframe acts as a rigid body.
i.e. there is no flexure
- 2) The GPS antennas and the CASI scanner are rigidly fixed to the airframe.
i.e. they experience the same attitude changes
- 3) The AHRS is rigidly attached to the ATM.
- 4) The offsets between the airframe and AHRS axes vary only by small and random amounts.
- 5) The rates of change of angular rates are constant.

All deviations from this assumed behaviour are treated as noise in the system. It is known that these assumptions are not correct, but without sufficient information to model how the real situation differs from this ideal state, any 'non-conformity' can only be treated as noise. For example, there is known to be flexure in the airframe, but it cannot be modelled sufficiently well to be treated as an additional state, therefore it is treated as noise.

Numerous conventions of notation and terminology have been used since the Kalman filter was first derived in Kalman (1960), the following description adopts the conventions used by Cross (1983), chosen for their familiarity to surveyors.

The system behaviour can be described by nine states, forming a state vector, \hat{x} , as follows:

$$\hat{x}^T = [\psi_{afrm} \quad \theta_{afrm} \quad \gamma_{afrm} \quad \dot{\psi}_{afrm} \quad \dot{\theta}_{afrm} \quad \dot{\gamma}_{afrm} \quad \delta\psi \quad \delta\theta \quad \delta\gamma] \quad (7.1)$$

Where,

$\psi_{afrm}, \theta_{afrm}, \gamma_{afrm}$ are the pitch, roll and heading components of the airframe attitude respectively,

$\dot{\psi}_{afrm}, \dot{\theta}_{afrm}, \dot{\gamma}_{afrm}$ are the pitch, roll and heading components of the airframe angular rates respectively, and,

$\delta\psi, \delta\theta, \delta\gamma$ are the offsets in the pitch, roll and heading components between the airframe and the AHRS.

As the offset between the airframe and AHRS axes has been shown to vary, particularly in turbulent flight conditions, three additional states, representing the rate of change of offset, were added to the state vector \hat{x} , to form a twelve state filter. However, tests of this filter, showed that any trends were not well enough defined to be successfully modelled. The nine state filter, which assumes any changes in offset are of a small and random nature, was considered preferable as it prevents any predictable and regular behaviour being assumed for what is in fact a random process.

7.3.2.2 Measurement Model

The Kalman filter combines two models, the measurement and dynamic models, to compute the elements of the state vector, and their precisions, at a particular epoch. In the current application, the measurement model relates the attitude values from the AHRS and GPS solutions to the filter states, and is expressed as:

$$A\hat{x} = b + v \quad (7.2)$$

where,

A is the design matrix relating the measurements to the estimated parameters, described in equation 7.4,

\hat{x} is the state vector described in equation 7.1,

b is the measurement vector described in equation 7.3, and,

v is the vector of measurement noise.

The precisions of the states are computed using the stochastic model of the measurement model, which is simply the covariance matrix of the measurements Cl , described in equation 7.5.

There are nine measurements available from which to estimate the states, forming a measurement vector, b , as follows:

$$b^T = [\psi_{ahrs} \quad \theta_{ahrs} \quad \gamma_{ahrs} \quad \dot{\psi}_{ahrs} \quad \dot{\theta}_{ahrs} \quad \dot{\gamma}_{ahrs} \quad \psi_{afrm} \quad \theta_{afrm} \quad \gamma_{afrm}] \quad (7.3)$$

where,

$\psi_{ahrs}, \theta_{ahrs}, \gamma_{ahrs}$ are the pitch, roll and heading attitude components measured by the AHRS,

$\dot{\psi}_{ahrs}, \dot{\theta}_{ahrs}, \dot{\gamma}_{ahrs}$ are the pitch, roll and heading components of the angular rate, measured by the AHRS, and,

$\psi_{afrm}, \theta_{afrm}, \gamma_{afrm}$ are the pitch, roll and heading attitude components derived from GPS data. As the GPS antennas are rigidly fixed to the airframe, attitude values derived from GPS data are treated as direct measurement of the airframe attitude.

The design matrix, A , is formed based on the measurements available at each epoch. In this application there are four possible scenarios describing the measurements available at an epoch.

- 1) No measurements. This is the case at a scan-line epoch, that does not coincide with an AHRS or GPS measurement epoch. The filtered states will be based on predictions only.
- 2) Six measurements, at an AHRS recording epoch. Three AHRS attitude components, $\psi_{ahrs}, \theta_{ahrs}, \gamma_{ahrs}$, and three components of AHRS angular rate, $\dot{\psi}_{ahrs}, \dot{\theta}_{ahrs}, \dot{\gamma}_{ahrs}$.
- 3) Three measurements, at a GPS recording epoch. Three components of airframe attitude, $\psi_{afm}, \theta_{afm}, \gamma_{afm}$.
- 4) Nine measurements at a coincident GPS and AHRS recording epoch, i.e. all nine elements of the measurement vector, b , are recorded.

If all nine measurements are available at an epoch, then the A matrix is:

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (7.4)$$

The stochastic model of the measurement model, or measurement covariance matrix, Cl , uses the precisions associated with each measurement, to determine the weight that should be given to the measured estimates of the states. In this application, no quality measures are provided with the AHRS or ADU2 values, and those from GRAPE are of limited use, hence, measurement weightings are assigned by the user as an initial step, and held fixed for all epochs. These weightings, assigned by choosing a value for the standard deviation of each measurement, are based on the equipment specifications and the typical accuracies that have been found in previous airborne tests. It is important that these values should properly reflect the relative accuracies of the two measurement devices, and should allow measured states to be properly weighted relative to predicted states.

If it is confirmed that the ATM scanner imagery reflects the AHRS attitude solution more closely than it reflects the airframe motion measured by the GPS antenna array, then the AHRS data can be given greater weighting than the GPS solution. In the case of the CASI scanner, which is rigidly fixed to the airframe, and is detached from the AHRS due to the damping devices on the ATM scanner, it may be appropriate to attach greater weight to the GPS measurements.

A simple covariance matrix, Cl , has been formed, assuming no correlation between individual measurements. This is due to a lack of information from which to properly estimate any correlations. The resulting covariance matrix is therefore:

$$Cl = \begin{bmatrix} \sigma_{\psi_{ahrs}}^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \sigma_{\theta_{ahrs}}^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_{\gamma_{ahrs}}^2 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma_{\dot{\psi}_{ahrs}}^2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \sigma_{\dot{\theta}_{ahrs}}^2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \sigma_{\dot{\gamma}_{ahrs}}^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \sigma_{\psi_{afrm}}^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \sigma_{\theta_{afrm}}^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \sigma_{\gamma_{afrm}}^2 \end{bmatrix} \quad (7.5)$$

where,

$\sigma_{\psi_{ahrs}}^2, \sigma_{\theta_{ahrs}}^2, \sigma_{\gamma_{ahrs}}^2$ are the variances of the pitch, roll and heading attitude components measured by the AHRS,

$\sigma_{\dot{\psi}_{ahrs}}^2, \sigma_{\dot{\theta}_{ahrs}}^2, \sigma_{\dot{\gamma}_{ahrs}}^2$ are the variances of the pitch, roll and heading components of the angular rate, measured by the AHRS, and,

$\sigma_{\psi_{afrm}}^2, \sigma_{\theta_{afrm}}^2, \sigma_{\gamma_{afrm}}^2$ are the variances of the pitch, roll and heading attitude components derived from GPS data.

7.3.2.3 Dynamic Model

The dynamic model describes how the states are expected to change over time, based on knowledge of the system's behaviour, it can be expressed as,

$$\hat{x}_i = M_{i-1,i} \hat{x}_{i-1} + y_{i-1,i} \quad (7.6)$$

where,

\hat{x}_i is the state vector at epoch i , described in equation 7.1,

$M_{i-1,i}$ is the transition matrix describing how the states change between epochs $i-1$ and i , described in equation 7.8.

\hat{x}_{i-1} is the state vector at the previous epoch $i-1$, and,

$y_{i-1,i}$ is a vector representing the unknown errors in the dynamic model.

In this application, and for most other practical problems, it is convenient to consider the $y_{i-1,i}$ vector as being given by (Cross, 1983):

$$y_{i-1,i} = Tg \quad (7.7)$$

where,

g is the vector of quantities which cause the dynamic model to be in error, often referred to as the *driving noise*, or *forcing function*, and,

T is a coefficient matrix relating the error components to the states.

To model the assumptions of the system's dynamic behaviour, the transition matrix $M_{i-1,i}$ is formed as follows:

$$M_{i-1,i} = \begin{bmatrix} 1 & 0 & 0 & \delta t & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & \delta t & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & \delta t & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \quad (7.8)$$

where,

δt is the time interval between epochs, which in this application varies.

In the current filter design, g is assumed to be a vector of random quantities with a zero mean and covariance matrix C_g , described in equation 7.11. In this application the driving noise is considered to be due to the rate of change of the angular rates, $\ddot{\psi}_{afrm}, \ddot{\theta}_{afrm}, \ddot{\gamma}_{afrm}$, and the rate of change of the airframe-to-AHRS offset, $\delta\dot{\psi}, \delta\dot{\theta}, \delta\dot{\gamma}$. Vector g is therefore,

$$g = \begin{bmatrix} \ddot{\psi}_{afrm} \\ \ddot{\theta}_{afrm} \\ \ddot{\gamma}_{afrm} \\ \delta\dot{\psi} \\ \delta\dot{\theta} \\ \delta\dot{\gamma} \end{bmatrix} \quad (7.9)$$

To relate these quantities to the states, the matrix of coefficients, T , is,

$$T = \begin{bmatrix} \delta t^2/2 & 0 & 0 & 0 & 0 & 0 \\ 0 & \delta t^2/2 & 0 & 0 & 0 & 0 \\ 0 & 0 & \delta t^2/2 & 0 & 0 & 0 \\ \delta t & 0 & 0 & 0 & 0 & 0 \\ 0 & \delta t & 0 & 0 & 0 & 0 \\ 0 & 0 & \delta t & 0 & 0 & 0 \\ 0 & 0 & 0 & \delta t & 0 & 0 \\ 0 & 0 & 0 & 0 & \delta t & 0 \\ 0 & 0 & 0 & 0 & 0 & \delta t \end{bmatrix} \quad (7.10)$$

The stochastic model of the dynamic model, or process covariance matrix, C_y , determines the weight that should be given to states estimated using the predicted motion, and is given by:

$$C_y = TC_gT^T$$

where,

T is the coefficient matrix described in equation 7.10, and,

C_g is the covariance matrix containing the variances of the quantities in vector g , and is therefore,

$$C_g = \begin{bmatrix} \sigma_{\ddot{\psi}_{airm}}^2 & 0 & 0 & 0 & 0 & 0 \\ 0 & \sigma_{\ddot{\theta}_{airm}}^2 & 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_{\ddot{\gamma}_{airm}}^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma_{\delta\dot{\psi}}^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & \sigma_{\delta\dot{\theta}}^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & \sigma_{\delta\dot{\gamma}}^2 \end{bmatrix} \quad (7.11)$$

Weights are used to describe the ‘predictability’ of the airframe motion. The weights assigned to the predicted motion relative to the measurements will dictate the nature of the filtered solution. If greater weight is assigned to predicted motion, then a smooth solution that does not vary significantly with each new measurement will be produced. If greater weight is assigned to measurements then the final solution will tend to follow the pattern of the original measurements more closely.

7.3.2.4 Filter Equations

The measurement model and dynamic model, with their associated stochastic models, are used to form a set of equations in which all the available information is used to produce an optimal estimate of the states at an epoch, i . The Kalman filter computes states in three separate operations, namely; ‘prediction’, ‘filtering’ and ‘smoothing’. The equations for these three processes are presented below. The terms (-), (+), and (s) have been introduced in the following equations,

- (-) denotes predicted values, i.e. based on information recorded up to, but not including, the current epoch.
- (+) denotes filtered values, i.e. based on previous and current information.
- (s) denotes smoothed values, i.e. based on previous and current information, as well as values recorded after the current epoch.

Prediction Equations

Assuming $y_{i-1,i}$ to be zero, as any unknown errors are as likely to be positive as negative, the predicted state vector is given by,

$$\hat{x}_i(-) = M_{i-1,i} \hat{x}_{i-1}(+) \quad (7.12)$$

and the predicted covariance of the state is,

$$C\hat{x}_i(-) = M_{i-1,i} C\hat{x}_{i-1}(+) M_{i-1,i}^T + C_y. \quad (7.13)$$

Filtering Equations

Firstly, the Kalman filter gain matrix, G_i , is computed. The gain matrix controls the amount that the filtered state vector, $\hat{x}_i(+)$, and the covariance matrix of the filtered states, $C\hat{x}_i(+)$, are affected by the predicted states, $\hat{x}_i(-)$, and the covariance matrix of the predicted states, $C\hat{x}_i(-)$, relative to the new measurements, b_i , and the covariance matrix of these measurements, Cl_i , at epoch i .

$$G_i = C\hat{x}_i(-) A_i^T (Cl_i + A_i C\hat{x}_i(-) A_i^T)^{-1} \quad (7.14)$$

The filtered state is then computed as,

$$\hat{x}_i(+)=\hat{x}_i(-)+G_i(b_i-A_i\hat{x}_i(-)) \quad (7.15)$$

with the covariance of the filtered states,

$$C\hat{x}_i(+)=(I-G_iA_i)C\hat{x}_i(-) \quad (7.16)$$

In the current filter, a more stable solution for the filtered covariance matrix, which should always maintain symmetry, has been used. This is given by:

$$C\hat{x}_i(+)=(I-G_iA_i)C\hat{x}_i(-)(I-G_iA_i)^T+G_iC_lG_i^T \quad (7.17)$$

Smoothing Equations

If there are a total of n epochs resulting in n state vectors (i.e. $i = n$), then the smoothing solution will produce results at epochs $n-1, n-2, \dots, 1$. The smoothing algorithm will use the filtered solution at the final epoch as a starting point.

Initially letting,

$$\hat{x}_n(s)=\hat{x}_n(+)\quad \&\quad C\hat{x}_n(s)=C\hat{x}_n(+)$$

then for ($i = n, n-1, n-2, \dots$ etc)

firstly, a smoothing matrix, S_i is computed from,

$$S_i=C\hat{x}_{i-1}(+)M_{i-1}^TC\hat{x}_i^{-1}(-) \quad (7.18)$$

the smoothed state vector is then given by,

$$\hat{x}_{i-1}(s)=\hat{x}_{i-1}(+)+S_i(\hat{x}_i(s)-\hat{x}_i(-)) \quad (7.19)$$

with the covariance matrix of the smoothed state,

$$C\hat{x}_{i-1}(s)=C\hat{x}_{i-1}(+)+S_i(C\hat{x}_i(s)-C\hat{x}_i(-))S_i^T \quad (7.20)$$

Smoothing back recursively over the desired number of epochs in this way is known as fixed-interval smoothing, or as the R-T-S method. These algorithms have restrictions in the amount of data that has to be stored for all the epochs that require smoothing. This includes the predicted and filtered state vectors, with the appropriate covariance matrices, for all epochs, and the transition matrices at all update points.

Alternative smoothing algorithms requiring less storage may be used, e.g. smooth back over a short period only, and then move forward. In this application a 5 epoch smoother has been used. This routine smoothes back over the 5 previous epochs only, so that when the filtered solution is computed at epoch n a smoothed solution is produced for epoch $n-4$. Tests on all the data-sets available to date have shown that smoothing has a negligible effect on the final solution. It is therefore not applied as standard as it increases the processing time but does not produce significantly different values to the filtered solution.

In practice it is not uncommon, to omit the smoothing process from a filter. In many applications, smoothing may significantly increase the computational load whilst having a negligible affect on the final results. It is clearly unsuitable for real-time applications, in which filters are commonly used, as it relies on measurements made after the current epoch.

7.3.2.5 Implementation of the Filter

In the course of this research, these algorithms have been implemented in a program called PRAYER, (Pitch, Roll And Yaw³ Estimation Routine). This program has been delivered to the Natural Environment Research Council's Airborne Remote Sensing Facility, and is currently used as one element of a larger software package that integrates navigation and scanner data, and applies geometric and radiometric corrections, to produce georeferenced imagery. This section briefly describes some of the main features of the PRAYER software. Appendix D provides a more detailed description of this software and shows some examples of file formats, etc.

Inputs

In the test flights described in this thesis, there were two sources of attitude data available. Firstly, there are GPS data recorded by the four independent receivers of the aircraft. This data was processed using the direct attitude determination procedure implemented in the GRAPE program, to compute attitude values. The second source of attitude data comes from the gyro based AHRS, which records angular rates as well as attitude components. The AHRS is installed permanently on the survey aircraft.

The four independent GPS receivers were installed on the survey aircraft for the two flights to gather data specifically for this project. In general, the only attitude solution from GPS data for this aircraft comes from the ADU2, a dedicated GPS attitude determination unit. The PRAYER software was designed to process attitude data recorded on these specific trials, but also to be of future use. It therefore accepts a GPS based attitude solution from either GRAPE, or the ADU2. In addition it reads in a scanner file, from the ATM or CASI, which includes the scan numbers and times at which an attitude solution must be computed.

³ The program was originally designed to use yaw values from GRAPE, hence the name. In practice though, yaw values have been reduced to heading components before filtering.

Weighting

Weighting is implemented by assigning standard deviations to each measurement and each motion prediction. Standard deviations are squared to produce variances, which are then used in the covariance matrices C_l and C_g , as described in equations 7.5 and 7.11 respectively. Weights should be selected so that the influence of new measurements and predicted behaviour are appropriately balanced, the measurement weights should also reflect the relative accuracies of the available solutions from AHRS and GPS.

When no reference solution of sufficiently high accuracy is available, it is difficult to judge the appropriate weightings. If scanner imagery is shown to be affected by high frequency motions, then greater weighting should be assigned to the measurements, particularly those from the GPS solution. If, on the other hand, the imagery shows only longer term trends, then greater weight can be assigned to the predicted aircraft behaviour.

Outlier Detection

An outlier detection routine has been implemented in the filter design using the predicted residuals, or innovation sequence. Predicted residuals show the difference between the actual system measurements (AHRS and/or GPS attitude values) and the predicted output, based on the predicted state, $\hat{x}_i(-)$, as computed in equation 7.12. The vector of predicted residuals is computed from,

$$v_i(-) = b_i - A_i \hat{x}_i(-) \quad (7.21)$$

If the predicted residuals exceed a pre-defined threshold level then the measurements from the b_i vector are not used in the computation of the filtered state, i.e. it uses only the values from the predicted state. In the current filter, thresholds can be defined for each of the nine available measurements, and they can be varied depending on the flight conditions. When conditions are generally stable smaller deviations from the general trend can be considered as outliers, in

turbulent conditions thresholds may have to be increased to prevent significant, high frequency attitude changes being considered as erroneous data. If an extended series of measurements is rejected, the filtered solution may fail to detect genuine trends in the data.

7.3.3 Filter Results

This filter has been tested using data collected on the two test flights described previously, and data collected during other missions carried out by the NERC ARSF survey aircraft. The results presented in this section have been computed by running the filter with GPS attitude solutions from the ADU2, and gyro attitude solutions from the AHRS⁴. Although the filter has also been used to integrate attitude solutions from the independent GPS receiver approach and AHRS data collected during the two test flights in October 1998, because the filter performance can be described equally well using either GPS attitude solution, only results from combining an ADU2 solution with an AHRS solution are presented here.

Figures 7.1 to 7.3 show the result of applying this filter to data recorded over a three minute flight line in which the ADU2 and AHRS measured attitude information and the ATM scanner recorded imagery. In each plot the original ADU2 solution is shown as a magenta line, the original AHRS solution as a blue line, and the filtered solution as a green line. Figures 7.1, 7.2 and 7.3 show these three solutions for the roll, pitch and heading components respectively.

⁴ These data were received from the NERC ARSF to test the filter during its development at UCL.

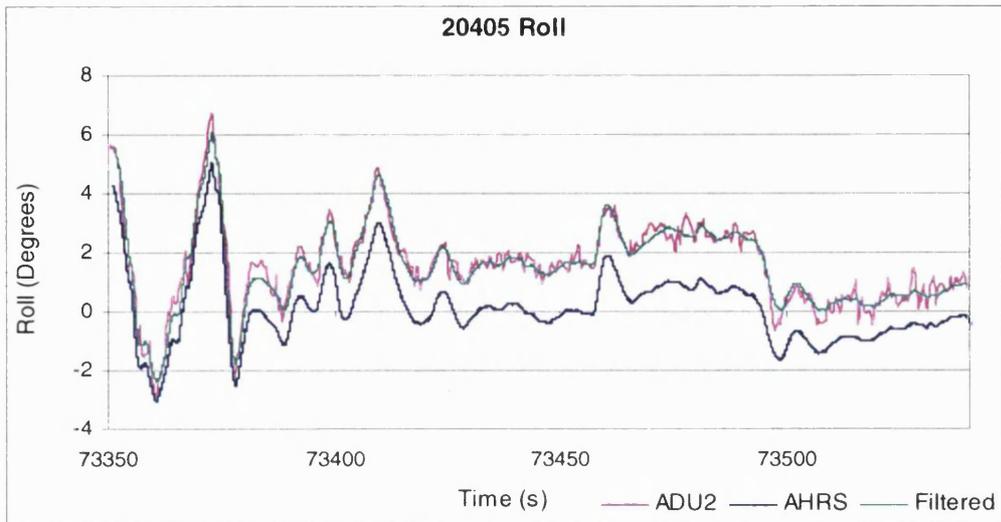


Figure 7.1: Filtered solution from ADU2 and AHRS roll data. 20405

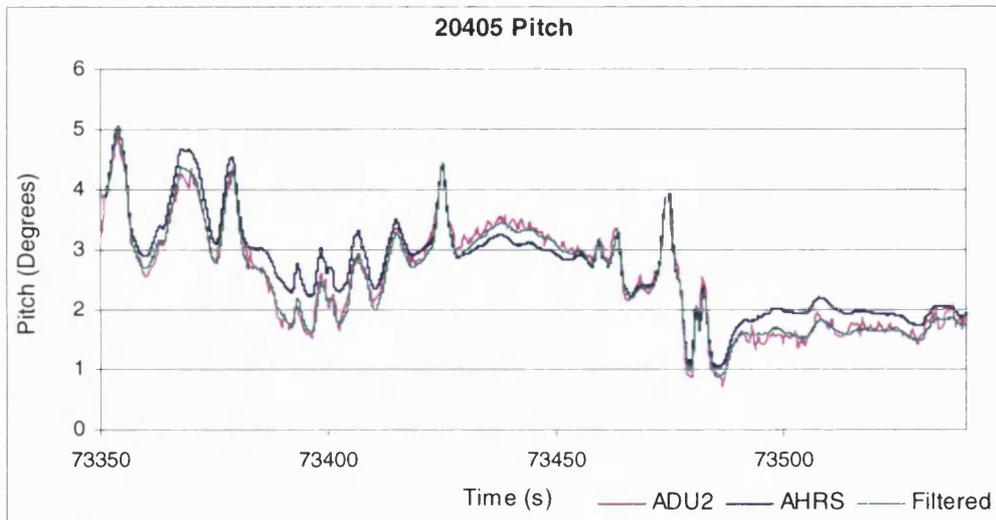


Figure 7.2: Filtered solution from ADU2 and AHRS pitch data. 20405

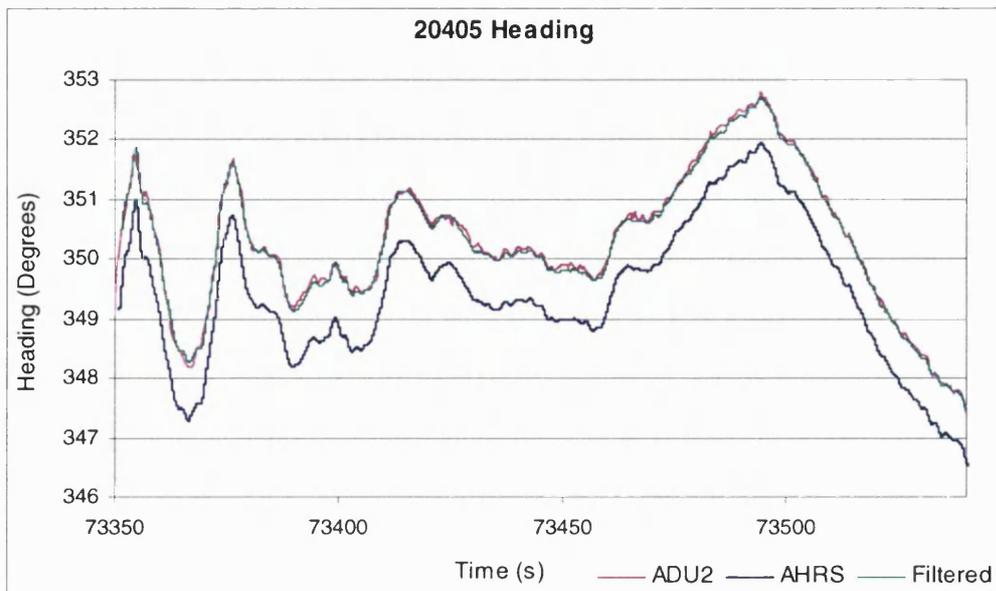


Figure 7.3: Filtered solution from ADU2 and AHRS heading data. 20405

The following features can be seen in these figures:

- In each case the short-term shape of the AHRS solution is preserved, but any drifts have been removed by constraining it to the ADU2 solution, which is not subject to the same time dependent biases. When the same filter is used to integrate GRAPE and AHRS data in a similar way it again accounts for any offsets in the solutions.
- High frequency variations in the ADU2 attitude solution are not reflected in the filtered solution as the weights have been assigned so that the attitude values for each ATM scan line follow the trends measured by the AHRS.
- For this flight line the AHRS compass control unit and magnetic sensor unit, described in section 4.2.4, have been operated, so in fact the original AHRS solution has a reasonably good orientation. In all the other flight data considered, including that recorded by the AHRS during the October 1998 flights, the AHRS heading has been assigned a purely arbitrary orientation. By estimating the offsets between the AHRS axes and the airframe, a precise heading solution can be achieved, even if the compass control unit and magnetic sensor unit are not in operation.

To date, all the flight data processed at the NERC ARSF and UCL with this filter have produced successful solutions for each attitude component over the entire flight line, using a set of default weightings (listed in appendix D), with one exception. For the pitch component of flight line 20402 (Figure 7.4) there is a period from 44225 to 44250 during which the standard filtered solution (shown in green) does not reflect the trends in either of the original solutions. The reason for this is the significant divergence in the two original data sources. However, by adjusting the model weights and rejection criteria, a better solution (shown in yellow) can be produced. In this case, the weight assigned to the prediction that the rate of change of pitch remains constant has been reduced, and a rejection threshold has been chosen so that some extremes of the ADU2 solution are ignored but not to the extent that any trends are missed.

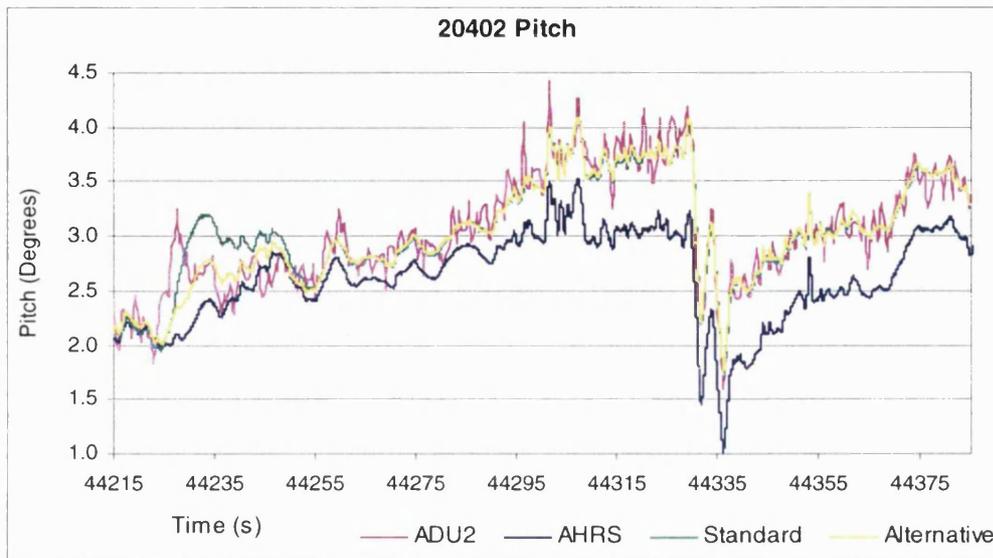


Figure 7.4: Filtered solutions from ADU2 and AHRS pitch data. 20402.

If the ADU2 and AHRS measurements for an attitude component over a single flight line show variations of this nature there is no fully automated way to produce a good filtered solution for the entire flight line. User intervention is required to make decisions as to which of the instrument outputs should be trusted, and which motion predictions are likely to be the most valid. Without knowing the true attitude values, or having quality measures from the instruments for these periods, the only option is for a user to make this judgement.

This filter successfully integrates GPS and AHRS data to compute attitude values for each scan line. This can usually be achieved using a set of standard processing parameters, but there are situations in which a degree of user intervention may be required due to inconsistencies in the original data.

7.4 Concluding Remarks

This chapter has described a method for improving the efficiency, and potentially the robustness, of the GPS direct attitude determination method, by using AHRS values to reduce or even eliminate the ambiguity search space. This integrated ambiguity resolution technique makes the GPS attitude determination process a genuinely single epoch solution.

A Kalman filter that integrates attitude solutions from AHRS and GPS data was described in terms of the input data, the measurement and dynamic models, and the filter equations. This filter provides a convenient means of combining these two data sources with assumptions regarding the aircraft's dynamics, in order to produce optimal estimates of the scanner attitude for each scan line. The filter has been implemented in a program called PRAYER, which was developed during the course of this research.

PRAYER is now used routinely in the NERC ARSF integrated data system. Results from processing a variety of data sets, and comparing the georeferenced imagery against ground control, have shown that PRAYER successfully computes appropriate attitude components for the scanner imagery.

Chapter 8 Conclusions and Suggestions for Further Work

8.1 Conclusions

The research described in this thesis has focused on the design and testing of a GPS-based direct georeferencing system to provide position and attitude information for airborne scanning. A trial system has been described in terms of the equipment and the processing methods used, and an analysis of the results have demonstrated the suitability of this approach for a range of remote sensing applications. A number of conclusions can be drawn from this research, and are summarised below. The first four describe the important results to come directly from investigations undertaken during this research, the final two place this work in a wider context.

1) Overall Performance

The GPS-based direct georeferencing method developed during this research meets the typical accuracy requirements for remote sensing applications using airborne line scanners. It computes position and attitude values at a high enough rate and with sufficient accuracy to provide all six exterior orientation parameters for the type of scanners commonly used for environmental monitoring and resource mapping.

The system consists of four GPS antennas mounted on a survey aircraft, each connected to an independent receiver. A local reference station also needs to be established in the survey area to allow very high accuracy positioning. Two trial flights were carried out over a surveyed test site to gather data using this configuration.

Positions were determined using two methods, both of which produce a fixed ambiguity carrier phase solution. The first of these methods resolves ambiguities using only a single epoch of GPS data, the second employs a multiple epoch method for ambiguity resolution. The accuracy of the positions determined in this application, from either solution, is generally between 5 and 10mm in plan, and 10 to 30mm in height.

The accuracy of the attitude solution produced when GPS data from the four independent aircraft receivers are processed using the direct attitude method is approximately 7 arc minutes in the pitch component, 20 arc minutes in roll, and 3 arc minutes in heading.

At a flying height of 600m this direct georeferencing system has the potential to position imaged features in a mapping co-ordinate system with an accuracy of around 4m, using little or no ground control. This is well within the accuracy requirements of 10 to 20m that the earth science community generally require. It therefore meets the objectives of this project, and would be a suitable means of providing georeferencing information for the majority of tasks undertaken by survey aircraft in the environmental and earth sciences.

2) **Robustness**

A fixed ambiguity, carrier phase solution can be used to position each aircraft antenna, relative to the local reference station, at every epoch during the selected flight lines using a multiple epoch ambiguity resolution method. Using a single epoch ambiguity resolution technique, ambiguities can be resolved between the local reference station and the aircraft antennas for the majority of the time. For both these methods to be successful, noisy data from specific satellites recorded at specific receivers must be removed from the solution. The high levels of noise found are almost certainly due to the affects of multipath at the aircraft antennas.

For attitude determination, ambiguities can be successfully resolved at every epoch during the selected flight lines using the direct attitude algorithm. In order for a single baseline positioning method to replicate this success, specific noisy measurements, almost certainly caused by multipath, must be removed from the solution. The direct attitude method is therefore more robust than a single baseline method. This improved robustness is due to the increased redundancy of the direct attitude method in which attitude parameters are computed directly from GPS phase measurements without explicitly determining the relative antenna positions.

3) **Recording Rates**

A GPS recording rate of 1Hz or lower, with a relatively simple interpolation method, should be adequate to provide positions for the scanners commonly used on small survey aircraft.

A higher recording rate appears to be necessary for the attitude components as the airframe may be subject to some significant high frequency vibrations. The exact rate required cannot be confirmed for the current application without scanner imagery, but it may be as high as 10Hz.

4) **Data Integration**

The efficiency, and potentially the robustness, of the GPS direct attitude determination method, has been improved by using attitude values from a low-cost gyro to reduce or even eliminate the ambiguity search space. This integrated ambiguity resolution technique makes the GPS attitude determination process a genuine single epoch solution.

A Kalman filter has been developed to integrate attitude solutions from a gyro and the GPS processing. The filter combines these two data streams with assumptions regarding the aircraft's dynamics, in order to produce optimal estimates of the scanner attitude for each scan line. This filter has been implemented in a piece of software, and is now used routinely in an operational environment. Results from processing a variety of data-sets have shown that it successfully computes appropriate attitude components for scanner imagery.

5) Context of this Research within the Field of Direct Georeferencing

The use of on-board navigation instruments to determine the exterior orientation parameters of airborne remote sensing devices is becoming increasingly common, and in numerous applications is now a standard procedure. Direct georeferencing of this kind can reduce or even eliminate the need for expensive and time-consuming ground control surveys, and increases operational flexibility. Systems have evolved to meet the full spectrum of accuracy requirements for airborne photogrammetry and remote sensing. These systems generally incorporate inertial systems and GNSS technology, and integrate the measurement data using a Kalman filter. Although many of the systems have much in common, they are increasingly being adapted to meet the needs of particular applications, providing ad-hoc solutions.

One of the major trends in direct georeferencing over the past few years has been the move towards low-cost systems, suited to lower accuracy remote sensing applications. These systems no longer use a medium or high quality INS as the primary navigation instrument, instead they incorporate GPS with a low-cost INS, or rely on GPS data alone to determine the scanner's position and attitude. The system developed in the course of this research project fits well with this trend as it uses GPS for determining both position and attitude. Although it can operate with only GPS data, it also has the flexibility to incorporate additional high rate attitude data from a low cost gyro.

6) General Applicability of GPS Attitude Determination from Independent Receivers

The direct attitude algorithm using GPS data from independent receivers also has the potential to be a useful tool for applications outside the field of airborne remote sensing. Attitude determination is important in many applications on land, sea, air and in space, from hydrographic surveys to spacecraft rendezvous and docking operations. As GPS receivers and antennas have become smaller, more rugged, and more affordable, they have become a practical alternative to traditional gyroscopic instruments.

For example, the orientation of an acoustic sounder needs to be known to make accurate depth measurements from a survey vessel. As the sounder is rigidly fixed to the vessel, determining the vessel's attitude will also determine the sounder's attitude. GPS antennas connected to independent receivers can be spaced widely apart on the vessel to improve the potential accuracy, and the robustness of the solution can be maximised by using all available GPS measurements to directly determine the attitude parameters.

8.2 Suggestions for Further Work

The following suggestions for future work have been divided into three sections. In section 8.2.1, additional methods that could be used to more rigorously calibrate the direct georeferencing system are discussed. Each of these relies on additional data that has not been available in the course of the current research. Section 8.2.2 describes how some of the processing methods used in this research, primarily those used in GPS attitude determination and filtering, could potentially be modified to improve the system for the current application and for possible future work. In section 8.2.3, some more general suggestions for similar work in the general field of direct georeferencing, are made.

8.2.1 Calibration

To provide a more comprehensive assessment of the direct georeferencing system developed in this project, additional validation procedures should be employed.

- (a) To adequately assess the on-the-ground accuracy with which the scanner imagery is finally georeferenced, clearly requires scanner data. Using this imagery, the combined effect of errors in position and attitude determination can be assessed, together with errors due to image measurement and the scanner properties.
- (b) An attitude solution from the ADU2, dedicated GPS attitude determination unit, would provide a calibration to at least the accuracy of that instrument. A better attitude reference solution could be obtained by operating a high accuracy INS during flights.
- (c) Using an aerial survey camera with digital input and output, the instant of exposure could be related to the position and attitude solution using GPS time. Determining the exterior orientation parameters of each photograph, either by establishing sufficient ground control in the coverage area of a single photo, or by performing a block bundle adjustment, would provide a sample of precisely known position and attitude values along the trajectory with which to calibrate the system.

- (d) If subsequent flights were undertaken with equipment that allowed a 10Hz GPS solution for a greater period of the flight, the possible high frequency vibration described in section 6.5.2 could be further investigated. Scanner imagery could also be used to determine whether or not this airframe motion is experienced by either of the scanners.

8.2.2 Algorithms

For the current application the GPS attitude determination was found to be adequate but a number of modifications could be made to improve the efficiency and applicability of the method. Similarly, the Kalman filter described in chapter 7 is relatively simple and appears to work well, but it too could potentially be improved.

- (a) Investigations have shown that the additional redundancy of the direct attitude method has allowed it to successfully resolve ambiguities using noisy data that causes a single baseline method to fail. If the noise levels increased further, and more phase measurements were affected, it may become necessary to identify and remove outliers before the direct attitude method can successfully resolve ambiguities. An outlier detection routine should therefore be introduced at the ambiguity resolution stage. This routine could be based on the cosine of the phase residuals, determined during the ambiguity function value computation, as described in section 5.5.7.1. Alternatively, if the least squares ambiguity searching technique (LSAST) was introduced, a similar routine to that which is used in the original GRAPE software, based on the Tau statistic, could be used.

- (b) The efficiency of the ambiguity resolution procedure could be improved in a number of ways. Reductions in the processing time may make the technique suitable for real-time applications.
- (i) If an additional source of attitude data is available, this can be used to improve the starting approximation on which the search space is centred, allowing the search space to be reduced or even eliminated, as described in section 7.2. Any reliance on an external source of data clearly limits the applicability of the method for general use.
 - (ii) The starting approximation for the search space could also be improved by implementing a simple prediction routine, or perhaps a Kalman filter. This too might allow a reduction in the search space, and hence the processing time.
 - (iii) Replacing the ambiguity function method (AFM) with the least squares ambiguity searching technique (LSAST) should improve efficiency because the processing time reduces as the number of observations increases with this method. This approach however may not be as successful using a single epoch solution when observation conditions are poor, or when the baseline length exceeds 5km. The benefits and limitations of LSAST using a single epoch of GPS data require further investigation.
 - (iv) Combining the ambiguity searching techniques of the single epoch positioning software, GASP, with the GPS attitude determination program, GRAPE, could improve the efficiency of the solution and allow the chosen ambiguity set to be validated. The number of trial attitudes at which computations are carried out can be reduced by only working with candidates that lead to different ambiguity combinations, and which produce a baseline estimate which satisfies an *a priori* constraint (see section 5.5.7.2).

- (c) Multipath modelling as described by Barnes *et al.* (1998) could be incorporated into GASP and GRAPE in the future. At present however, the models developed (based on the phase signal to noise ratio) are extremely complicated to apply and need significant human intervention. To date this method has only proved successful when static tests were conducted using known baselines, allowing reverse engineered stochastic models to be developed. The implementation of some form of multipath model has the potential to improve the accuracy of both the position and attitude solutions.

- (d) The Kalman filter described in section 7.3 makes some simple assumptions regarding the relative weights of measurements and predictions. It may be possible to ‘tune’ the filter by analysing the predicted residuals produced when the filter is run with different data. Improvements may also be possible if an adaptive filtering mechanism is introduced, as described in section 2.3.3.3. To judge whether changes in the filtered solution actually represent an improvement requires a reference solution.

8.2.3 General Suggestions

- (a) The robustness of the direct GPS attitude determination method has been compared to that of a single baseline system, but to date no assessment has been made of the relative accuracy of the two approaches. To compare the attitude solution from each system requires the single baseline method to be developed further. The single baseline method should use antenna positions in the body frame system to determine baseline constraints for ambiguity resolution. Once the relative antenna positions are determined a constrained least squares adjustment should be used to derive attitude.

- (b) Throughout this work the airframe of the aircraft has been treated as a rigid body. In reality of course, there will be some degree of flexure or deformation. Further studies are needed to establish whether this flexure exhibits any systematic behaviour that could be modelled, perhaps by adding one or more flexure terms to the state vector of the Kalman filter.
- (c) The maximum distance over which a fixed ambiguity carrier phase solution can be used to position the aircraft has not been investigated, nor have any alternative GPS positioning methods. Further investigations are required to establish the potential positioning accuracy when the aircraft operates at greater distances from the local reference station.
- (d) Provided that sufficient satellite observations are available, a single frequency solution could be used to derive attitude with the direct attitude method, and possibly with a single baseline system. The use of L1 only receivers would considerably reduce the cost of the system, although one dual frequency receiver may be required to provide sufficiently accurate position information.

References and Bibliography

- Abidin, H. Z. (1993). On-the-fly Ambiguity Resolution: Formulation and Results. *Manuscripta Geodetica* (1993) 18: 380-405
- Ackermann, F. (1999). Airborne Laser Scanning – Present Status and Future Expectations. *ISPRS Journal of Photogrammetry and Remote Sensing*, 54 (1999) 64-67.
- Ackermann, F. and H. Schade (1993). Application of GPS for Aerial Triangulation. *Photogrammetric Engineering and Remote Sensing*, Vol. 59, No. 11, November 1993, 1625-1632.
- Al-Haifi, Y.M.H. (1996). Short Range GPS Single Epoch Ambiguity Resolution. Ph.D. thesis, Department of Geomatics, University of Newcastle, Newcastle upon Tyne, 292 pp.
- Al-Haifi, Y. M., S. J. Corbett, and P. A. Cross (1998). Performance Evaluation of GPS Single Epoch on-the-fly Ambiguity Resolution. *Navigation: Journal of the Institute of Navigation*, Vol. 44, No. 4: 479-487.
- Applanix (2000). Applanix Corporation Website: <http://www.applanix.com>. Contains all product information and technical specification referred to in text.
- Ashtech (1996). Attitude Determination Unit 2TM: User's Guide.
- Barnes, J. B., N. Ackroyd and P. A. Cross (1998). Stochastic Modelling for Very High Precision Real-time Kinematic GPS in an Engineering Environment. *Proceedings of FIG XXI, International Congress*, Brighton, July 19-25, 1988, Commission 6, Engineering Surveys, 61-76.
- Barnes, G. P., and M. J. Little (1997). Airborne GPS/INS Capability for the Australian Defence Force: The FMS-800 Integrated with LTN-92. *Proceedings of the Satellite Division of the Institute of Navigation 10th International Technical Meeting: ION GPS-97*, September 16-19, 1997, 1019-1027.
- Bastos, L., S. Cuhna, T. Cuhna and P. Tome (1999). Real Time Aircraft Position + Attitude of Airborne Sensors. *Proceedings of the Satellite Division of the Institute of Navigation 12th International Technical Meeting: ION GPS-99*, September 14-17, 1999, 1999-2007.
- Baustert, G., E. Cannon, E. Dorrer, G. Hein, H. Krauss, H. Landau, K. P. Schwarz, Ch. Schwiertz (1989). German-Canadian Experiment in Airborne INS-GPS Integration for Photogrammetric Applications. In Y. Bock and N. Leppard (Eds) *Proceedings of the IAG Symposia. Vol. 102. Global Positioning System: Overview*. Edinburgh, UK, August, 1989, Springer Verlag, 328-333.

Becker, R. D. and J. P. Barriere (1993). Airborne GPS for Photo Navigation and Photogrammetry: An Integrated Approach. *Photogrammetric Engineering and Remote Sensing*, Vol. 59, No. 11, November 1993, 1659-1665.

Brown, R.G (1983). *Introduction to Random Signals and Applied Kalman Filtering*. Second Edition. John Wiley & Sons, Inc.

Brown, R.G, and P.Y.C Hwang (1992). *Introduction to Random Signals and Applied Kalman Filtering*. Second Edition. John Wiley & Sons, Inc.

Bruton, A.M., K-P. Schwarz and J. Skaloud (1999). The use of Wavelets for the Analysis and De-noising of Kinematic Geodetic Measurements. In K.-P. Schwarz (Ed) *Proceedings of the IAG Symposia. Vol. 121 Geodesy Beyond 2000 - The Challenges of the First Decade*. Birmingham, UK, July 19-30, 1999, Springer Verlag, 227-232.

Burman, H. (2000). Calibration and Orientation of Airborne Image and Laser Scanning Data using GPS and INS. PhD Thesis, Royal Institute of Technology, Department of Geodesy and Photogrammetry, Stockholm, Sweden.

Campana, R., L. Marradi and S. Bonfanti (1999). GPS Attitude Determination by Adaptive Kalman Filtering. *Proceedings of the Satellite Division of the Institute of Navigation 12th International Technical Meeting: ION GPS-99*, September 14-17, 1999, 1979-1988.

Cohen, C. E., B. W. Parkinson and B. D. McNally (1994). Flight Tests of Attitude Determination Using GPS Compared Against an Inertial Navigation Unit. *Navigation: Journal of the Institute of Navigation*, Vol. 41, No.1, Spring 1994, 83-97.

Colombo, O. L., C. Rizos and B. Hirsch (1995). Long-Range Carrier Phase DGPS: The Sydney Harbour Experiment. *Proceedings of the 4th International Conference of Differential Satellite Navigation Systems DSNS 95*, 24-28 April, 1995.

Colomina, I. (1993). A Note on the Analytics of Aerial Triangulation with GPS Aerial Control. *Photogrammetric Engineering and Remote Sensing*, Vol. 59, No. 11, November 1993, 1619-1624.

Corbett, S. J. (1993). GPS for Attitude Determination and Positioning in Airborne Remote Sensing. *Proceedings of the Satellite Division of the Institute of Navigation 6th International Technical Meeting: ION GPS-93*, September, 1993, 789-796.

Corbett, S. J. (1994). GPS Single Epoch Ambiguity Resolution for Airborne Positioning and Orientation. PhD Thesis, University of Newcastle upon Tyne, UK.

- Corbett, S.J. and P.A. Cross (1995) GPS Single Epoch Ambiguity Resolution. *Survey Review*, Vol. 33, No. 257: 149-160.
- Corbett, S.J. and T.M. Short (1994). Positions and Orientations of Airborne Sensors from Kinematic GPS Data. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS94)*, Banff, Canada, June, 349-357.
- Corbett, S.J. and T.M. Short (1995). Development of an Airborne Positioning System. *Photogrammetric Record*, 15(85):3-15 (April 1995).
- Counselman, C. C., and S. A. Gourevitch (1981). Miniature Interferometer Terminals for Earth Surveying: Ambiguity and Multipath in Global Positioning System. *IEEE Transactions on Geoscience and Remote Sensing*.GE-19 (No.4): 244-252.
- Cramer, M., D. Stallmann, and N. Halla (1997). High Precision Georeferencing Using GPS/INS and Image Matching. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 453-462.
- Cramer, M., and N. Haala (1999). Direct Exterior Orientation of Airborne Sensors. *GIM International*, Vol. 13, No. 9, September 1999: 46-49.
- Creamer, N. G., G. C. Kirby, R. E. Weber, A. B. Bosse, S. Fisher, and F. A. Tasker (1998). An Integrated GPS/Gyro/Smart Structures Architecture for Attitude Determination and Baseline Metrology. *Navigation: Journal of the Institute of Navigation*, Vol. 45, No.4, Winter 1998-1999.
- Cross, P.A. (1983). Advanced Least Squares Applied to Position-Fixing. Working Paper No. 6, Department of Land Surveying, University of East London, London, 205 pp.
- Cross, P. A., S. J. Corbett, and M.R. Mahmud (1997). Benchmarking of Commercial Offshore GPS-Based Attitude Determination Systems. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 379-388.
- Crowl, J. and D. Merchant (1995). Airborne GPS Photogrammetric Calibration and Test Ranges at Madison County, Ohio. *In Proceeding of ASPRS/ACSM*, Vol. 3, 603-611.
- Curry, S. and K. Schuckman (1993). Practical Considerations for the Use of Airborne GPS for Photogrammetry. *Photogrammetric Engineering and Remote Sensing*, Vol. 59, No. 11, November 1993, 1611-1617.
- Da, R. (1997). Analysis and Test Results of AIMS GPS/INS System. *Proceedings of the Satellite Division of the Institute of Navigation 10th International Technical Meeting: ION GPS-97*, September 16-19, 1997, 771-779.

Da, R., G. Dedes and K. Shubert (1996). Design and Analysis of a High Accuracy Airborne GPS/INS System. *Proceedings of the Satellite Division of the Institute of Navigation 9th International Technical Meeting: ION GPS-96*, September 17-20, 1996, 955-963.

van Diggelen (1998). GPS Accuracy: Lies, Damn Lies, and Statistics. *GPS World*, January 1998, 41-45.

Edwards, S.J., P.A. Cross, J.B. Barnes, D. Betaille (1999). A Methodology for Benchmarking Real Time Kinematic GPS. *The Hydrographic Journal* No. 94, October 1999, 11-15.

El-Mowafy, A. and K.P. Schwarz (1994). Epoch by Epoch Ambiguity Resolution for Real-Time Attitude Determination Using a GPS Multi-Antenna System. *Navigation: Journal of the Institute of Navigation*, Vol. 42, No.2 Summer 1995, 391-408.

El-Sheimy, A. and K.P. Schwarz (1999). Navigating Urban Areas by VISAT – A Mobile Mapping System Integrating GPS/INS/Digital Cameras for GIS Applications. *Navigation: Journal of the Institute of Navigation*, Vol. 45, No. 4 Winter 1998-1999, 275-285.

Erickson, C., (1992) An Analysis of Ambiguity Resolution Techniques for Rapid Static GPS Surveys using Single Frequency Data. *Proceedings of the Satellite Division of the Institute of Navigation 5th International Technical Meeting: ION GPS-92*, Albuquerque, NM, 1992, 453-461.

Euler, H, and C. Hill (1995). Attitude Determination: Exploiting all Information for Optimal Ambiguity Resolution. *Proceedings of the Satellite Division of the Institute of Navigation 8th International Technical Meeting: ION GPS-95*, September 12-15, 1751 - 1757.

Feng, Y. (1997). A Multiple Outlier Diagnostics Approach Towards Sub-Centimetre GNSS Kinematic Positioning. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 113-124.

Ford, T., W. Kunysz, R. Morris, J. Neumann, J. Rooney, T. Smit (1996). Beeline RT20 a Compact, Medium Precision Positioning System with an Attitude. *Proceedings of the Satellite Division of the Institute of Navigation 9th International Technical Meeting: ION GPS-96*, September 17-20, 1996, 687-695.

Fraser, C.S. (1995). Observational weighting Considerations in GPS Aerial Triangulation. *Photogrammetric Record* 15(86) 263-275, October 1995.

Frei, E. and Beutler, G. (1990). Rapid Static Positioning Based on the Fast Ambiguity Resolution Approach "FARA": Theory and First Results. *Manuscripta Geodetica*, Vol. 15, 1990, 325-56.

- Fricker, P., R. Sandau and A. S. Walker (1999). Airborne Digital Three-lines Sensors. *GIM International*, June 1999, 49-51.
- Fuller, R., S. Gomez, L. Marradi and J. Rodden (1996). GPS Attitude Determination from Double Difference Differential Phase Measurements. *Proceedings of the Satellite Division of the Institute of Navigation 9th International Technical Meeting: ION GPS-96*, September 17-20, 1996, 1073-1080.
- Gelb, A. (1974). *Applied Optimal Estimation*. M.I.T Press.
- Gleason, D. M. (1992). Extracting Gravity Vectors from the Integration of Global Positioning System and Inertial Navigation System Data. *Journal of Geophysical Research*, Vol. 97, No. B6, June 10, 1992, 8853-8864
- Goldstein, H. (1980). *Classical Mechanics*. 2nd edition. Addison-Wesley Publishing Company, Reading, Massachusetts, USA, 672 pp.
- Gourevitch, S., F. van Diggelen, X. Qin and M. Kuhl (1995). Ashtech RTZ, Real-Time Kinematic GPS with OTF Initialization. *Proceedings of the Satellite Division of the Institute of Navigation 8th International Technical Meeting: ION GPS-95*, September 12-15, 1995, 167-170.
- van Graas, F. and M.S. Braasch (1992). Real-Time Attitude and Heading Determination Using GPS. *GPS World*, Vol. 3, No. 3, pp. 32-39.
- Graydon, O. (1997). Integrated Gyro is set to Reduce Cost of Navigation. *Opto and Laser Europe*, Issue 46, 23-26, December 1997.
- Grejner-Brzezinska, D.A., R. Da and C. Toth (1998). GPS Error Modelling and OTF Ambiguity Resolution for High-Accuracy GPS/INS Integrated System. *Journal of Geodesy* (1998) 72: 626-638.
- Gruen, A., M. Cocard and H.-G. Kahle (1993). Photogrammetry and Kinematic GPS: Results of a High Accuracy Test. *Photogrammetric Engineering and Remote Sensing*, Vol. 59, No. 11, November 1993, 1633-1636.
- Gurtner, W. (1994). RINEX: Receiver-Independent Exchange Format. *GPS World*, Vol. 5, No. 7, July 1994, 48-52.
- Han, S., K. Wong and C Rizos (1997). Instantaneous Ambiguity Resolution for Real-Time GPS Attitude Determination. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 409-416.
- Harvey, R. (1998). Development of an Integrated GPS Dual-Antenna Fixed Baseline Precision Pointing System. *Proceedings of the Satellite Division of the Institute of Navigation 11th International Technical Meeting: ION GPS-98*, September, 1998, 1801-1810.

Hatch, R. (1989). Ambiguity Resolution in the Fast Lane. *Proceedings of the Satellite Division of the Institute of Navigation 2nd International Technical Meeting: ION GPS-89*, 45-50.

Hatch, R. (1990). Instantaneous Ambiguity Resolution. *Proceeding of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS90)*, Banff, Canada, 12-19.

Hayward, R. C. and J. D. Powell (1998). Real Time Calibration of Antenna Phase Errors for Ultra Short Baseline Attitude Systems. *Proceedings of the Satellite Division of the Institute of Navigation 11th International Technical Meeting: ION GPS-98*, September, 1998, 1753 - 1762.

He, X.F., Y.Q. Chen and H.B. Iz (1997). An Improved Filter for High Order Integrated GPS/INS Based on U-D Factorization. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 257-264.

He, X.F., Y.Q. Chen and B. Vik (1999). Design of Minimax Robust Filtering for an Integrated GPS/INS System. *Journal of Geodesy* (1999) 73: 407-411.

Heipke, C. (1999). Automatic Aerial Triangulation: Results of the OEEPE-ISPRS Test and Current Developments. *Photogrammetric Week in Stuttgart 1999*.

Hill, C. and T. Moore (2000). A Free Receiver Upgrade ! GPS Surveying without Selective Availability. *Surveying World*, Vol. 8, July/August 2000, 21-22.

Hofmann, O. (1988). A Digital Three Line Stereo Scanner System. *Archives of the International Society of Photogrammetry and Remote Sensing, Commission II, WGII/6*, 206-213.

Hofmann-Wellenhof, B., H. Lichtenegger, and J. Collins (1994). *Global Positioning System: Theory and Practice*. 3rd edition. Springer-Verlag, Wien and New York, 355 pp.

Howell, G. and W. Tang (1994). A Universal GPS/INS Kalman Filter Design. *Proceedings of the Satellite Division of the Institute of Navigation 7th International Technical Meeting: ION GPS-94*, September, 1994, 443-451.

Husti, G. J. and P. G. Sluiter (1999). Airborne Kinematic GPS with OTF Ambiguity Resolution. *The Hydrographic Journal*, No. 94, October 1999, 17-23.

Jacobsen, K. (1993). Experiences in GPS Photogrammetry. *Photogrammetric Engineering and Remote Sensing*, Vol. 59, No. 11, November 1993, 1651-1658.

Jia, M, X. Ding and M. Tsakiri (1997). A New Robust Algorithm for Kinematic Surveying Systems. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 313-320.

- de Jong, K. (2000). Airborne Kinematic GPS. *Geo Informatics/ Volume 3 June 2000*. 32-35.
- Kalman, R. E., (1960). A New Approach to Linear Filtering and Prediction Problems. ASME. *Journal of Basic Engineering*, Volume 82, March 1960, 35-45.
- Keong, J.H. and G. Lachapelle (1998). Heading and Pitch Determination using GPS/GLONASS with a Common Clock. *Proceedings of the Satellite Division of the Institute of Navigation 11th International Technical Meeting: ION GPS-98*, September, 1998, 2053-2061.
- Kleusberg, A.(1995). Mathematics of Attitude Determination with GPS. *GPS World*, Vol. 6, No. 9: 72-78.
- Kleusberg, A. and P.J.G. Teunissen, eds. (1996). *GPS for Geodesy*. Springer-Verlag, Berlin and Heidelberg, 407 pp.
- Koch, K. R. and Y. Yang (1998). Robust Kalman Filter for Rank Deficient Observation Models. *Journal of Geodesy* (1998) 72: 436-441.
- Kornfeld, R. P., J. Hansman and J. J. Deyst (1998). Single Antenna GPS-Based Aircraft Attitude Determination. *Navigation: Journal of the Institute of Navigation*, Vol. 45, No. 1 Spring 1998.
- Krabill, W. B., C. W. Wright, R. N. Swift, E. B. Frederick, S. S. Manizade, J. K. Yungel, C.F. Martin, J. G. Sonntag, M. Duffy, W. Hulslander and J. C. Brock (2000). Airborne Laser Mapping of Assateague National Seashore Beach. *Photogrammetric Engineering and Remote Sensing*. Vol. 66, No. 1, January 2000, 65-71.
- Krabill, W. B., R. H. Thomas, C. F. Martin, R. N. Swift and E. B. Frederick (1995). Accuracy of Airborne Laser Altimetry over the Greenland Ice Sheet. *International Journal of Remote Sensing*, Vol. 16, No.7, 1211-1222.
- Kusevic, K. and P. Mrstik (1997). The Effects of Varying Epoch Intervals on Non Linear Interpolation Solutions for Airborne GPS Trajectories. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 329-336.
- Lachapelle, G. and K. P. Schwarz (1989). Kinematic Applications of GPS and GPS/INS Algorithms, Procedures, and Equipment Trends. In Y. Bock and N. Leppard (Eds) *Proceedings of the IAG Symposium Vol. 102. Global Positioning System: Overview*. Edinburgh, UK, August, 1989, Springer Verlag, 298-315.
- Lapucha, D., R. Barker and Z. Liu (1996). High-Rate Precise Real-Time Positioning using Differential Carrier Phase. *Navigation: Journal of the Institute of Navigation*, Vol. 43, No.3, Fall 1996, 295-305

Leica, (1994). SKI – Static Kinematic Software: User Manual. Revised June 1994. Leica AG, Heerbrugg (Switzerland).

Leick, A. (1995). *GPS Satellite Surveying*. John Wiley & Sons, New York, USA.

Levy, L.J. (1997). The Kalman Filter: Navigation's Integration Workhorse. *GPS World*, Vol. 8, No. 9, September 1997, pp 65-71.

LH Systems, (2000a). ASCOT Aerial Survey Control Tool. Brochure available to download from www.lh-systems.com

LH Systems, (2000b). Flykin Suite++ for Post-Processing of Airborne GPS Data. Brochure available to download from www.lh-systems.com

LH Systems, (2000c). ORIMA Orientation Management. Brochure available to download from www.lh-systems.com

Litef (1997). LCR-92 uAHRS[®]: Attitude and Heading Reference System. Technical data supplied by Litef GmbH Freiburg, Germany. February 1997.

Lithopoulos, E., B. Reid, and J. Hutton (1999). Direct Georeferencing with Minimal Ground Control. *GIM International*, 58-61, June.

Liu, L. T., H. T. Hsu, Y. Z. Zhu and J. K. Ou (1999). A New Approach to GPS Ambiguity Decorrelation. *Journal of Geodesy* (1999) 73: 478-490.

Liu, Z. W., M. D. Reed and D. R. Lapucha (1997). GPS Gyro Integration for Airborne Attitude Reference. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 215-221.

Logan, S. A. (1997). Integration of Fibre Optic Gyroscopes and other Sensors with GPS Phase Measurements for Navigation and Route Mapping. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 441-452.

Longstaff, I., L. Burton, D. Tieben and C. Koop (1996). Multi-application GPS/Inertial Navigation Software. *Proceedings of the Satellite Division of the Institute of Navigation 9th International Technical Meeting: ION GPS-96*, September 17-20, 1996, 983-988.

Lucas, J.R. (1987). Aerotriangulation Without Ground Control. *Photogrammetric Engineering and Remote Sensing*, Vol. 53, No. 3, March 1987, 311-314.

Mader, G. L. (1990). Ambiguity Function Techniques for GPS Phase Initialisation and Kinematic Solutions. *Proceeding of the Second International Symposium on Precise Positioning with the Global Positioning System*. Ottawa, September 3-7, The Canadian Institute of Surveying and Mapping, 1233-1247.

- Mader, G. L. (1992). Rapid Static and Kinematic Global Positioning System Solutions using the Ambiguity Function Technique. *Journal of Geophysical Research*, Vol. 97, No. B3, March 1992, 3271-3283.
- Mader, G. L. and J. R. Lucas (1989). Verification of Airborne Positioning Using Global Positioning System Carrier Phase Measurements. *Journal of Geophysical Research*, Vol. 94, No. B8, August 1989, 10,175-10,181.
- Mahmud, M.R., (1999). Precise Three-Dimensional Attitude Estimation from Independent GPS Arrays. PhD Thesis, University of London, UK.
- Malys, S., J. A. Slater, R. W. Smith, L. E. Kunz, S. C. Kenyon (1997). Status of the World Geodetic System 1984. *Proceeding of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 25-34.
- Martin, M.K., and B.C. Detterich (1997a). C-Migits™ II Design and Performance. *Proceedings of the Satellite Division of the Institute of Navigation 10th International Technical Meeting: ION GPS-97*, September 16-19, 1997, 95-102.
- Martin, M.K., and B.C. Detterich (1997b). Digital Quartz IMU-Navigation Processor (DQI-NP) Design and Performance. *Proceedings of the Satellite Division of the Institute of Navigation 10th International Technical Meeting: ION GPS-97*, September 16-19, 1997, 161-166.
- Masson, A., D. Burtin and M. Sebe (1996). Kinematic DGPS and INS Hybridization for Precise Trajectory Determination. *Proceedings of the Satellite Division of the Institute of Navigation 9th International Technical Meeting: ION GPS-96*, September 17-20, 1996, 965-973.
- Mathes, A. and E. Groten (1997). Kinematic Levelling and Positioning with an Integrated GPS-GLONASS-INS-LASER System. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 177-184.
- Maybeck, P. S. (1979). *Stochastic Models, Estimation and Control*. Volume 1. Academic Press, New York, San Francisco, London, 1979.
- Merchant, D.C.(1993).GPS Controlled Aerial Photogrammetry. *Photogrammetric Engineering and Remote Sensing*, Vol. 59, No. 11, November 1993, 1633-1636.
- Mohamed, A.H., and K.P. Schwarz (1999). Adaptive Kalman Filtering for INS/GPS. *Journal of Geodesy* (1999) 73: 193-203.
- Mok, E. (1996). A New Ambiguity Function Algorithm for Short Baseline GPS Engineering Surveying Applications. Ph.D. thesis, Department of Geomatics, University of Newcastle, Newcastle upon Tyne, 189 pp.

Mostafa, M.M.R, K.P. Schwarz and P. Gong (1997). A Fully Digital System for Airborne Mapping. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 463-471.

Mostafa, M.M.R., K.P. Schwarz and M.A. Chapman (1998). Development and Testing of an Airborne Remote Sensing Multi-Sensor System. . *Proceedings of ISPRS Commission II Symposium*, Cambridge, UK, July 13-17, 217-222.

Mostafa, M.M.R. and K.P. Schwarz (1999). A GPS/INS/Imaging System for Kinematic Mapping in Fully Digital Mode. In K.-P. Schwarz (Ed) *Proceedings of the IAG Symposia. Vol. 121 Geodesy Beyond 2000 - The Challenges of the First Decade*. Birmingham, UK, July 19-30, 1999, Springer Verlag, 331-336.

Murphy, J. G. and W. V. Cottrell (1997). Airborne Testing of GPS+GLONASS Positioning Sensor Against a Proven Flight Test Truth Source. *Proceedings of the Satellite Division of the Institute of Navigation 10th International Technical Meeting: ION GPS-97*, September 16-19, 1997, 1047-1054.

Nadler, A. and I. Y. Bar-Itzhack (1998). An Efficient Algorithm for Attitude Determination using GPS. *Proceedings of the Satellite Division of the Institute of Navigation 11th International Technical Meeting: ION GPS-98*, September, 1998, 1783-1789.

Napier, M.E. and V. Ashkenazi (1987). Modern Navigation and Positioning Techniques. *The Journal of Navigation*, Vol. 40, No. 2.

Pio, R. L. (1966). Euler Angle Transformations. *IEEE Transactions on Automatic Control*, Vol. Ac-11, No. 4, October, 1966, 707-715.

Pope, A.J, (1976). The Statistics of Residuals and the Detection of Outliers. *NOAA Technical Report NOS 65 NGS 1*, National Oceanic and Atmospheric Administration, Rockville, Maryland, USA, 133pp.

Ray, J. K. (1999). Use of Multiple Antennas to Mitigate Carrier Phase Multipath in Reference Stations. *Proceedings of the Satellite Division of the Institute of Navigation 12th International Technical Meeting: ION GPS-99*, September 14-17, 1999, 269-279.

Ray, J. K., O. S. Salychev and M. E. Cannon (1999). The Modified Wave Estimator as an Alternative to a Kalman filter for real-time GPS/GLONASS-INS Integration. *Journal of Geodesy* (1999) 73: 568-576.

Remondi, B. W. (1985). Performing Centimeter-Level Surveys in Seconds with GPS Carrier Phase: Initial results. *Navigation: Journal of the Institute of Navigation*, 32(4), 386-400

Remondi, B. (1991). Kinematic GPS Results Without Static Initialisation. *NOAA Technical Memorandum NOS NGS-55*, National Geodetic Information Center, Rockville, MD, USA, 1991.

Rossbach, U. and G. W. Hein (1996). Treatment of Integer Ambiguities in DGPS/DGLONASS Double Difference Carrier Phase Solutions. *Proceedings of the Satellite Division of the Institute of Navigation 9th International Technical Meeting: ION GPS-96*, September 17-20, 1996, 909-916.

Santerre, R. (1991). Impact of GOS Satellite Sky Distribution. *Manuscripta Geodetica* (1991) 16: 28-53.

Schelpe, J. B. (1997). A Real-Time Attitude Determination System using a Quaternion Parameterization. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 395-408.

Scherzinger, B. M. (1997). A Position and Orientation Post-Processing Software Package for Inertial/GPS Integration (POSProc). *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 197-204.

Schwarz, K.P., M. A. Chapman, M.W. Cannon and P. Gong (1993). An Integrated INS/GPS Approach to the Georeferencing of Remotely Sensed Data. *Photogrammetric Engineering and Remote Sensing*, Vol. 59, No. 11, November 1993, 1667-1674.

Schwarz, K. P., M. A. Chapman, M. W. Cannon, P. Gong and D. Cosandier (1994) A Precise Positioning/Attitude System in Support of Airborne Remote Sensing. *ISPRS, Commission II, WGII/1 – Session 2*, Ottawa, Canada, June 6-10, 1994, 191-201.

Schwarz, K.P. and J. Sun (1998). Development and Testing of a Real-Time Integrated INS/DGPS. *Proceedings of the ISPRS Commission II Symposium*, Cambridge, UK, July 13-17, 262-268.

Seeber, G. (1993). *Satellite Geodesy*. de Gruyter, Berlin and New York

Shaw, M., D. Turner and K. Sandhoo (2000). GPS Modernization. Presentation made at *GNSS 2000*, Edinburgh, May 1-4, 2000. In press.

Sheridan, K.F., P.A. Cross and M.R. Mahmud (1999) GPS Based Attitude Determination for Airborne Remote Sensing. In K.-P. Schwarz (Ed) *Proceedings of the IAG Symposia Vol. 121 Geodesy Beyond 2000 - The Challenges of the First Decade*. Birmingham, UK, July 19-30, 1999, Springer Verlag, 337-342.

Sheridan, K.F. and P.A. Cross (2000). A Robust and Reliable GPS-based System for the Direct Georeferencing of Airborne Line Scanning Sensors. Presented at *GNSS 2000*, Edinburgh, UK, May 1-4, 2000. 675-684.

Skaloud, J. (1995). Strapdown INS Orientation Accuracy with GPS Aiding. MSc thesis, Department of Geomatic Engineering, University of Calgary (UCGE Report No. 20079), 1995.

Skaloud, J. (1998). Reducing the GPS Ambiguity Search Space by including Inertial Data. *Proceedings of the Satellite Division of the Institute of Navigation 11th International Technical Meeting: ION GPS-98*, September, 1998, 2073-2080.

Skaloud, J., M. Cramer and K.P. Schwarz (1996). Exterior Orientation by Direct Measurement of Camera Position and Attitude. *International Archives of Photogrammetry and Remote Sensing*. Vol. XXXI, Part B3. Vienna 1996.

Skaloud, J., Y.C. Li, and K.P. Schwarz (1997). Airborne Testing of a C-MIGITS II – Low Cost Integrated GPS/INS. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 161-166.

Skaloud, J., and K.P. Schwarz (1998). Accurate Orientation for Airborne Mapping Systems. *Proceedings of the ISPRS Commission II Symposium*, Cambridge, UK, July 13-17, 283-290.

Song, I. (1997). Measurement Noise and Multipath Reduction with Combined Observables. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 79-87.

Sun, H. (1994). GPS/INS Integration for Airborne Applications. MSc thesis, Department of Geomatic Engineering, University of Calgary (UCGE Report No. 20069), 1994.

Tait, D.A. (1990). *North Seeking Instruments and Inertial Systems*. In T.J.M. Kennie and G. Petrie (Eds) *Engineering Surveying Technology*. Blackie, 1990.

Thompson, E. H. (1969). *An Introduction to the Algebra of Matrices with some Applications*. Adam Hilger, London.

Teunissen, P. J. G. (1993). Least-Squares Estimation of the Integer GPS Ambiguities. Invited Lecture, Section IV Theory and Methodology, *IAG General Meeting*, Beijing.

Teunissen, P. J. G. (1997a). GPS Double Difference Statistics: with and without using Satellite Geometry. *Journal of Geodesy* (1997) 71: 137-148.

Teunissen, P. J. G. (1997b). On the Sensitivity of the Location, Size and Shape of the GPS Ambiguity Search Space to Certain Changes in the Stochastic Model. *Journal of Geodesy* (1997) 71: 541-551.

Tiberius, C. C. J. M. (1998). Quality Control in Positioning. *The Hydrographic Journal*, No. 90, October 1998, 3-8.

Tiberius, C., N. Jonkman and F. Kenselaar (1999). The Stochastics of the GPS Observables. *GPS World*, February, 1999, 49-54.

Toth, C.K., and D.A. Grejner-Brzezinska (1998). Performance Analysis of the Airborne Integrated Mapping System (AIMSTM). *Proceedings of the ISPRS Commission II Symposium*, Cambridge, UK, July 13-17, 320-326.

Toth, C.K., and D.A. Grejner-Brzezinska (1999). Sensor Integration for Improved Surface Modelling. *GIM International*, June 1999.

Ulmer, K. and S. Deines (1997). Stingray: Precise GPS Position, Navigation and Time with Azimuth/Attitude Determination Capability. *Proceedings of the Satellite Division of the Institute of Navigation 10th International Technical Meeting: ION GPS-97*, September 16-19, 1997, 151-160.

Vinnins, M. and L. D. Gallop (1997). GPS as a Method of Heading Determination in a Short Baseline Application in the Canadian Arctic. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 389-394.

Vinnins, M. and L. D. Gallop (1998). A Real-Time Short-Baseline Azimuth Determination System Using GPS. *Proceedings of the Satellite Division of the Institute of Navigation 11th International Technical Meeting: ION GPS-98*, September, 1998, 1763-1772.

Wang, J., M.P. Stewart and M Tsakiri (1998). A Discrimination Test Procedure for Ambiguity Resolution On-the-fly. *Journal of Geodesy* (1998) 72: 644-653.

Wang, J., M.P. Stewart and M Tsakiri (1999). Adaptive Kalman Filtering for Integration of GPS with GLONASS and INS. In K.-P. Schwarz (Ed) *Proceedings of the IAG Symposia, Vol. 121 Geodesy Beyond 2000 - The Challenges of the First Decade*. Birmingham, UK, July 19-30, 1999, Springer Verlag, 325-330.

Wei, M. and K. P. Schwarz (1989). A Discussion of Models for GPS/INS Integration. In Y. Bock and N. Leppard (Eds) *Proceedings of the IAG Symposia, Vol. 102. Global Positioning System: Overview*. Edinburgh, UK, August, 1989, Springer Verlag, 316-327.

Wei, M. and K. P. Schwarz (1998). Flight Test Results from a Strapdown Gravity System. *Journal of Geodesy* (1998) 72: 323-332.

Westrop, J. M., M. E. Napier and V. Ashkenazi (1989). The Use of Phase for Kinematic Positioning by GPS. In Y. Bock and N. Leppard (Eds) *Proceedings of the IAG Symposia, Vol. 102. Global Positioning System: Overview*. Edinburgh, UK, August, 1989, Springer Verlag, 334-339.

Wilson, A. K., (1995). NERC Scientific Services Airborne Remote Sensing Facility User Guide Handbook. Version 1, NERC, 1995.

Wilson, A.K., (1997). An Integrated Data System for Airborne Remote Sensing. *International Journal of Remote Sensing*, Vol. 18, No.9: 1889-1901.

Wolf, P.R, (1983). *Elements of Photogrammetry*. Second Edition. McGraw-Hill.

Wolf, R., B. Eissfeller, G. Hein (1997) . A Kalman Filter for the Integration of a Low Cost INS and Attitude System. *Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97)*, Banff, Canada, June 3-6, 143-150.

Zhodzishsky, M., P. Zagnetov, A. Zinoviev and J. Ashjaee (1999). On Reducing the Impact of GDOP Discontinuities on Receiver Positioning Accuracy. *Proceedings of the Satellite Division of the Institute of Navigation 12th International Technical Meeting: ION GPS-99*, September 14-17, 1999, 1629-1634.

Appendix A Partial Derivatives for Modified Double Differencing

This section presents the formulae for computing the partial derivatives needed to construct the design matrix A, for the modified double differenced phase observation equation, used to relate GPS phase observations directly to attitude components. All the notation used is described in chapter 3. This information comes from Mahmud (1999).

$$\begin{aligned}
 \frac{\partial \rho_{AB}^{ij}}{\partial \psi} = & \left\{ \frac{1}{2} \left[(x_A^i)^2 + (y_A^i)^2 + (z_A^i)^2 \right]^{-1/2} \frac{\partial}{\partial \psi} \left[(x_A^i)^2 + (y_A^i)^2 + (z_A^i)^2 \right] \right\} - \\
 & \left\{ \frac{1}{2} \left[(x_B^i)^2 + (y_B^i)^2 + (z_B^i)^2 \right]^{-1/2} \frac{\partial}{\partial \psi} \left[(x_B^i)^2 + (y_B^i)^2 + (z_B^i)^2 \right] \right\} - \\
 & \left\{ \frac{1}{2} \left[(x_A^j)^2 + (y_A^j)^2 + (z_A^j)^2 \right]^{-1/2} \frac{\partial}{\partial \psi} \left[(x_A^j)^2 + (y_A^j)^2 + (z_A^j)^2 \right] \right\} + \\
 & \left\{ \frac{1}{2} \left[(x_B^j)^2 + (y_B^j)^2 + (z_B^j)^2 \right]^{-1/2} \frac{\partial}{\partial \psi} \left[(x_B^j)^2 + (y_B^j)^2 + (z_B^j)^2 \right] \right\} \quad (A1)
 \end{aligned}$$

$$\begin{aligned}
 \frac{\partial \rho_{AB}^{ij}}{\partial \psi} = & \left\{ \frac{1}{2 \left[(x_A^i)^2 + (y_A^i)^2 + (z_A^i)^2 \right]^{1/2}} \left[2(x_A^i) \frac{\partial}{\partial \psi} (x_A^i) + 2(y_A^i) \frac{\partial}{\partial \psi} (y_A^i) + 2(z_A^i) \frac{\partial}{\partial \psi} (z_A^i) \right] \right\} - \\
 & \left\{ \frac{1}{2 \left[(x_B^i)^2 + (y_B^i)^2 + (z_B^i)^2 \right]^{1/2}} \left[2(x_B^i) \frac{\partial}{\partial \psi} (x_B^i) + 2(y_B^i) \frac{\partial}{\partial \psi} (y_B^i) + 2(z_B^i) \frac{\partial}{\partial \psi} (z_B^i) \right] \right\} - \\
 & \left\{ \frac{1}{2 \left[(x_A^j)^2 + (y_A^j)^2 + (z_A^j)^2 \right]^{1/2}} \left[2(x_A^j) \frac{\partial}{\partial \psi} (x_A^j) + 2(y_A^j) \frac{\partial}{\partial \psi} (y_A^j) + 2(z_A^j) \frac{\partial}{\partial \psi} (z_A^j) \right] \right\} + \\
 & \left\{ \frac{1}{2 \left[(x_B^j)^2 + (y_B^j)^2 + (z_B^j)^2 \right]^{1/2}} \left[2(x_B^j) \frac{\partial}{\partial \psi} (x_B^j) + 2(y_B^j) \frac{\partial}{\partial \psi} (y_B^j) + 2(z_B^j) \frac{\partial}{\partial \psi} (z_B^j) \right] \right\} \quad (A2)
 \end{aligned}$$

which leads to

$$\begin{aligned}
\frac{\partial \rho_{AB}^{ij}}{\partial \psi} = & \left\{ \frac{1}{\rho_A^i} \left[x_A^i \frac{\partial}{\partial \psi} (x_A^i) + y_A^i \frac{\partial}{\partial \psi} (y_A^i) + z_A^i \frac{\partial}{\partial \psi} (z_A^i) \right] \right\} - \\
& \left\{ \frac{1}{\rho_B^i} \left[x_B^i \frac{\partial}{\partial \psi} (x_B^i) + y_B^i \frac{\partial}{\partial \psi} (y_B^i) + z_B^i \frac{\partial}{\partial \psi} (z_B^i) \right] \right\} - \\
& \left\{ \frac{1}{\rho_A^j} \left[x_A^j \frac{\partial}{\partial \psi} (x_A^j) + y_A^j \frac{\partial}{\partial \psi} (y_A^j) + z_A^j \frac{\partial}{\partial \psi} (z_A^j) \right] \right\} + \\
& \left\{ \frac{1}{\rho_B^j} \left[x_B^j \frac{\partial}{\partial \psi} (x_B^j) + y_B^j \frac{\partial}{\partial \psi} (y_B^j) + z_B^j \frac{\partial}{\partial \psi} (z_B^j) \right] \right\}
\end{aligned} \tag{A3}$$

Similarly,

$$\begin{aligned}
\frac{\partial \rho_{AB}^{ij}}{\partial \theta} = & \left\{ \frac{1}{\rho_A^i} \left[x_A^i \frac{\partial}{\partial \theta} (x_A^i) + y_A^i \frac{\partial}{\partial \theta} (y_A^i) + z_A^i \frac{\partial}{\partial \theta} (z_A^i) \right] \right\} - \\
& \left\{ \frac{1}{\rho_B^i} \left[x_B^i \frac{\partial}{\partial \theta} (x_B^i) + y_B^i \frac{\partial}{\partial \theta} (y_B^i) + z_B^i \frac{\partial}{\partial \theta} (z_B^i) \right] \right\} - \\
& \left\{ \frac{1}{\rho_A^j} \left[x_A^j \frac{\partial}{\partial \theta} (x_A^j) + y_A^j \frac{\partial}{\partial \theta} (y_A^j) + z_A^j \frac{\partial}{\partial \theta} (z_A^j) \right] \right\} + \\
& \left\{ \frac{1}{\rho_B^j} \left[x_B^j \frac{\partial}{\partial \theta} (x_B^j) + y_B^j \frac{\partial}{\partial \theta} (y_B^j) + z_B^j \frac{\partial}{\partial \theta} (z_B^j) \right] \right\}
\end{aligned} \tag{A4}$$

and

$$\begin{aligned}
\frac{\partial \rho_{AB}^{ij}}{\partial \gamma} = & \left\{ \frac{1}{\rho_A^i} \left[x_A^i \frac{\partial}{\partial \gamma} (x_A^i) + y_A^i \frac{\partial}{\partial \gamma} (y_A^i) + z_A^i \frac{\partial}{\partial \gamma} (z_A^i) \right] \right\} - \\
& \left\{ \frac{1}{\rho_B^i} \left[x_B^i \frac{\partial}{\partial \gamma} (x_B^i) + y_B^i \frac{\partial}{\partial \gamma} (y_B^i) + z_B^i \frac{\partial}{\partial \gamma} (z_B^i) \right] \right\} - \\
& \left\{ \frac{1}{\rho_A^j} \left[x_A^j \frac{\partial}{\partial \gamma} (x_A^j) + y_A^j \frac{\partial}{\partial \gamma} (y_A^j) + z_A^j \frac{\partial}{\partial \gamma} (z_A^j) \right] \right\} + \\
& \left\{ \frac{1}{\rho_B^j} \left[x_B^j \frac{\partial}{\partial \gamma} (x_B^j) + y_B^j \frac{\partial}{\partial \gamma} (y_B^j) + z_B^j \frac{\partial}{\partial \gamma} (z_B^j) \right] \right\}
\end{aligned} \tag{A5}$$

whereas,

$$\frac{\partial \rho_{AB}^{ij}}{\partial N_{AB}^{ij}} = \lambda \quad (\text{A6})$$

The partial derivatives that are associated with the foregoing equations are as follows :

$$\frac{\partial m_{11}}{\partial \psi} = 0 \quad (\text{A7})$$

$$\frac{\partial m_{11}}{\partial \theta} = -\sin \theta \cos \gamma = m_{13} \cos \gamma \quad (\text{A8})$$

$$\frac{\partial m_{11}}{\partial \gamma} = -\cos \theta \sin \gamma = -m_{12} \quad (\text{A9})$$

$$\frac{\partial m_{12}}{\partial \psi} = 0 \quad (\text{A10})$$

$$\frac{\partial m_{12}}{\partial \theta} = -\sin \theta \sin \gamma = m_{13} \sin \gamma \quad (\text{A11})$$

$$\frac{\partial m_{12}}{\partial \gamma} = \cos \theta \cos \gamma = m_{11} \quad (\text{A12})$$

$$\frac{\partial m_{13}}{\partial \psi} = 0 \quad (\text{A13})$$

$$\frac{\partial m_{13}}{\partial \theta} = -\cos \theta \quad (\text{A14})$$

$$\frac{\partial m_{13}}{\partial \gamma} = 0 \quad (\text{A15})$$

$$\frac{\partial m_{21}}{\partial \psi} = \cos \psi \sin \theta \cos \gamma + \sin \psi \sin \gamma = m_{31} \quad (\text{A16})$$

$$\frac{\partial m_{21}}{\partial \theta} = \sin \psi \cos \theta \cos \gamma = m_{23} \cos \gamma \quad (\text{A17})$$

$$\frac{\partial m_{21}}{\partial \gamma} = -\sin \psi \sin \theta \sin \gamma - \cos \psi \cos \gamma = -(\sin \psi \sin \theta \sin \gamma + \cos \psi \cos \gamma) = -m_{22} \quad (\text{A18})$$

$$\frac{\partial m_{22}}{\partial \psi} = \cos \psi \sin \theta \sin \gamma - \sin \psi \cos \gamma = m_{32} \quad (\text{A19})$$

$$\frac{\partial m_{22}}{\partial \theta} = \sin \psi \cos \theta \sin \gamma = m_{23} \sin \gamma \quad (\text{A20})$$

$$\frac{\partial m_{22}}{\partial \gamma} = \sin \psi \sin \theta \cos \gamma - \cos \psi \sin \gamma = m_{21} \quad (\text{A21})$$

$$\frac{\partial m_{23}}{\partial \psi} = \cos \psi \cos \theta = m_{33} \quad (\text{A22})$$

$$\frac{\partial m_{23}}{\partial \theta} = -\sin \psi \sin \theta = m_{13} \sin \psi \quad (\text{A23})$$

$$\frac{\partial m_{23}}{\partial \gamma} = 0 \quad (\text{A24})$$

$$\frac{\partial m_{31}}{\partial \psi} = -\sin \psi \sin \theta \cos \gamma + \cos \psi \sin \gamma = -(\sin \psi \sin \theta \cos \gamma - \cos \psi \sin \gamma) = -m_{21} \quad (\text{A25})$$

$$\frac{\partial m_{31}}{\partial \theta} = \cos \psi \cos \theta \cos \gamma = m_{33} \cos \gamma \quad (\text{A26})$$

$$\frac{\partial m_{31}}{\partial \gamma} = -\cos \psi \sin \theta \sin \gamma + \sin \psi \cos \gamma = -(\cos \psi \sin \theta \sin \gamma - \sin \psi \cos \gamma) = -m_{32} \quad (\text{A27})$$

$$\frac{\partial m_{32}}{\partial \psi} = -\sin \psi \sin \theta \sin \gamma - \cos \psi \cos \gamma = -(\sin \psi \sin \theta \sin \gamma + \cos \psi \cos \gamma) = -m_{22} \quad (\text{A28})$$

$$\frac{\partial m_{32}}{\partial \theta} = \cos \psi \cos \theta \sin \gamma = m_{33} \sin \gamma \quad (\text{A29})$$

$$\frac{\partial m_{32}}{\partial \gamma} = \cos \psi \sin \theta \cos \gamma + \sin \psi \sin \gamma = m_{31} \quad (\text{A30})$$

$$\frac{\partial m_{33}}{\partial \psi} = -\sin \psi \cos \theta = -m_{23} \quad (\text{A31})$$

$$\frac{\partial m_{33}}{\partial \theta} = -\cos \psi \sin \theta = m_{13} \cos \psi \quad (\text{A32})$$

$$\frac{\partial m_{33}}{\partial \gamma} = 0 \quad (\text{A33})$$

Substituting the foregoing equations in the partial derivatives of x_A^i , y_A^i and z_A^i with respect to ψ , θ and γ results in the following :

$$\frac{\partial}{\partial \psi}(x_A^i) = 0 \quad (\text{A34})$$

$$\frac{\partial}{\partial \theta}(x_A^i) = -m_{13}x_A^{\text{BFS}} \cos \gamma - m_{13}y_A^{\text{BFS}} \sin \gamma + z_A^{\text{BFS}} \cos \theta \quad (\text{A35})$$

$$\frac{\partial}{\partial \gamma}(x_A^i) = m_{12}x_A^{\text{BFS}} - m_{11}y_A^{\text{BFS}} \quad (\text{A36})$$

$$\frac{\partial}{\partial \psi}(y_A^i) = -m_{31}x_A^{\text{BFS}} - m_{32}y_A^{\text{BFS}} - m_{33}z_A^{\text{BFS}} \quad (\text{A37})$$

$$\frac{\partial}{\partial \theta}(y_A^i) = -m_{23}x_A^{\text{BFS}} \cos \gamma - m_{23}y_A^{\text{BFS}} \sin \gamma - m_{13}z_A^{\text{BFS}} \sin \psi \quad (\text{A38})$$

$$\frac{\partial}{\partial \gamma}(y_A^i) = m_{22}x_A^{\text{BFS}} - m_{21}y_A^{\text{BFS}} \quad (\text{A39})$$

$$\frac{\partial}{\partial \psi}(z_A^i) = m_{21}x_A^{\text{BFS}} + m_{22}y_A^{\text{BFS}} + m_{23}z_A^{\text{BFS}} \quad (\text{A40})$$

$$\frac{\partial}{\partial \theta}(z_A^i) = -m_{33}x_A^{\text{BFS}} \cos \gamma - m_{33}y_A^{\text{BFS}} \sin \gamma - m_{13}z_A^{\text{BFS}} \cos \psi \quad (\text{A41})$$

$$\frac{\partial}{\partial \gamma}(z_A^i) = m_{32}x_A^{\text{BFS}} - m_{31}y_A^{\text{BFS}} \quad (\text{A42})$$

The partial derivatives for the satellite/antenna combinations iB, jA and jB can be derived following identical procedures.

Appendix B GPS Data Description

This section provides information on various aspects of the GPS data recorded during the selected flight lines. For both days on which test flights were conducted, a period which encompasses all the selected flight lines for that day, has been considered. For day 288 (15-10-98), the chosen period is 387900 to 390000, expressed as GPS seconds of the week, or 11:45:00 to 12:20:00. For day 289 (16-10-98), the chosen period is 467700 to 468900, or 9:55:00 to 10:15:00.

Tables B.1 and B.2 present details of the data recorded by each of the independent aircraft receivers on days 288 and 289 respectively. The number of satellites tracked at each epoch is summarised, and the 'data' and 'no data' figures indicate the percentage of the complete observation period at which any satellite data were recorded. The 'from' figure is the number of epochs in a complete data-set. In both tables the 'from' figure for receiver 4 is half that of the other receivers as GPS data is recorded at 5Hz on this receiver and 10Hz on the other three.

Total Tracked Satellites	Receiver 1		Receiver 2		Receiver 4	
	Epochs	%	Epochs	%	Epochs	%
7	1	0.005	0	0	0	0
8	482	2.30	160	0.83	97	0.92
9	1957	9.32	1750	9.04	1234	11.75
10	13378	63.71	17442	90.13	8133	77.46
11	0	0	0	0	0	0
Data	15818	75.33	19532	100	9464	90.13
No Data	5182	24.67	0	0	1036	9.87
From	21000		19532 ¹		10500	

Table B.1: GPS data logging summary. Day 288 (15-10-98)

¹ The complete data-set for receiver 2 on day 288 does not cover the entire observation period as this receiver stopped logging at 389835.1 (12:17:15.1) when the memory card was full. Data are available from this receiver for all of the flight lines but the period chosen to assess observation conditions extends beyond the final flight line (L008).

Total Tracked Satellites	Receiver 1		Receiver 2		Receiver 3		Receiver 4	
	Epochs	%	Epochs	%	Epochs	%	Epochs	%
7	9	0.07	3	0.03	1	0.01	0	0
8	1346	11.22	1312	10.93	485	4.04	1218	20.30
9	5750	47.92	5328	44.40	3851	32.09	3724	62.07
10	4218	35.15	5242	43.68	1166	9.72	885	14.75
11	7	0.06	115	0.96	1	0.01	0	0
Data	11330	94.42	12000	100	5504	45.87	5827	97.12
No Data	670	5.58	0	0	6496	54.13	173	2.88
From	12000		12000		12000		6000	

Table B.2: GPS data logging summary. Day 289 (16-10-98)

The completeness of these data sets clearly affects the frequency with which positions and attitudes can be determined. In order to compute three dimensional attitude components, data must be available from at least three receivers at an epoch. On day 288, there are no GPS data available from antenna 3, the port winglet antenna, due to logging problems with the Ashtech Z-Sensor and PC configuration. As a result, a full attitude solution is not available at 10Hz as data from the 5Hz Ashtech Z-12 receiver connected to antenna 4 must be used in the solution.

Figures B.1 to B.3 show the total number of satellites from which observations were recorded by each of the three usable aircraft receivers during day 288. Figures B.4 to B.7 show the same information for day 289 when data from all four antennas were logged some of the time. The gaps noted above do not show up on these plots except in figure 6.6 where the diagonal lines are caused by missing data.

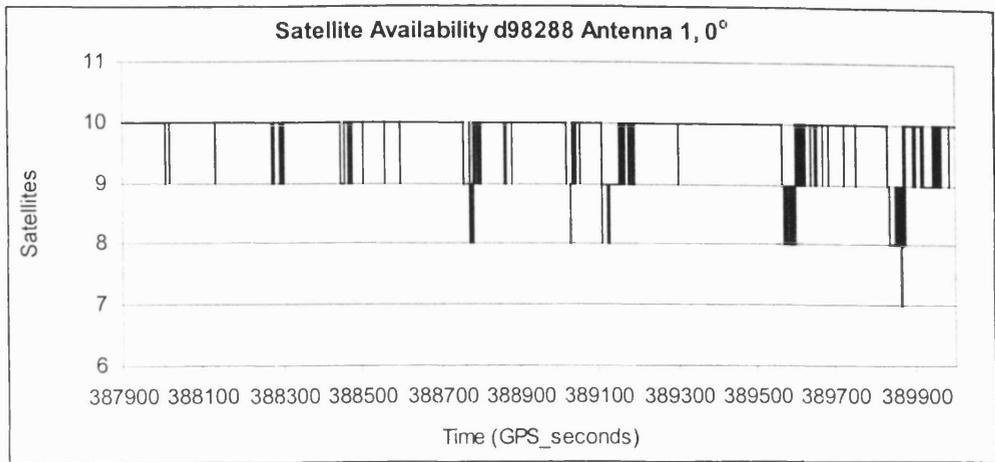


Figure B.1: Number of satellites from which observations were recorded at antenna 1 during the observation period on day 288 (15-10-98)

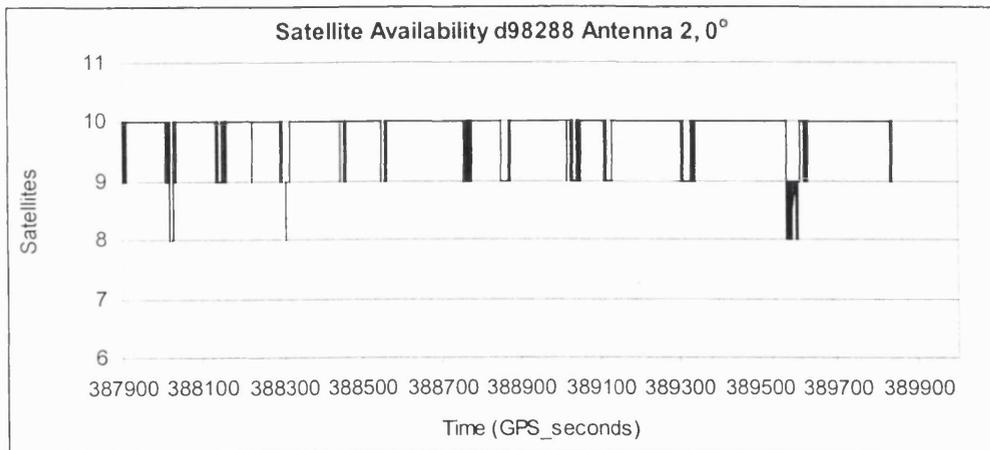


Figure B.2: Number of satellites from which observations were recorded at antenna 2 during the observation period on day 288 (15-10-98)

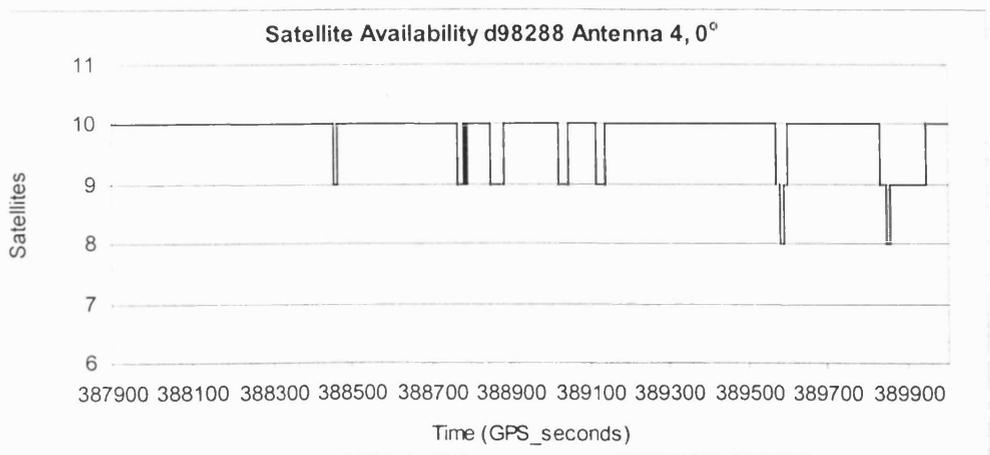


Figure B.3: Number of satellites from which observations were recorded at antenna 4 during the observation period on day 288 (15-10-98)

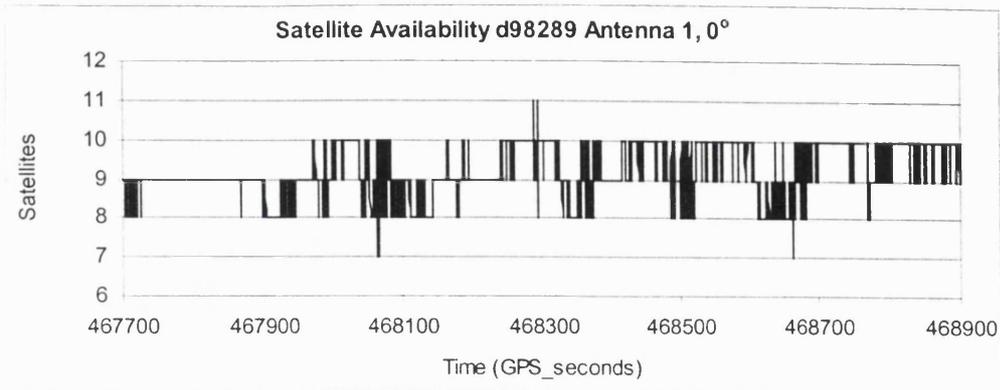


Figure B.4: Satellite observations recorded by receiver 1. Day 289 (16-10-98)

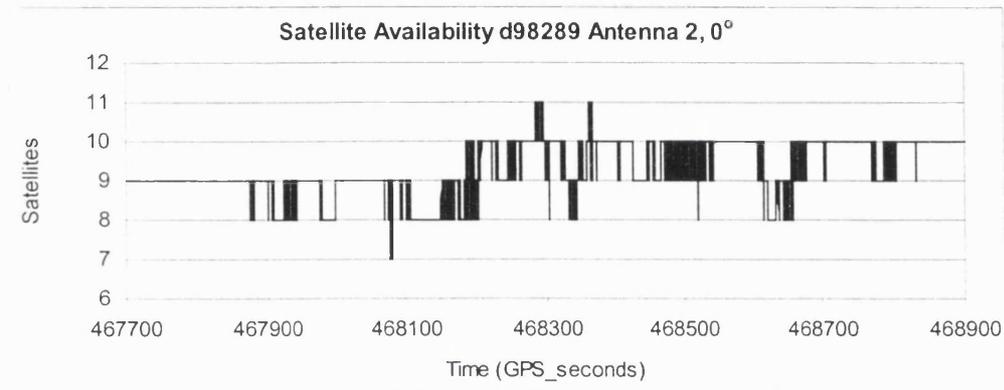


Figure B.5: Satellite observations recorded by receiver 2. Day 289 (16-10-98)

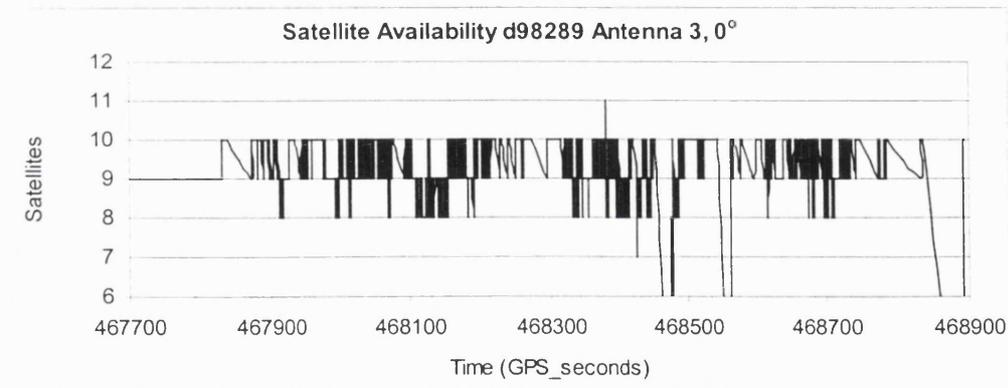


Figure B.6: Satellite observations recorded by receiver 3. Day 289 (16-10-98)

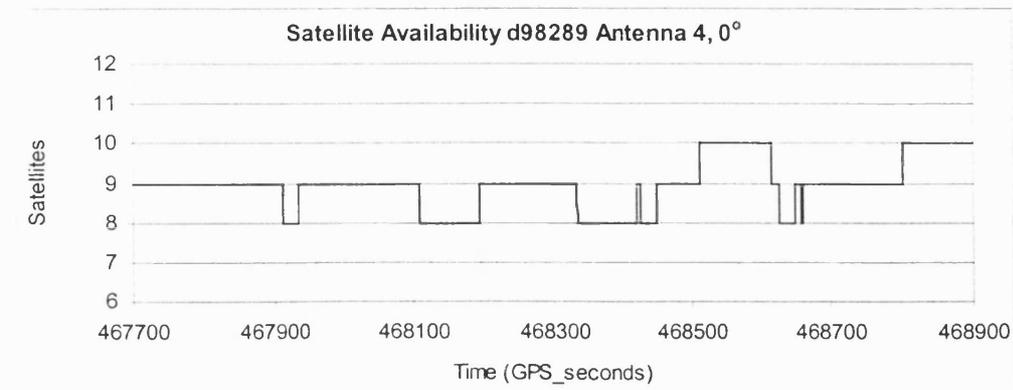


Figure B.7: Satellite observations recorded by receiver 4. Day 289 (16-10-98)

At the ground reference station, GPS data were recorded at 10Hz for a period covering the entire flight on each day. The Leica MC1000 receiver operates with a minimum elevation mask of 10°. No data gaps occurred in either set of reference data. Figure B.8 shows the number of satellites tracked at the reference station on day 288, for the selected observation period. For 87.7% of this period, observations from eight satellites were recorded, nine satellites were tracked for the remainder of the time. Figure B.9 shows the number of satellites tracked at the reference station on day 289, for the selected observation period. For 63.9% of this period, observations from seven satellites were recorded, eight satellites were tracked for the remaining 36.1% of the time.

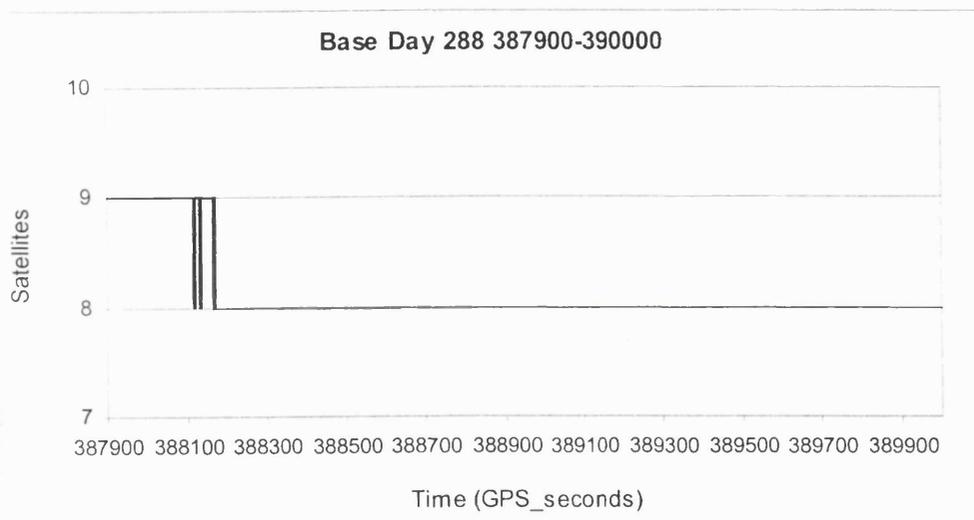


Figure B.8: Number of satellites tracked at the reference station on day 288.

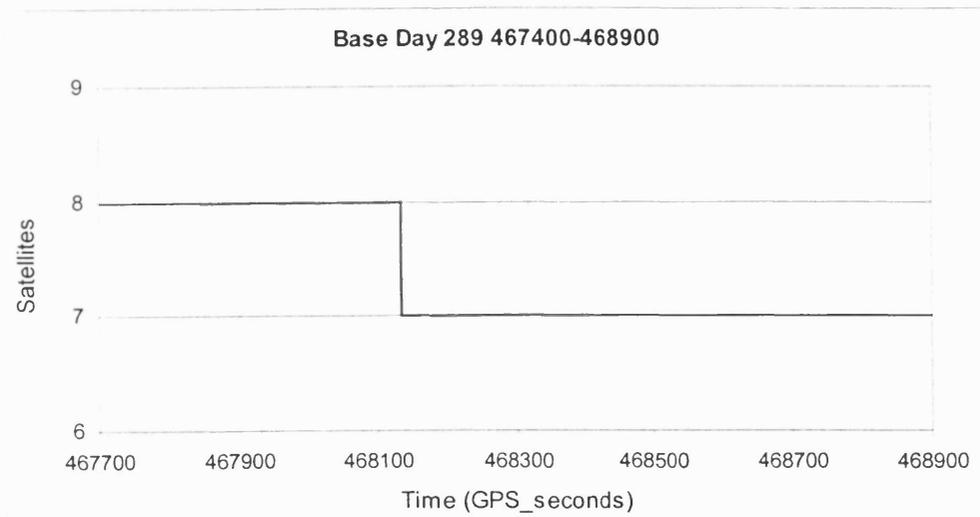


Figure B.9: Number of satellites tracked at the reference station on day 289.

Figure B.10 and table B.3 show the elevation and azimuth of the satellites tracked during the observation period on day 288. In table B.3, the elevation and azimuth of each satellite are given for the first and last epochs of the observation period (in GPS seconds of week). The sky plot (B.10) uses the average satellite azimuth and elevation for the observation period. Due to the short observation period, the satellite geometry does not change significantly, therefore this ‘snap-shot’ should adequately describe the satellite geometry for each flight line recorded during this period.

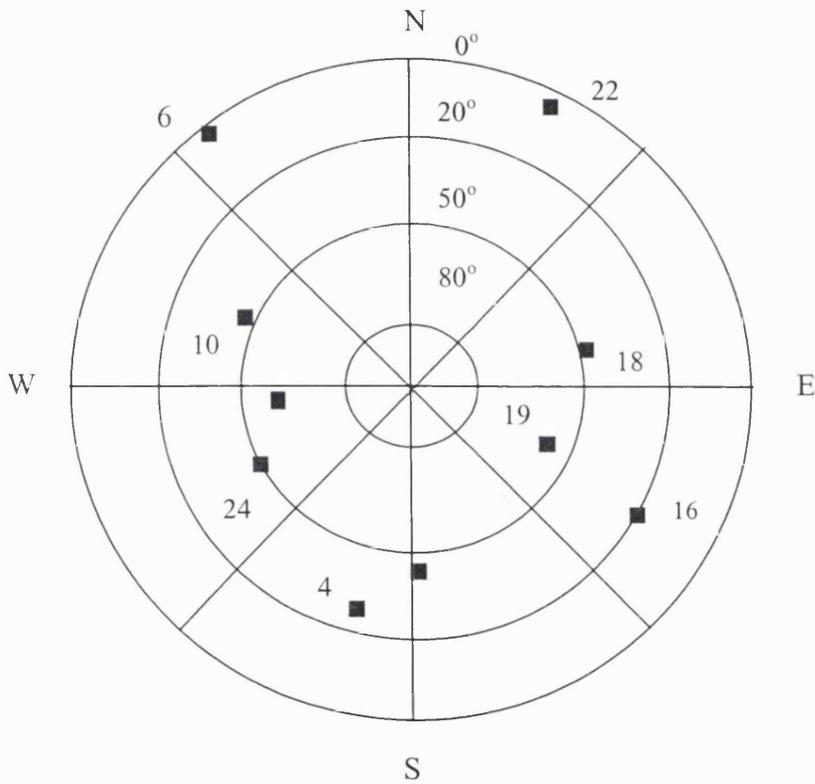


Figure B.10: Sky plot, day 288.

PRN	Elevation		Azimuth	
	387900	390000	387900	390000
19	59	63	133	104
18	56	42	73	77
13	56	70	250	262
24	54	46	257	236
10	38	52	296	296
4	37	23	199	194
27	35	50	181	177
16	27	14	118	124
22	10	5	32	21
6	9	7	328	316

Table B.3: Satellite elevations & azimuths at start and end of day 288 observation period.

Figure B.11 and table B.4 show the elevation and azimuth of the satellites tracked during the observation period on day 289². In table B.4, the elevation and azimuth of each satellite are given for the first and last epochs of the observation period (in GPS seconds of week). The sky plot (B.11) uses the average satellite azimuth and elevation for the observation period.

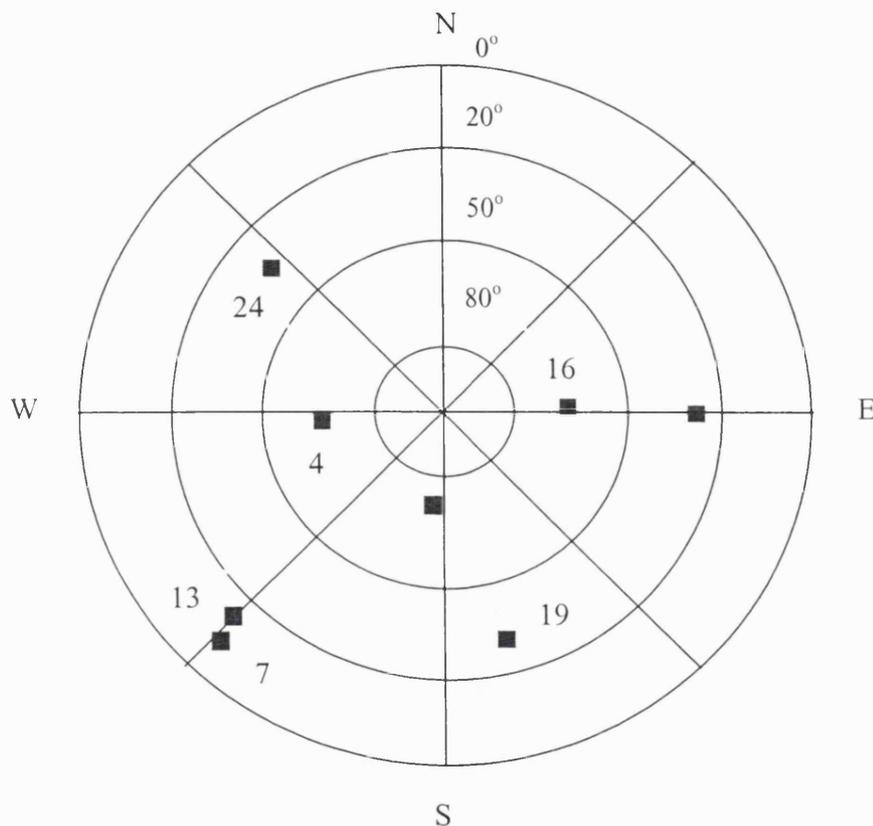


Figure B.11: Sky plot, day 289.

PRN	Elevation		Azimuth	
	467700	468900	467700	468900
16	72	64	83	92
18	70	79	202	189
4	65	67	274	251
24	29	37	310	307
14	29	21	93	98
19	19	25	166	163
7	12	5	223	219
13	10	19	226	230

Table B.4: Satellite elevations & azimuths at start and end of day 289 observation period.

² Three additional satellites are tracked intermittently during this period at elevations below 5°. They are not detailed here as they are not used in further processing

Figure B.12 and B.13 show the elevations of the tracked satellites during the observation periods on day 288 and 289 respectively.

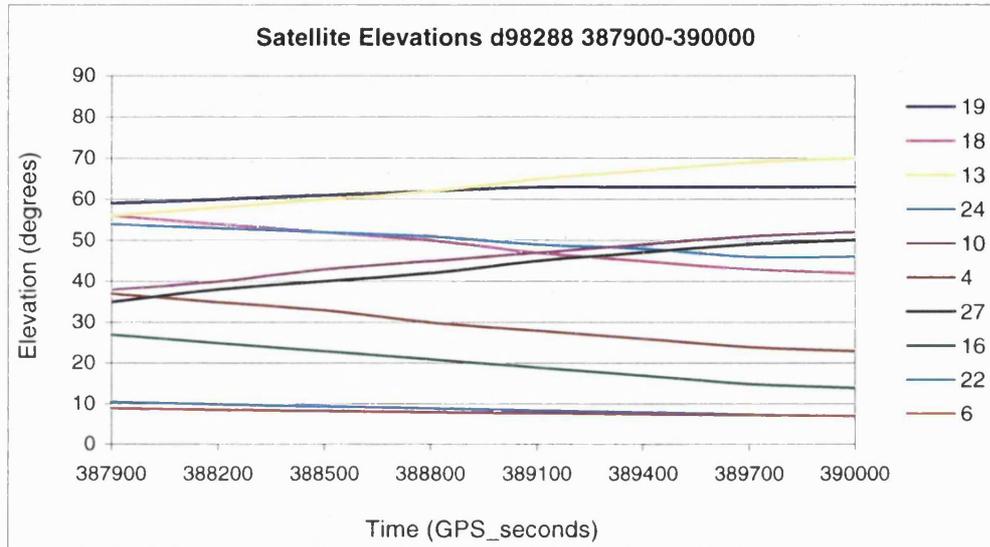


Figure 6.12: Satellite elevations for the observation period on day 288 (15-10-98)

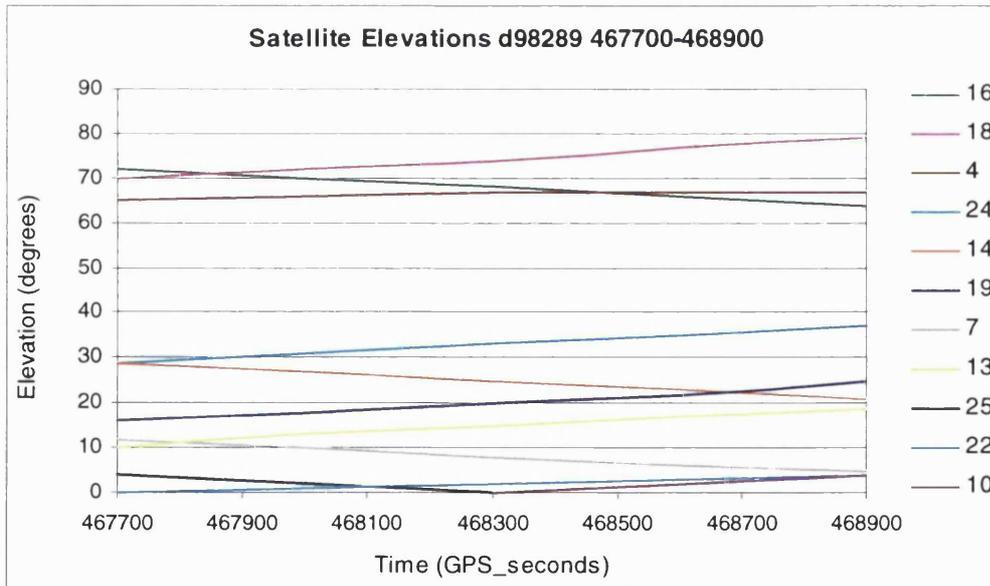


Figure 6.13: Satellite elevations for the observation period on day 289 (16-10-98)

Figures B.14 and B.15 show the PDOP for the observation periods on day 288 and 289 respectively, using a 10° and 20° elevation mask. The ‘steps’ in PDOP values are the results of satellites being included and excluded from the solution as their elevations change over time.

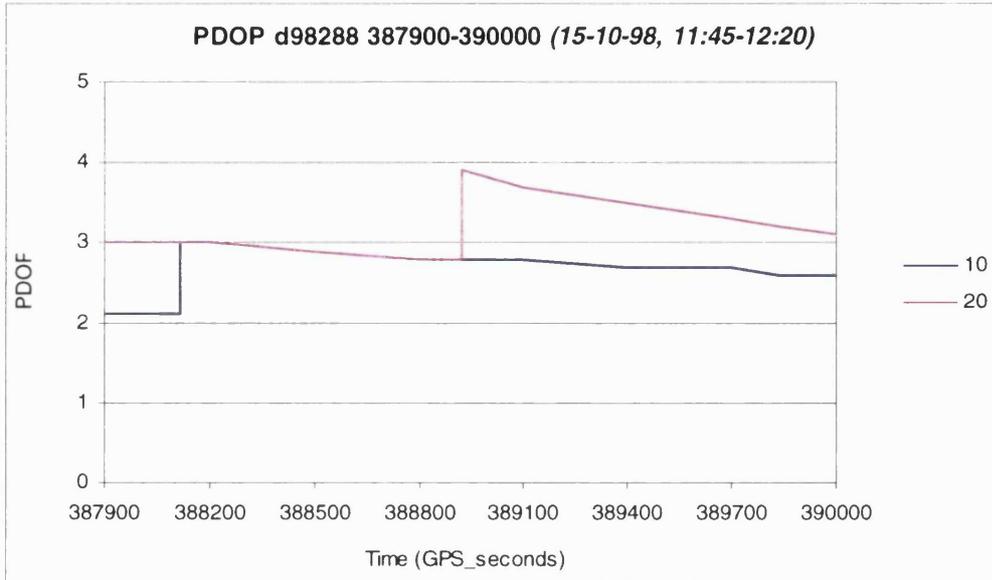


Figure B.14: PDOP for the observation period on day 288 (15-10-98)

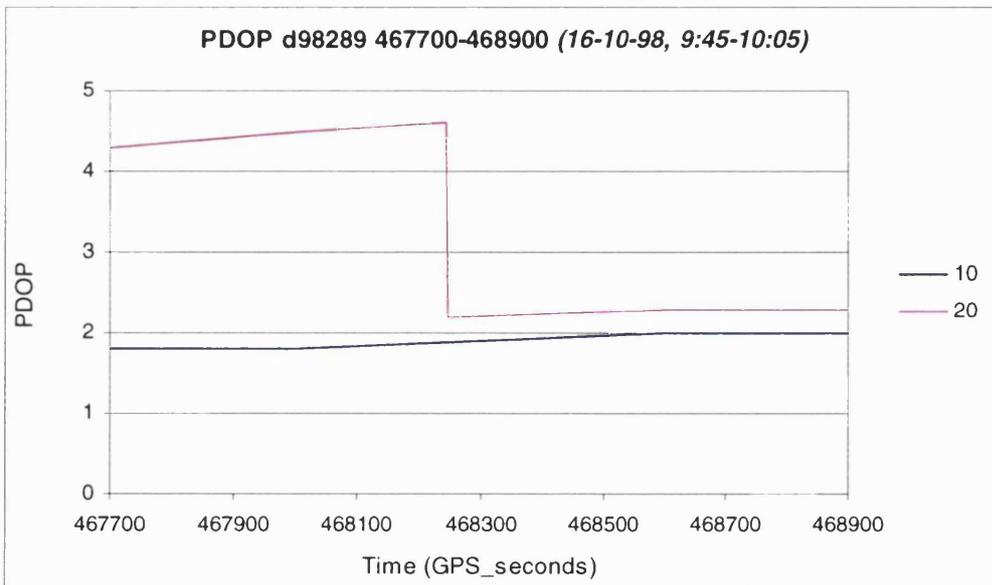


Figure B.15: PDOP for the observation period on day 289 (16-10-98)

The PDOP is usually between 2 and 4 during both observation periods, indicating that the satellite geometry is generally good during the flight lines.

Appendix C GPS Phase Residuals

This appendix consists of twelve tables which give details of the single differenced phase residuals during the flight lines L003, L006 and L008 on day 288. These values were computed by SKI when aircraft antennas were positioned relative to the local reference station. For each flight line, the L1 and L2 phase residuals for single differences between the reference station and antennas 2 and 4 are shown. Each table uses the following structure:

Row 1	PRN number of the satellite used in the single difference,
Row 2	satellite elevation in degrees,
Row 3	mean phase residual per flight line, in metres,
Row 4	minimum phase residual per flight line, in metres,
Row 5	maximum phase residual per flight line, in metres,
Row 6	standard deviation (σ) of phase residuals per flight line, in metres,
Row 7	range of phase residuals per flight line, in metres.

This record is repeated for each satellite used in the solution.

Tables C.1 and C.2 present details of the L1 and L2 phase residuals for the single differenced observations between the reference station and antenna 2 during L003. Tables C.3 and C.4 present details of the L1 and L2 phase residuals for the single differenced observations between the reference station and antenna 4 during L003.

Tables C.5 and C.6 present details of the L1 and L2 phase residuals for the single differenced observations between the reference station and antenna 2 during L006. Tables C.7 and C.8 present details of the L1 and L2 phase residuals for the single differenced observations between the reference station and antenna 4 during L006.

Tables C.9 and C.10 present details of the L1 and L2 phase residuals for the single differenced observations between the reference station and antenna 2 during L008. Tables C.11 and C.12 present details of the L1 and L2 phase residuals for the single differenced observations between the reference station and antenna 4 during L008.

PRN	13	19	10	27	24	18	4	16
Elevation	62	62	52	51	45	42	31	22
Mean	0.000	0.003	-0.003	-0.002	-0.003	-0.005	0.010	0.016
Min	-0.004	-0.001	-0.007	-0.006	-0.006	-0.010	0.004	0.008
Max	0.004	0.007	-0.001	0.002	0.001	0.000	0.016	0.024
σ	0.001	0.002	0.001	0.001	0.001	0.002	0.002	0.003
Range	0.008	0.008	0.006	0.008	0.007	0.010	0.012	0.016

Table C.1: L1 phase residuals between reference and antenna 2 during L003.

PRN	13	19	10	27	24	18	4	16
Elevation	62	62	52	51	45	42	31	22
Mean	0.005	0.001	0.004	0.000	-0.001	-0.010	0.002	-0.007
Min	-0.001	-0.004	-0.001	-0.006	-0.005	-0.016	-0.009	-0.018
Max	0.010	0.005	0.010	0.005	0.005	-0.001	0.010	0.006
σ	0.002	0.002	0.002	0.002	0.002	0.003	0.003	0.004
Range	0.011	0.009	0.011	0.011	0.010	0.015	0.019	0.024

Table C.2: L2 phase residuals between reference and antenna 2 during L003.

PRN	13	19	10	27	24	18	4	16
Elevation	62	62	52	51	45	42	31	22
Mean	-0.002	-0.007	0.000	0.000	0.002	-0.002	0.015	0.002
Min	-0.006	-0.014	-0.003	-0.008	-0.002	-0.005	0.009	-0.007
Max	0.002	0.001	0.004	0.008	0.008	0.004	0.022	0.013
σ	0.002	0.002	0.001	0.003	0.002	0.002	0.002	0.004
Range	0.008	0.015	0.007	0.016	0.010	0.009	0.013	0.020

Table C.3: L1 phase residuals between reference and antenna 4 during L003.

PRN	13	19	10	27	24	18	4	16
Elevation	62	62	52	51	45	42	31	22
Mean	-0.001	0.009	0.000	0.000	-0.002	0.000	-0.011	-0.007
Min	-0.009	0.003	-0.005	-0.007	-0.008	-0.008	-0.019	-0.019
Max	0.005	0.016	0.008	0.010	0.006	0.005	-0.003	0.007
σ	0.003	0.002	0.003	0.004	0.003	0.002	0.003	0.005
Range	0.014	0.013	0.013	0.017	0.014	0.013	0.016	0.026

Table C.4: L2 phase residuals between reference and antenna 4 during L003.

PRN	13	19	10	27	24	18	4
Elevation	67	63	50	48	48	44	26
Mean	-0.003	0.002	0.004	-0.003	-0.006	0.001	0.013
Min	-0.012	-0.002	-0.002	-0.012	-0.011	-0.003	0.004
Max	0.004	0.009	0.009	0.006	0.002	0.006	0.022
σ	0.003	0.002	0.002	0.003	0.002	0.002	0.003
Range	0.016	0.011	0.016	0.016	0.017	0.012	0.026

Table C.5: L1 phase residuals between reference and antenna 2 during L006.

PRN	13	19	10	27	24	18	4
Elevation	67	63	50	48	48	44	26
Mean	0.005	0.001	0.000	0.002	-0.003	-0.006	-0.001
Min	-0.003	-0.003	-0.007	-0.007	-0.012	-0.012	-0.016
Max	0.011	0.008	0.009	0.009	0.005	0.000	0.010
σ	0.003	0.002	0.003	0.003	0.003	0.002	0.004
Range	0.005	0.011	0.016	0.016	0.017	0.012	0.026

Table C.6: L2 phase residuals between reference and antenna 2 during L006.

PRN	13	19	10	27	24	18	4
Elevation	67	63	50	48	48	44	26
Mean	-0.003	-0.010	0.002	0.005	-0.004	0.011	0.005
Min	-0.012	-0.015	-0.005	-0.003	-0.011	-0.001	-0.001
Max	0.003	-0.003	0.010	0.015	0.002	0.017	0.013
σ	0.003	0.003	0.003	0.003	0.003	0.003	0.003
Range	0.015	0.012	0.015	0.018	0.013	0.018	0.014

Table C.7: L1 phase residuals between reference and antenna 4 during L006.

PRN	13	19	10	27	24	18	4
Elevation	67	63	50	48	48	44	26
Mean	0.007	-0.005	0.000	0.012	-0.002	-0.004	-0.023
Min	0.000	-0.017	-0.009	0.005	-0.009	-0.016	-0.035
Max	0.014	0.010	0.007	0.018	0.005	0.010	-0.010
σ	0.002	0.005	0.003	0.002	0.003	0.005	0.005
Range	0.014	0.027	0.016	0.013	0.014	0.026	0.025

Table C.8: L2 phase residuals between reference and antenna 4 during L006.

PRN	13	19	10	27	24	18	4
Elevation	70	63	52	50	45	42	24
Mean	-0.005	0.003	0.005	-0.004	-0.003	0.001	0.013
Min	-0.011	-0.004	0.000	-0.009	-0.008	-0.004	0.005
Max	0.001	0.008	0.009	0.002	0.007	0.005	0.025
σ	0.002	0.002	0.001	0.002	0.003	0.002	0.003
Range	0.012	0.012	0.009	0.011	0.015	0.009	0.020

Table C.9: L1 phase residuals between reference and antenna 2 during L008.

PRN	13	19	10	27	24	18	4
Elevation	70	63	52	50	45	42	24
Mean	0.002	0.005	-0.001	0.002	-0.003	-0.008	0.001
Min	-0.003	-0.002	-0.006	-0.004	-0.009	-0.012	-0.015
Max	0.008	0.010	0.007	0.008	0.005	-0.003	0.014
σ	0.002	0.002	0.002	0.002	0.002	0.002	0.006
Range	0.011	0.012	0.013	0.012	0.014	0.009	0.029

Table C.10: L2 phase residuals between reference and antenna 2 during L008.

PRN	13	19	10	27	24	18	4
Elevation	70	63	52	50	45	42	24
Mean	-0.003	-0.005	-0.001	0.002	-0.002	0.010	0.003
Min	-0.007	-0.008	-0.005	-0.004	-0.008	0.006	-0.004
Max	0.001	0.000	0.004	0.007	0.005	0.014	0.014
σ	0.001	0.001	0.002	0.002	0.002	0.002	0.004
Range	0.008	0.008	0.009	0.011	0.013	0.008	0.018

Table C.11: L1 phase residuals between reference and antenna 4 during L008.

PRN	13	19	10	27	24	18	4
Elevation	70	63	52	50	45	42	24
Mean	0.002	-0.004	0.002	0.009	0.002	-0.005	-0.021
Min	-0.003	-0.011	-0.006	0.003	-0.004	-0.010	-0.038
Max	0.006	0.002	0.006	0.015	0.010	0.000	-0.011
σ	0.002	0.003	0.002	0.002	0.002	0.002	0.004
Range	0.009	0.013	0.012	0.012	0.014	0.010	0.027

Table C.12: L2 phase residuals between reference and antenna 4 during L008.

Appendix D PRAYER Software Description

The information in this section has been extracted from a ‘user manual’ which was written in conjunction with the PRAYER code when it was adapted at University College London for the use of the Natural Environment Research Council. As it was written for general use by the NERC ARSF, the only GPS attitude solution referred to is from the ADU2, dedicated GPS attitude determination unit. The independent receiver configuration was used specifically on the two test flights for this project, and is not generally available.

Introduction

This document describes how to install and run the ‘PRAYER’ software written at the Department of Geomatic Engineering. The program reads attitude information from an ADU2 file and an AHRS file and produces an optimal estimate of the attitude of each scan line (read from a scanner file) using a Kalman filter. The following document describes the structure and linking of the source code files and libraries, and how input and output files should be set up. It also provides information on the various processing parameters that the user may wish to change, and suggests appropriate default values that should give a usable solution for any standard data-set.

Contents

D.1	Software Installation	292
D.2	Input Data	292
D.3	Running the Software	293
D.4	Initialisation File Contents	294
D.5	Output Files	297
D.6	Changes to Source Code and Re-compiling	298
D.7	Weighting	299
Appendices	D.A Input File Examples	301
	D.B Output File Examples	303

D.1 Software Installation

This software has been written in the C programming language and currently compiles using 'cc' or 'c89' compilers. To date it has been tested on Silicon Graphics and SUN UNIX workstations in the Department of Geomatic Engineering, UCL. The files are stored in the following sub-directories which are all within 'Attitude_Filter' directory.

bin	contains the executable 'PRAYER' and any initialisation files,
data	contains the AHRS, ADU2 and scanner files in the format described in section D.2,
include	contains library files in which functions, data structures and global constant definitions are given,
obj	contains the object files created when the source code is compiled,
source	contains the program functions in a series of separate files and a <i>makefile</i> which links all source and library files for compiling,
output	contains the standard output file and any additional files of processing information that the user requests,

D.2 Input Data

To run the program three space delimited text files are required, containing the following information:

AHRS	time : pitch : roll : heading : pitch-rate : roll-rate : heading-rate
ADU2	time : pitch : roll : heading
ATM/CASI	scan-number : scan-time

At present the input data files also contain additional values that are not required for the program and are therefore not retained once each line is read. Examples of the current format are given in Appendix D.A. If the format of the input data files changes in the future the reading statements can be changed in the source code relatively easily (files *rd_ahrs.c*, *rd_gps.c* and *rd_scanner.c*).

D.3 Running the Software

Before running the program, an initialisation file must be created in the **bin** directory. The three input data files (AHRS, ADU2 and Scanner files) must be put in the directory stated in this initialisation file and an output directory needs to be created. Output files are created within the program and given the names specified in the initialisation file. The initialisation file also contains a number of processing parameters that allow the user to assign weights to measurements and predictions, to define outlier rejection thresholds, and to specify what processing information should be produced when the program is run. More details on the content of the initialisation file are given in section D.4.

The program is run by changing to the **bin** directory and typing the name of the executable followed by the name of an initialisation file.

eg. prompt% PRAYER PRAYER.ini

A progress indicator is printed to the screen if the initialisation is found and opened successfully.

reading initialisation file ...

If there are any problems opening the files specified in the initialisation file, the program will stop running and an appropriate error message will be displayed. If an error message is displayed it will probably instruct the user to check the name and path of one or more of the input and output files. If the initialisation file is read, and no problems are found, the following progress indicator will be printed to the screen,

initialisation file read

The program will now run with no additional prints on the screen, and once it is complete the prompt will reappear. All the output files specified will be in the directory stated in the initialisation file.

D.4 Initialisation File Contents

An intialisation file must be created and specified each time the program is run. It contains information regarding the input and output files and allows the user to specify a range of processing parameters. It is divided into five sections, as follows,

- 1) Input and Output Data
- 2) Standard Deviations of Measurements
- 3) Standard Deviations of Motion Predictions
- 4) Processing Information
- 5) Rejection Thresholds

1) Input and Output Data

Input and Output Data

AHRS filename (& path)	../data/AHRS20402.vec
ADU2 filename (& path)	../data/ADU20402.vec
Scanner filename (& path)	../data/SC20402.vec
ATM or CASI ? (type A or C)	a
ADU2 Scaling factor for CASI	1
Output filename (attitude per scanline)	../output/20402.out
AHRS residuals output filename	../output/20402.ahr
ADU2 residuals output filename	../output/20402.adu
Summary output filename	../output/20402.sum
Smoothed output filename	../output/20402.smt

These are the names and paths of all input and output files. The paths do not have to be set up as indicated here, but depending on the operating system, long paths should generally be avoided. The content and structure of input and output files are outlined in sections D.2 and D.5 respectively. The options to choose the type of scanner and to define a scaling factor have not been implemented in the current program.

2) Standard Deviations of Measurements

Standard Deviations of Measurements

AHRS Pitch (degrees)	0.5
AHRS Roll (degrees)	0.5
AHRS Heading (degrees)	0.5
AHRS Rate of Pitch (degrees/second)	0.3
AHRS Rate of Roll (degrees/second)	0.3
AHRS Rate of Heading (degrees/second)	0.3
ADU2 Pitch (degrees)	0.5
ADU2 Roll (degrees)	0.5
ADU2 Heading (degrees)	0.3

Each measurement is assigned a standard deviation that should reflect its accuracy as closely as possible. If possible, good approximations of each instrument's absolute accuracy should be used. The key requirement however, is to define values which adequately reflect the *relative* accuracies of the ADU2 and AHRS, and also the extent to which the final solution should accommodate individual measurements or predicted motion. This point is discussed further in section D.7.

3) Standard Deviations of Motion Predictions

Standard Deviations of Motion Predictions

Airframe angular acceleration	
Pitch (degrees/second_squared)	1
Roll (degrees/second_squared)	1
Heading (degrees/second_squared)	1
Rate of change of airframe-ATM offset	
Pitch offset (degrees/second_squared)	1
Roll offset (degrees/second_squared)	1
Heading offset (degrees/second_squared)	1

These values dictate the weight given to the predicted motion of the system. Standard deviations of the airframe angular acceleration effectively assign confidence to the assumption of constant angular rates. A low standard deviation equates to a high weighting, and the model will therefore hold the angular rates relatively constant. If a larger standard deviation is used the weighting is reduced and the final

solution will be more sensitive to any changes in the angular rate indicated by the measurements. The standard deviations of the rate of change of offset act in a similar way. A low standard deviation will cause the model to give weight to the assumption that the offset is constant, making the final solution less sensitive to any changes indicated by the measurements. The critical issue of weighting is discussed in greater detail in section D.7.

4) Processing Information

```

Processing Information
-----
(type Y or N)
Summary (.sum)                Y
Residuals at measurement epochs (.res) y
Apply smoothing                n

```

These options allow the user to choose whether a processing summary file is produced, whether files of predicted residuals are created, and whether or not smoothing is applied. In tests to date smoothing has had a negligible effect on the final solution and should not be used by default.

5) Rejection Thresholds

```

Rejection Thresholds
-----

AHRS Pitch Threshold          1
AHRS Roll Threshold           1
AHRS Heading Threshold        1

AHRS Pitch-rate Threshold     3
AHRS Roll-rate Threshold      3
AHRS Heading-rate Threshold   3

ADU2 Pitch Threshold          1
ADU2 Roll Threshold           1
ADU2 Heading Threshold        1

```

Rejection thresholds allow measurements that do not correspond closely enough to the predicted values at an epoch to be rejected. Low thresholds may remove a greater proportion of outlying data but may lead to so many measurements being rejected that the solution no longer reflects the original measurement trends sufficiently. Information on any rejection that has taken place during processing is contained in the summary and residual files. If any measurement at an epoch exceeds the threshold value then all the measurements will be rejected as they are considered as similarly outlying.

D.5 Output Files

The outputs of primary interest are the attitude values for each scanline. These are contained in the file named as 'output file' in the initialisation file. This file is a space delimited text file containing five pieces of information for each scanline, i.e.

scanline number	scanline time	roll	pitch	heading
-----------------	---------------	------	-------	---------

In addition to this standard output file other processing information can be stored and output to a file. The briefest of these is a processing summary file which states the number of epochs of ADU2 and AHRS used in the filter, and the number of these epochs that were rejected after being identified as outliers.

Files listing the predicted residuals at each measurement epoch can also be created. The predicted residual is the difference between a predicted value at an epoch and the measured value of the same quantity. The magnitude and pattern of these residuals gives an indication of how well the original measurements agree with the predicted motion of the system. Processing parameters should be adjusted until the best possible residuals are produced. Two files of residuals are produced, one for ADU2 measurements and one for AHRS.

If smoothing is used when processing, a file similar to the main output file will be produced but the values will be the smoothed rather than filtered solutions per scanline. This file could be used in preference to the standard output, or could be used simply to check whether smoothing is significantly improving the final solution.

Examples of each of the current formats are given in Appendix D.B.

D.6 Changes to Source Code and Re-compiling

The present program should work for all flight lines and produce attitude values for each scanline. The quality of the final filtered solution will depend on the quality of the original data and the appropriateness of the user-specified weightings for measurements and motion predictions.

If, for any reason, changes need to be made to the source code (e.g. if the format of the input or output files needs to be altered) this should be done with care. It is a wise precaution to make copies of any original files before making edits. Once edits have been made the files should be saved and if any re-naming has been necessary the *makefile* must be updated accordingly. Similarly, some edits to the source code may require updates to some of the files in the *include* directory (probably PRAYER.h, if any). To recompile the code, change to the *source* directory and type:

Prompt% make clean

Followed by

Prompt% make

This will use the *makefile* to link all the necessary files and to compile the code. The executable will be created in the *bin* directory, unless the path has been changed in the *makefile*. The current *makefile* specifies the compiler 'cc' although this can be changed to 'c89'. To compile this code the appropriate function libraries must be in place.

D.7 Weighting

Appropriate weighting must be applied in a Kalman filter to produce an optimal solution. The filter produces a best estimate of the states based on the new measurements available at an epoch and the predicted states, based on the previously estimated states and assumptions describing how the system changes with time.

If, for example, the measurements from a particular instrument are known to be of high quality (accurate and trustworthy) then greater weights can be assigned to these values using appropriately small standard deviations. The filtered solution will then follow the pattern of the original measurements closely. If on the other hand, measurements of a vehicle's behaviour are known to be noisy, but it is known that in fact the vehicle is moving in a regular and predictable manner then greater weight should be attached to the predictions.

Ideally a reference solution that accurately describes the system's motion should be available for calibrating the system. The relative weightings of the measurements and predictions can then be adjusted to the appropriate levels and these can be applied in cases when no reference solution is available. With the currently available data it is not possible to determine with any certainty what the true attitudes of the scanners were as each line was imaged. The best estimate must therefore be based on assumptions about the system's behaviour and the quality of the measurements from the ADU2 and AHRS instruments. The weightings then describe the confidence that can be given to these assumptions.

In this program the estimated states are:

- i) the airframe attitude (pitch, roll and heading),
- ii) the angular rates of the airframe (pitch-rate, roll-rate, heading-rate),
and,
- iii) the axis offset between the AHRS and ADU2 reference frames (pitch, roll, heading).

The angular acceleration of the airframe and the rate of change of offset are considered to be of a small and random nature, and can therefore not be successfully modelled as additional states. With insufficient information to model any angular acceleration or change of offset, these values are set to zero, as they are as likely to be positive as negative. The standard deviations that the user sets for these values in the initialisation file effectively indicate how good these assumptions of constant angular velocity and constant offset are believed to be. If a small standard deviation is assigned to the angular acceleration estimate of zero, then the filtered solution will show a relatively constant angular velocity even if the measurements suggest that this should not be the case. Overall, the greater the weighting assigned to the predictions relative to the measurements, the smoother the filtered solution will be.

The user can assign standard deviations, and hence weights, to all of the system measurements in this application. As described above, the weights assigned to these measurements should reflect their quality relative to the quality of the model predictions. The weights should also reflect the differences in the accuracy of the ADU2 and AHRS instruments, and describe the relative quality of different attitude components. For example, the ADU2 accuracy is a function of the antenna separation, therefore the roll component which is based on a 1.2m baseline is approximately three times less accurate than the pitch component which is based on a 3.5m baseline. The heading component is more accurate than both roll and pitch because it does not rely on the height component of GPS positioning.

Appendix D.A Input File Examples

For each flight line there are three input data files; AHRS, ADU2 and scanner. These are extracted from hdf files as space delimited text files.

The ADU2 file contains the following data:

file line no.	time (s)	roll (°)	pitch (°)	heading (°)
0	44209.0	-0.384	2.574	353.818
1	44209.5	-0.699	2.438	353.821
2	44210.0	-0.602	2.628	353.979
3	44210.5	-0.305	2.471	353.850
4	44211.0	-0.534	2.392	353.495
5	44211.5	-0.643	2.340	353.175
6	44212.0	-0.431	2.430	353.201
7	44212.5	-1.105	2.475	353.329
8	44213.0	-0.893	2.393	353.294
9	44213.5	-0.435	2.172	353.226
10	44214.0	-0.444	2.345	353.392

The AHRS data file contains the following information:

file line no.	time	Attitude (°)			Angular Rates (°/s)		
		roll	pitch	heading	roll	pitch	heading
0	44213.018	-0.359	2.129	344.004	0.164	0.059	-0.119
1	44213.033	-0.349	2.129	344.004	0.509	0.119	-0.119
2	44213.049	-0.359	2.129	343.999	0.209	0.044	-0.119
3	44213.065	-0.349	2.129	343.993	0.119	0.029	-0.164
4	44213.080	-0.340	2.129	343.993	0.509	-0.059	-0.149
5	44213.096	-0.340	2.129	343.988	0.419	-0.089	-0.179
6	44213.112	-0.340	2.129	343.988	0.059	0.014	-0.209
7	44213.127	-0.319	2.129	343.982	0.329	-0.074	-0.164
8	44213.143	-0.329	2.129	343.977	0.435	-0.074	-0.179
9	44213.159	-0.340	2.129	343.977	-0.029	-0.029	-0.194

The scanner data file contains only,

file	scanline	
line no.	number	scan-time
0	7128	44214.503
1	7129	44214.523
2	7130	44214.543
3	7131	44214.563
4	7132	44214.583
5	7133	44214.603
6	7134	44214.623
7	7135	44214.643
8	7136	44214.663

To produce a good solution from the filtering program the times from each of the three files must be synchronised before processing. The accuracy of this time synchronisation will have a significant impact on the accuracy of the final solution. The filtering program cannot automatically determine and apply a time offset.

Appendix D.B Output File Examples

Up to five output files can be created each time the program is run, these are:

- i) a standard output file containing the filtered attitude values for each scanline,
- ii) an output file containing smoothed attitude values for each scanline,
- iii) predicted residuals for ADU2 measurement epochs,
- iv) predicted residuals for AHRS measurement epochs,
- v) a summary file indicating the measurement data used, and the amount of data rejected.

- i) The standard output file contains:

scanline number	scanline time	filtered roll (°)	filtered pitch (°)	filtered heading (°)
7128	44214.503	-0.167	2.329	353.267
7129	44214.523	-0.160	2.328	353.263
7130	44214.543	-0.152	2.327	353.259
7131	44214.563	-0.144	2.326	353.255
7132	44214.583	-0.138	2.324	353.248
7133	44214.603	-0.130	2.322	353.243
7134	44214.623	-0.124	2.321	353.238
7135	44214.643	-0.117	2.319	353.233
7136	44214.663	-0.112	2.316	353.224

- ii) The smoothed output file has the same structure as the standard filtered output file.

- iii) The predicted residuals file for ADU2 measurements contains:

Predicted Residual Attitude (°)

time	pitch	roll	heading
44357.5	0.173	-0.130	0.032
44358.0	-0.147	-0.401	-0.204
44358.5	0.088	-0.814	-0.307
44359.0	0.223	-0.498	-0.252
44359.5	0.132	-0.692	-0.157
44360.0	0.123	0.651	-0.151
44360.5	-0.191	0.451	-0.309
44361.0	-0.214	-0.059	-0.334

iv) The predicted residuals file for AHRS measurements contains:

time	Predicted Residual					
	Attitude (°)			Angular Rates (°/s)		
	pitch	roll	heading	pitch	roll	heading
44357.501	-0.107	0.103	0.005	-0.229	0.260	-0.385
44357.517	-0.117	0.089	0.000	-0.124	-0.159	-0.385
44357.533	-0.117	0.086	-0.005	-0.168	-0.264	-0.369
44357.548	-0.117	0.102	-0.010	-0.168	0.110	-0.385
44357.564	-0.127	0.089	-0.016	-0.243	-0.010	-0.368
44357.579	-0.117	0.075	-0.020	-0.123	-0.414	-0.398
44357.595	-0.127	0.093	-0.026	-0.123	-0.084	-0.398
44357.611	-0.127	0.078	-0.037	-0.212	0.049	-0.398
44357.626	-0.127	0.064	-0.041	-0.122	-0.474	-0.444
44357.642	-0.137	0.082	-0.046	-0.212	-0.294	-0.427

v) The summary file contains a very brief account of the measurements used in the filter.

10971 epochs of AHRS measurements were used in the filter of which 0 were rejected as outliers

343 epochs of ADU2 measurements were used in the filter of which 1 were rejected as outliers