UCGE Reports
Number 20145

Department of Geomatics Engineering

Improving the Accuracy and Resolution of SINS/DGPS Airborne Gravimetry

(URL: http://www.geomatics.ucalgary.ca/GradTheses.html)

by

Alexander Mark Bruton

December 2000

UNIVERSITY OF CALGARY

Improving the Accuracy and Resolution of SINS/DGPS Airborne Gravimetry

by

Alexander Mark Bruton

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE

DEGREE OF DOCTOR OF PHILOSOPHY

DEPARTMENT OF GEOMATICS ENGINEERING

CALGARY, ALBERTA

December, 2000

© Alexander Mark Bruton 2000

Abstract

This dissertation describes improvements made to a system for airborne mapping of the gravity field of the Earth. The research is carried out using an airborne gravity system that is based on a Strapdown Inertial Navigation System (SINS) and receivers of the Global Positioning System in differential mode (DGPS). The objective of the research is to optimize the performance of the system, especially for geodesy and geophysics.

An introduction to the field of airborne gravimetry is given and the state of current research in the field is surveyed. Data from recent airborne gravity campaigns is used to provide a detailed analysis of the DGPS error budget for airborne positioning, providing a realistic evaluation of the accuracy of current kinematic carrier phase techniques. A fundamental consideration of the various processes of differentiation is given and particular differentiating filters are proposed for the determination of high precision velocity and acceleration. A detailed analysis is given in the frequency domain of the DGPS error budget for acceleration determination. This provides an understanding of the characteristics of each of the relevant error sources for spatial resolutions up to 500 m and forms the basis for a set of recommendations regarding acceleration determination for airborne gravimetry. The limitations of the SINS gravimeter that are imposed by the accelerometer biases are analyzed and quantified. A thorough analysis is provided of the dynamics experienced by survey aircraft. The high-frequency errors affecting airborne gravimetry are analyzed in detail and methods for reducing them are proposed and implemented with success.

An improvement to the performance of the system for medium-resolution applications is achieved and it is demonstrated for the first time that the SINS/DGPS system can be used for high-resolution applications. Major results include a demonstrated accuracy of 1.5 mGal for a spatial resolution of 2.0 km and an accuracy of 2.5 mGal for a resolution of 1.4 km. Improvements to processing methods have yielded slightly better performance than the LaCoste and Romberg gravimeter on a common flight. A method for removing the effect of the Phugoid motion has been proposed and implemented with success.

Acknowledgements

Graduate studies has been a wonderful experience for me. It has allowed me to learn, to travel to faraway places like Greenland and Brazil, and perhaps most importantly, to make some lasting friendships. I would like to express appreciation for the support I have received over the last four years:

- To Dr. Klaus-Peter Schwarz, thank you for the opportunity to work within such a world-class research group during my graduate studies. Your support, encouragement, trust and technical leadership have made it a truly enjoyable and invaluable experience.
- To those organizations whose financial support has made my graduate studies possible, especially the Natural Sciences and Research Council (NSERC) of Canada, whose support has been fundamental. Thanks to the Government of Alberta for support in the form of an Alberta Government Graduate Scholarship, to the University of Calgary for support in the form of the Helmut Moritz Graduate Scholarship and to the Department of Geomatics Engineering for financial support on several occasions. Also, thanks go to those organizations that improved my graduate experience in the form of awards and travel grants.
- To past and present members of our research group at the University of Calgary. To Dr. Ming Wei, thank you for your generosity and for the practical work that has created the framework upon which many students at the University of Calgary have built, myself included. Although there are many others, I would like especially to thank Jan Skaloud, Craig Glennie, Yecai Li, Mohamed Mostafa, Michael Kern, Kris Keller, Fadi Bayoud and Sandy Kennedy for their help and partnership on many occasions. Craig is also thanked for thoroughly testing the first version of GREATGUN in his thesis work.
- To my unofficial examining committee: Michael Kern, Jason Innes and especially Len Bruton. The quality of this dissertation was greatly improved as a result of

the discussions we had and as a result of your thoughtful criticism of the first draft. Your encouragement and the time you spent to help me is truly appreciated.

- Thanks goes to the Geomatics for Informed Decisions (GEOIDE) Network Centre of Excellence for financing my position at the University of Calgary for the last year and a half and for the great environment for collaboration that they helped create. Also, other members of the *Airborne Gravity for Exploration and Mapping* Project are thanked for enriching my learning experience.
- To Rene Forsberg of the Danish National Survey and Cadastre (KMS), for the invitation to participate in the workshop and field test in Greenland, in June 1998. Greenland was an unforgettable experience!
- To Hugh Martell and Darren Cosandier of Waypoint Consulting Ltd., for the loan of GrafNav 6.02. It made data processing a pleasure.
- To Keith Tennant, for your feedback and the useful discussions we had.
- To Intermap Technologies Corporation for the loan of the Honeywell Laseref III on many occasions and for the flight services that were provided.
- To Steve Ferguson for your hospitality and for always being willing to help.
- To Sander Geophysics Ltd. for the flight services that were provided during the two testing periods that I joined them in Ottawa.
- To Vladimir Argeseanu for providing the upward continued reference gravity field for the Kananaskis campaign.
- To Paul Mrstik for his generosity and help.
- To the members of my examining committee for their advice: J.M. Brozena, M.G. Sideris, M.E. Cannon and J.W. Haslett.
- To the staff in the main office of Geomatics Engineering, especially Anne Gehring. Your help is appreciated.

On a more personal note, I would like to express appreciation to my family and friends who have been so supportive over the last four years, especially:

- To Mum, Dad, Michelle, Nicole and Adrian, thank you for your friendship and your support.
- Finally, thanks to my best friend Krista for sharing and dreaming with me.

Dedication

This work is dedicated to my parents who teach me every day by example.

To Mum, whose arms are always open, making ours the luckiest family in the world, and to Dad, whose love for us is evident in everything that he does.

Thank you for your support, your strength, your patience, your friendship and your love.

Table of Contents

List of Figures

Notation, Symbols and Acronyms

Notation

In this section, mathematical notation and coordinate frames used in the dissertation are introduced. They are the same as those used in Schwarz and Wei (1997).

Vectors are represented by letters that are lowercase and boldface. The most common uses of vectors herein are for the representation of position, velocity and angular velocity.

The following orthogonal Cartesian **coordinate frames** are used in this dissertation:

a) The Operational Inertial Frame (i-frame) is a reference frame in which Newton's equations of motion apply. It is therefore not rotating or accelerating. For practical applications, it is defined as follows:

   origin: at the centre of mass of the Earth
   z-axis: parallel to the spin axis of the Earth
   x-axis: pointing towards the mean vernal equinox
   y-axis: completing a right-handed frame

b) The Conventional Terrestrial Frame (e-frame) is fixed to the Earth. It is defined as follows:

   origin: at the centre of mass of the Earth
   z-axis: parallel to the spin axis of the Earth
   x-axis: pointing towards the mean meridian of Greenwich
   y-axis: completing a right-handed frame

   The e-frame rotates with respect to the i-frame by a nearly constant angular rotation about the z-axis.

c) The Local-level (l-frame) is defined with respect to an Earth-fixed ellipsoid of revolution such that its axes point east, north and up, as follows:

   origin: at the point of interest (e.g. the origin of the sensor frame)
   z-axis: along the normal of the reference ellipsoid, pointing outward
   x-axis: completing a right-handed frame
   y-axis: pointing towards geodetic north

d) The Body Frame (b-frame) is fixed to the sensor platform. In principle, this coincides with the body of the host vehicle (or can be related to it by a translation

and rotation). Because it depends on the motion of the vehicle, it can have any orientation with respect to the Earth. It is defined as follows:

origin: at the origin of the sensor platform
z-axis: pointing upward
x-axis: completing a right-handed frame
y-axis: point in the direction of travel

A superscript is used to indicate the coordinate frame in which the components of a vector are given. For example, the position of an object with respect to the b-frame is given by:

$$\mathbf{r}^b = \begin{pmatrix} x^b \\ y^b \\ z^b \end{pmatrix}.$$

An angular velocity vector is defined using a superscript and two subscripts. The superscript is used, as for any vector, to represent the coordinate frame in which the information is expressed. The rotation is from the frame given by the second subscript to the frame given by the first subscript. For example, the angular velocity between the i-frame and the b-frame is expressed in the b-frame by:

$$\boldsymbol{\omega}_{ib}^b = \begin{pmatrix} \omega_x \\ \omega_y \\ \omega_z \end{pmatrix}.$$

Matrices are represented by letters that are uppercase and boldface. The most common use of matrices herein is to represent the rotation from one coordinate frame to another. In this case, the original frame is represented by a superscript and the new frame is represented by a subscript. For example, the representation of a coordinate in the l-frame, $\mathbf{r}^l$, can be computed from its representation in the b-frame, $\mathbf{r}^b$, as follows:

$$\mathbf{r}^l = \mathbf{R}_b^l \mathbf{r}^b .$$

It is useful to define the rotation matrix, $\mathbf{R}_b^l$, because it defines the attitude of a vehicle, i.e. the orientation of the vehicle with respect to the l-frame. It is given by:

$$\mathbf{R}_{\mathrm{b}}^{\mathrm{l}} = \begin{pmatrix} \cos\psi\cos\varphi - \sin\psi\sin\theta\sin\varphi & -\sin\psi\cos\theta & \cos\psi\sin\varphi + \sin\psi\sin\theta\cos\varphi \\ \sin\psi\cos\varphi + \cos\psi\sin\theta\sin\varphi & \cos\psi\cos\theta & \sin\psi\sin\varphi - \cos\psi\sin\theta\cos\varphi \\ -\cos\theta\sin\varphi & \sin\theta & \cos\theta\cos\varphi \end{pmatrix},$$

where φ, θ and ψ are three Euler angles known as the roll, pitch and yaw of the host vehicle. They are positive counter clockwise rotations about the y, x and z-axes of the b-frame, respectively.

The angular velocity is also commonly expressed in the so-called **skew symmetric form** using a matrix, as in the following example:

$$\Omega_{\mathrm{ib}}^{\mathrm{b}} = \begin{pmatrix} 0 & -\omega_{\mathrm{z}} & \omega_{\mathrm{y}} \\ \omega_{\mathrm{z}} & 0 & -\omega_{\mathrm{x}} \\ -\omega_{\mathrm{y}} & \omega_{\mathrm{x}} & 0 \end{pmatrix}.$$

Finally, a dot is used above a vector to indicate the time derivative of that vector. For example, $\dot{\mathbf{v}}^{\mathrm{l}}$ is the time derivative of $\mathbf{v}^{\mathrm{l}}$.

List of acronyms

The following acronyms are used in the dissertation.

Acronym	Meaning
AGEM	Airborne Gravity for Exploration and Mapping, GEOIDE
b-frame	Body Frame
DGPS	Global Positioning System in differential mode
DOP	Dilution of Precision (in reference to satellite geometry)
e-frame	Conventional Terrestrial (or Earth Fixed) Frame
EGM96	Earth Gravitational Model of 1996
FIR	Finite Impulse Response (for a discrete-time filter)
GEOIDE	Geomatics for Informed Decisions, Network Centre of Excellence
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
GSD	Geodetic Survey Division of Geomatics Canada
GREATGUN	Gravity Estimates from Airborne Techniques for Geoid Undulations

i-frame	Operational Inertial Frame
IFSAR	Interferometric Synthetic Aperture Radar
IIR	Infinite Impulse Response (for a discrete-time filter)
INS	Inertial Navigation System
Intermap	Intermap Technologies Corporation
ITC	Inertial Technology Center in Moscow
KMS	Kort & Matrikelstyrelsen (Danish National Survey and Cadastre)
LRF-III	Honeywell Laseref III inertial navigation unit
l-frame	Local-level Frame
mGal	milliGal = one thousandth of a Gal = 1×10^{-5} m/s^2
MUF	Maximum Undistorted Frequency
NCE	Network Centres of Excellence
NRL	United States Naval Research Laboratory
PDOP	Position Dilution of Precision (in reference to satellite geometry)
r-frame	Reference Frame (for aircraft motion, see Chapter 7)
RISG	Rotation Invariant Scalar Gravimetry
RMS	Root Mean Square
SAS	Stability Augmentation System (aircraft autopilot)
SI	Système Internationale (units)
SINS	Strapdown Inertial Navigation System
SISG	Strapdown Inertial Scalar Gravimetry
ZLS	Zero-length spring gravimeter (produced or modified)

List of symbols

The following symbols are used in the dissertation.

Symbol	**Meaning**
δa	deflection caused by the aileron (on an aircraft)
$\mathbf{a}$	unknown integer ambiguities
A	attenuation factor
$\mathbf{A}$	ambiguity design matrix, GPS linear equation, Section 2.3.3 only

β	inverse of the correlation period
$\mathbf{b}$	accelerometer bias
$\mathbf{B}$	baseline design matrix, GPS linear equation
d	error, with reference to a vector, e.g. d$\mathbf{v}$
$\mathbf{d}$	gyroscope biases
δe	deflection caused by the elevator (on an aircraft)
ε_{Φ}	measurement delay due to carrier phase noise
ε	misalignment
$\mathbf{e}$	measurement noise
f_c	cut-off frequency
f_s	sampling frequency
$\mathbf{f}$	specific force
$\mathbf{f}_p$	propulsive forces
$\mathbf{f}_a$	aerodynamic forces
$\mathbf{F}$	skew symmetric form of the specific force
$\mathbf{F}_r$	restoring force (of an accelerometer)
$\mathbf{F}_{lon}$	characteristic equation of the natural longitudinal motion of an aircraft
$\mathbf{F}_{lat}$	characteristic equation of the natural lateral motion of an aircraft
g	magnitude of gravity
δg	an element of the gravity disturbance vector
δg_{GM}	gravity disturbance from a global geopotential model
δg_r	relative gravity disturbance
$\delta\mathbf{g}$	gravity disturbance
$\mathbf{g}$	gravity
$\mathbf{G}_{lon}$	longitudinal control matrix (aircraft autopilot)
$\mathbf{G}_{lat}$	lateral control matrix (aircraft autopilot)

h	ellipsoidal height
$\mathbf{h}$	impulse response
H	frequency response
$H(\psi)$	Hotine kernal function
$\mathbf{H}$	angular momentum, Section 7.1 only
$\mathbf{H}$	measurement design matrix, Section 7.4 only
I	measurement delay due to the ionosphere
$\mathbf{I}$	instantaneous inertia tensor
L1, L2	GPS carrier signals
$\mathbf{L}$	torque
m	mass element (of an accelerometer)
m_{Φ}	measurement delay due to carrier phase multipath
$\mathbf{n}_f$	accelerometer noise
$\mathbf{n}_\omega$	gyroscope noise
N	ambiguous integer carrier phase cycle
N	geoid undulation
N_{GM}	geoid undulation from a global geopotential model
N_r	relative geoid undulation
$\mathbf{N}_f$	accelerometer non-orthogonalities
$\mathbf{N}_\omega$	gyroscope non-orthogonalities
ρ	true range between the satellite and the antenna
$\delta\rho$	ephemeris error
$\mathbf{P}$	covariance of the state vector
$\mathbf{Q}$	process noise matrix
δr	deflection caused by the rudder (on an aircraft)
$\mathbf{r}$	position
$\ddot{\mathbf{r}}$	inertial acceleration

$\mathbf{R}$	rotation matrix
$\mathbf{R}$	measurement noise covariance matrix, Section 7.4 only
s	speed
$\mathbf{S}_f$	accelerometer scale factor
$\mathbf{S}_\omega$	gyroscope scale factor
δt	deflection caused by the throttle (on an aircraft)
T	period (of time)
T	measurement delay due to the neutral atmosphere, Chapters 2, 3 and 5
$\mathbf{v} = \dot{\mathbf{r}}$	velocity
$\mathbf{v}$	eigenvector, Section 7.1 and Appendix D only
ω	frequency (radians)
ω_s	sampling frequency (radians)
ω_c	cut-off frequency (radians)
$\boldsymbol{\omega}$	angular velocity
$\mathbf{x}$	unknown misalignments
$\mathbf{z}$	measurement vector, Chapters 2 and 7
λ	eigenvalue, Section 7.1 and Appendix D only
λ	wavelength of the carrier, Chapters 2 and 3
σ	sigma, with reference to standard deviation, e.g. 1-σ
γ	normal gravity
φ	roll
θ	pitch
ψ	yaw
$\nabla\Delta$	double difference, between satellites and receivers
Φ	carrier phase measurement
Φ	transition matrix
Ω	skew symmetric form of the angular velocity vector

Introduction

This dissertation describes contributions and improvements made by the author to a system for airborne mapping of the gravity field of the Earth.

The research has been carried out using a system that is based on the combination of a Strapdown Inertial Navigation System (SINS) and receivers of the Global Positioning System (GPS) being used in differential mode (DGPS). Under the direction and supervision of Dr. K.P. Schwarz at the University of Calgary, the development of the SINS/DGPS system as a gravimeter has gone hand in hand with its development as a system for navigation and for image mapping.

From concept to implementation, the use of a SINS for airborne gravimetry was pioneered in the early 1990s by researchers at the University of Calgary. The first SINS/DGPS test results, that became available in 1995, demonstrated the feasibility of the system as an accurate and relatively small and low cost alternative to all other available airborne gravity systems. Over the last five years, the SINS/DGPS system has undergone exciting phases of development and testing that have led to its acceptance as an airborne mapping system for use in a wide variety of applications.

This dissertation describes some of the contributions that the author has made to the development and testing of that system since May 1996. The immediate goal of the reported research is to optimize the performance of the SINS/DGPS system, especially for applications in geodesy and geophysics. For these applications, the ultimate objective is to use airborne methods to measure the gravity field of the Earth with an accuracy of better than 1 mGal (which is about one part per million of the magnitude of the Earth's gravity) and a spatial resolution of 1 km.

This goal is shared by the Airborne Gravity for Exploration and Mapping (AGEM) project of the Canadian Network Centres of Excellence (NCE) known as Geomatics for Informed Decisions (GEOIDE). This NCE project was established in May 1999 and began with the objective of bringing together Canada's best expertise in airborne

gravimetry. It includes members from Canadian industry, government and academia. This project is relevant to this dissertation because much of the research presented herein contributes to the AGEM project and also because the work in Chapters 3 and 5 was carried out by the author as a part of the project.

Real SINS and DGPS data from a number of airborne gravity field campaigns is used throughout the research. These are described in Appendix A, where it is shown in Figure A-1, that they took place over the last four years in areas of the Canadian Rocky Mountains, Greenland and Southern Ontario and Quebec. They are the result of various national and international collaborative efforts and demonstrate the applicability of the University of Calgary system.

This dissertation is divided into three parts, each containing a number of chapters. The objectives in Part 1 are to introduce the field of airborne gravimetry and to set the stage for the contributions that follow. In meeting these objectives, it answers the basic question of why gravity data is needed and addresses the role that airborne gravimetry plays in its collection. The specific objectives of the research are also included in Part 1.

In Part 2, contributions are reported on the use of DGPS for positioning, velocity determination and motion compensation: three topics that are applicable to airborne gravity systems in general. The determination of position and velocity are also directly applicable to airborne navigation and mapping. The reported research includes a detailed study of the DGPS error budget, for both position and acceleration determination, and forms the basis for recommendations regarding the minimization of that error budget for use in an airborne gravity system.

In Part 3, contributions are described that specifically apply to airborne gravity systems that are based on the SINS concept. The objectives of this part of the research are to characterize the inertial sensor errors that currently limit the bandwidth of such a system and to propose ways to minimize them. These errors are twofold. First are the biases that influence the low-frequency performance of the system and second are the high-frequency errors that are induced by the dynamics of the aircraft.

PART 1: BACKGROUND AND JUSTIFICATION

The objective of the two chapters in this first part is to introduce the field of airborne gravity mapping in order to set the stage for the contributions presented in the balance of the dissertation.

In Chapter 1, a summary of the field is provided that includes a description of the needs for airborne gravity data as well as the accuracy and resolution requirements within each of the major application areas. It states the problems addressed herein and describes related research being carried out by others. The objectives of the research in this dissertation are described.

In Chapter 2, the principle of SINS airborne gravimetry is formalized in terms of the basic equations used in subsequent chapters. The observables and the systems that are used to measure them are described and the equations relating the errors in the latter to those in the former are presented. Finally, the contributions in this dissertation are put into context using these error equations.

No contributions to the field of airborne gravimetry are included in Part 1, other than the survey of the state of the art that is presented in Chapter 1 and a personal spin on the way material is presented in Chapter 2.

1 Background and Research Objectives

There has been considerable growth in the field of airborne mapping using SINSs over the last decade. Although they have been available for use in navigation since the mid 1970s, it was not until the GPS became fully operational in the early 1990s that many of the capabilities of SINSs could be exploited for mapping. The SINS/GPS sensor duo is now recognized as an important tool for accurate, rapid and cost efficient mapping from a moving platform such as an airplane.

In typical airborne mapping applications, the airplane may be host to a variety of sensors such as optical or digital cameras, multi-spectral scanners, interferometric synthetic aperture radar and scanning lasers. These sensors take advantage of the excellent vantage point afforded by an airborne platform to efficiently measure the surface of the Earth. The SINS and GPS sensors are used together to define the origin in space and time of the measurements made by the mapping sensors. This process is known as direct georeferencing and results in knowledge of the position and attitude (i.e. orientation) of the measured objects with respect to a known and useful coordinate frame. In principle, as long as initial conditions are available, the SINS provides the position and attitude of the survey aircraft (and therefore the mapping sensor) in this frame. Unfortunately, while excellent over short time periods, the accuracy of the relative position information provided by the SINS is dominated by long-term errors (biases) that severely limit its performance. This implies the need for the GPS that provides an accurate measure of the absolute position and velocity of the aircraft, especially over the long-term.

The overall accuracy of such a mapping system is a function of the accuracy of

a) the navigation parameters provided by the SINS/GPS sub-system,
b) the attitude provided by the SINS,
c) the measurements made by the mapping sensors within the frequency band of interest, and

d) the knowledge of the spatial and temporal relationships between the mapping sensors, the SINS and the GPS.

The research presented herein is based on such an airborne mapping system in which the goal is to make measurements of the gravity field.

1.1 Gravimetry, airborne gravimetry and the SINS gravimeter

In principle, gravity is measured using an accelerometer. The goal of this section is to define the basic operation of an ideal accelerometer and to describe some of the different practical forms it can take in systems designed to measure gravity, including the SINS gravimeter.

The basic components of an ideal accelerometer include a proof mass that is restricted to movement along a single sensitive axis, a restraining device (e.g. a spring), and a housing. The mass is supported by the restraining device and it displaces with respect to an equilibrium position when subject to acceleration. The output of the accelerometer is an electrical signal that is proportional to the displacement. Because the relationship is known between the displacement and the restoring force applied to the mass, the accelerometer provides a measure of the force required to counter the force due to accelerations acting on the mass. In Newtonian mechanics, this can be expressed using the time derivative of the law of conservation of linear momentum as follows:

$$\mathbf{F}_{\mathrm{r}} + m\mathbf{g} = m\ddot{\mathbf{r}}, \tag{1.1}$$

assuming that the only forces acting on the proof mass, m, are the restoring force, $\mathbf{F}_{\mathrm{r}}$, and the component along the sensitive axis of the forces due to its inertial acceleration, $\ddot{\mathbf{r}}$, and gravity, $\mathbf{g}$. This has units of force, e.g. kg m/s^2 in the units of the Système Internationale (SI). Because m is known, this can be expressed in terms of the output of the accelerometer as follows:

$$\mathbf{f} = \frac{\mathbf{F}_{\mathrm{r}}}{m} = \ddot{\mathbf{r}} - \mathbf{g}, \tag{1.2}$$

where $\mathbf{f}$ is the specific force that has units of acceleration, e.g. m/s^2 in the units of the SI. Specific force is therefore the observable of an accelerometer.

The simplest use of an accelerometer for measuring gravity arises when $\ddot{\mathbf{r}} = 0$, a condition that is roughly approximated when it is stationary on the surface of the Earth. If

it is also level (i.e. oriented such that the sensitive axis coincides with the direction of the gravity vector), then the observable is the magnitude of gravity. Most modern terrestrial gravimeters are based on this principle.

Despite the simplicity and relatively high accuracy of this approach, the concept of gravimetry from a moving platform such as an airplane is an attractive alternative because of the obvious potential for increased speed, greater range and significantly lower cost. However, the situation described above becomes considerably more complex when the accelerometer is moving. This is because the inertial acceleration, $\ddot{\mathbf{r}}$, is no longer zero and because the instrument is not easily kept level. In principle then, an airborne gravity system requires several components: an accelerometer for measuring the specific force, a system for leveling that accelerometer (or for computing its attitude) and a system that measures the inertial acceleration of the airplane. The gravity vector is estimated by subtracting the measurements made by the accelerometer from the measured inertial acceleration, a process known as motion compensation.

Although the concept of measuring the gravity field of the Earth using airborne techniques is not new, it has been revolutionized in recent years by the use of the GPS as an accurate and reliable system component and by the development of new gravity system concepts. For example, consider Figure 1.1 that summarizes the history of airborne gravity systems by showing their development over time since 1950 and how they have made use of enabling technologies for positioning and motion compensation. The accuracy of airborne gravity before the late 1980s was mainly limited by the inadequate positioning and motion compensation systems that were available. It is shown in Figure 1.1a that a number of such systems have been employed over time, many of which were of insufficient quality or limited geographical extent. Although these systems had varying levels of success prior to 1990, it is shown in Figure 1.1 that it was not until carrier phase DGPS became a reality that any airborne gravity system became fully operational (i.e. capable of large-scale surveys, anywhere in the world, at any time). As a result of the high accuracy and reliability afforded by the GPS, all airborne gravity systems now depend on carrier phase DGPS as their primary source of information for both positioning and motion compensation.

a) use of enabling technologies in airborne gravity

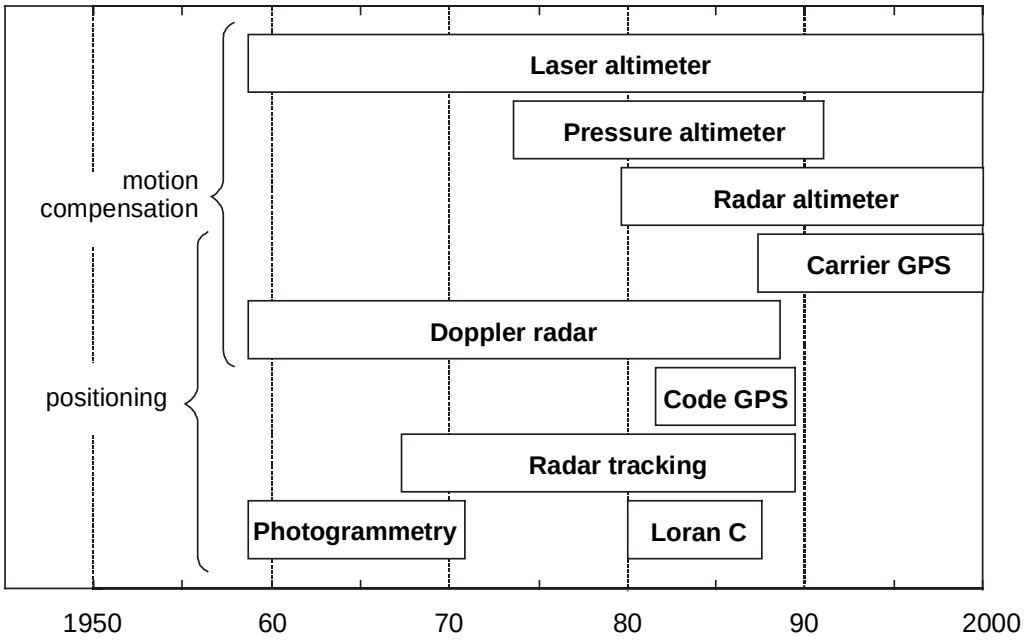

b) airborne gravity systems (solid lines indicate periods where fully operational)

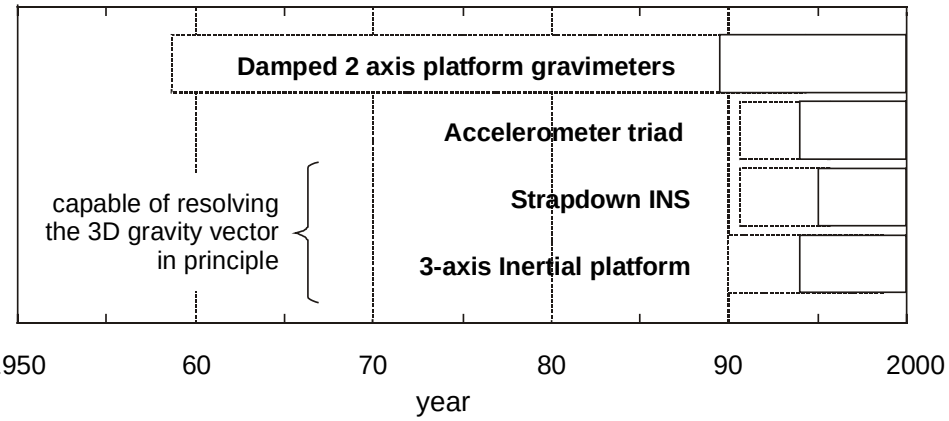

Figure 1.1 A brief history of airborne gravity mapping systems

As shown in Figure 1.1b, the challenge of georeferencing gravity measurements has traditionally been met by using damped 2-axis platform systems that attempt to physically maintain a level platform (equivalent to the l-frame) onto which a single specific force sensor is mounted. The most common approach along these lines is the use of a shipborne gravimeter modified to work in the more turbulent airborne environment. As shown, systems of this type have been under development and in use since the first prototypes of the late 1950s and finally became fully operational shortly before 1990

when carrier phase DGPS became available. It is also shown that a number of other airborne gravity system concepts have been implemented since the GPS entered the scene. Among them are those that make use of 3-axis inertial platform systems, SINS and accelerometer triads. An outline of these system concepts can be found in Schwarz and Li (1996b) and an interesting theoretical comparison between them can be found in Czompo and Ferguson (1995). Details about currently available systems of each of these types are given in Section 1.3.

The airborne gravity system currently used at the University of Calgary makes use of a SINS and DGPS. Its development as a gravity sensor has gone hand in hand with its developments as a system for navigation and for optical and digital image mapping over the last two decades. The Calgary system is the focus of the investigations reported in this dissertation.

A SINS contains two sensor triads. The first is an orthogonal set of accelerometers that measure specific force and the second is an orthogonal set of gyroscopes that measure angular velocity. These triads are strapped to the body of the host vehicle and therefore go through the full range of motion experienced by the aircraft. The angular rates sensed by the gyroscopes are integrated to compute the attitude of the SINS with respect to the l-frame. This permits the gravity measurements to be transformed numerically from the measurement frame (b-frame) to the l-frame, rather than the traditional approach in which a level platform is physically maintained. While the strapdown approach has the disadvantage of setting more stringent requirements for the performance of the sensors (e.g. larger dynamic range, higher resolution and better scale factor stability), it has several significant advantages including the facts that:

a) an off-the-shelf SINS can be used that has been designed and sold for navigation purposes, and
b) the gravity system can serve simultaneously as a system for georeferencing the measurements made by other mapping sensors.

These advantages lead to a much smaller size, lower cost, lower power consumption and lower failure rate than other alternatives, while providing considerable flexibility. The strapdown approach is described in more detail in Chapter 2.

The magnitude of the gravity vector can also be estimated from the triad of accelerometers without knowledge of their attitude. Because a SINS contains a triad of accelerometers, the Calgary system is also used in this mode for gravity estimation.

More details about airborne mapping and georeferencing can be found in Schwarz (1998) and a complete history of direct georeferencing is presented in Skaloud (1999). For more information about the history of airborne gravimetry and a more complete list of related references, the reader is referred to Schwarz and Li (1996b) and Gumert (1998). Schwarz and Li (1996b) also offers an introduction to the principles of airborne gravimetry using the SINS-based approach.

1.2 Why are airborne gravity measurements needed?

The gravity field of the Earth provides fundamental information for geoscientists. However, the current state of knowledge of the gravity field is insufficient for many applications. Generally speaking, low and medium-resolution information is essential in geodesy where the goal is to measure and represent the Earth and its gravity field. Medium and high-resolution components contain information that is used in geophysics to identify and characterize sub-terrain features that define the physics of the Earth. Each of these areas of application is treated in turn below.

Throughout this dissertation, the terms high, medium and low-resolution refer to the parts of the gravity spectrum corresponding to spatial resolutions below 5 km, between 5 and 150 km, and above 150 km, respectively. Further, because it is assumed that the reader is familiar with the concept of spatial resolution as it applies to an airborne gravity system, discussion of that topic is found in Appendix C.1. As discussed there, the terms *spatial resolution* and *half-wavelength* are used synonymously throughout this dissertation. More information about the spectral characteristics of the gravity field and of different sources of gravity information can be found in Schwarz (1984).

Major applications in geodesy

As mentioned above, the major task of geodesy is to measure and represent the Earth, including both its surface and its gravity field. Although these may seem like separate features of our planet, they are in fact inherently linked. For example, determination of the height of an object with respect to a meaningful reference surface depends on the

gravity field of the Earth. A major task of geodesy is therefore the determination of the equipotential surface of the gravity field that most closely corresponds to mean sea level. This surface, commonly known as the geoid, is used as a meaningful height reference.

Mapping agencies currently make use of surface measurements and a combination of data obtained from satellites to determine the geoid; the satellite data provides the low-resolution information and the surface data provides the high-resolution information. In the future, it is expected that low-resolution information will be provided by geophysical satellites that are dedicated to the collection of global gravity field information. This is discussed further in Section 6.1.

In Figure 1.2, a number of geodetic applications of gravity data are depicted by the dotted lines, in terms of the ranges of geoid accuracy and spatial resolution required over areas of the planet in which accurate terrestrial gravity data are not already available. The limits of the resolution and accuracy that can be achieved using current global models, data from future dedicated-gravity satellite missions and airborne gravimetry are shown by the solid lines. Several observations follow. First, the use of present global models by themselves clearly does not meet the requirements for many geodetic applications. While this problem is not as great in areas of the world where surface gravity data of medium to high-resolution is available (as in North America, Europe, Japan and Australia), it is severe in areas where they are unavailable. The second observation is that although the data from future satellite missions is expected to greatly improve the situation in many cases, it will not meet the requirements for all applications. Most notable of these are the determination of local geoids and leveling by GPS.

The major task of airborne gravimetry in geodesy is therefore summarized as a requirement to offer a fast, homogeneous, economical and potentially more accurate alternative to surface methods for the collection of high-resolution gravity data.

It is important to note that although airborne gravity has the potential to provide very high-resolution information, the minimum spatial resolution required for ‘cm-level’ geoid determination is roughly 5 km, even in mountainous areas (Li (2000)). Also, because future satellite missions will provide an accurate estimate of the geoid for spatial resolutions greater than about 150 km, an estimate of the gravity field across the whole

spectrum may be obtained by combining the data from these missions with that from airborne gravimetry. In turn, the expected data from the satellite missions also imply that providing low-resolution information will no longer be a stringent requirement of airborne gravity systems (except at the poles where the satellite missions cannot provide data). Taken together, these observations imply that the major goal of airborne gravity researchers working in geodesy is to increase the accuracy of the systems within the bandwidth corresponding roughly to spatial resolutions between 5 and 150 km.

The requirements and uses of airborne gravity data are discussed in more detail in Chapter 6. Further information can be found in Schwarz and Li (1996 a, b).

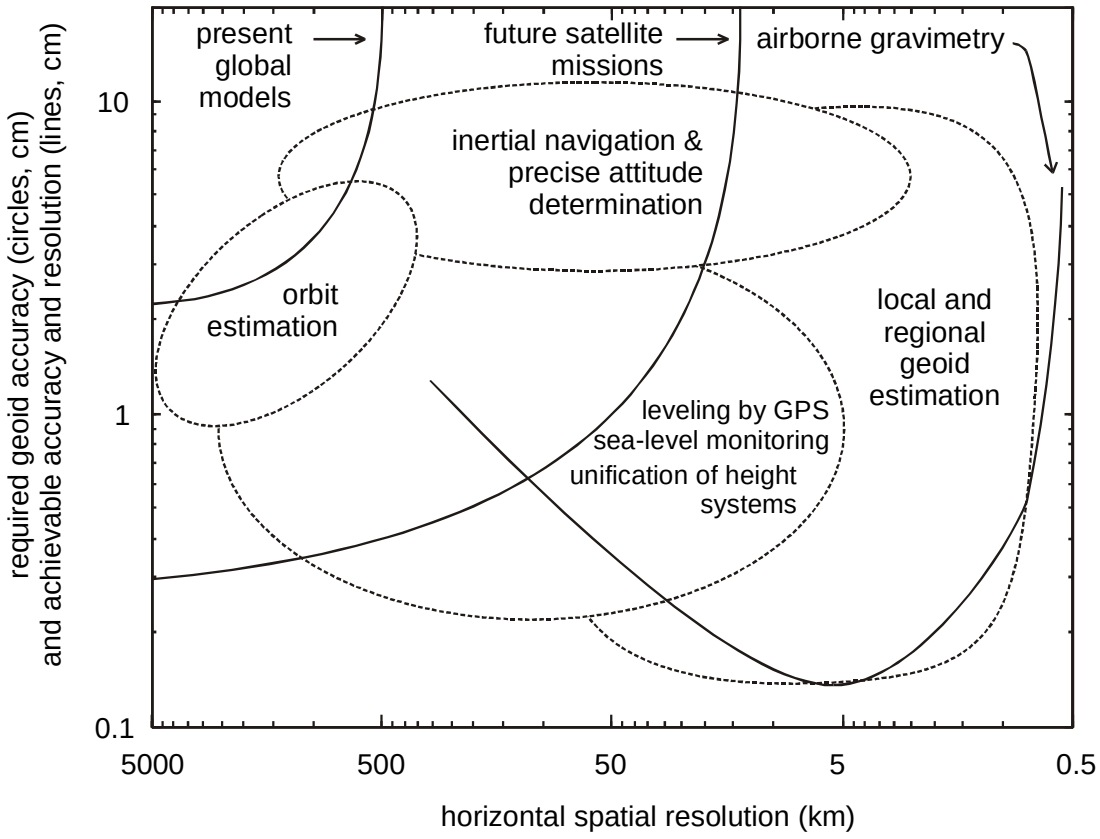

Figure 1.2 The required accuracy of the geoid for geodetic applications as a function of spatial resolution (After ESA (1999) and Schwarz and Li (1996a))

Major applications in geophysics

In the field of geophysics, a goal is often to investigate the crust and near-surface of the Earth for economic reasons. It has long been recognized that traditional methods of collecting gravity data for geophysical applications have limitations that can be overcome by airborne methods. Techniques that use stationary gravity meters to obtain data are relatively expensive and are usually of limited geographical extent. Applications of gravity data in geophysics are varied and depend on the accuracy and spatial resolution with which the data can be obtained. Figure 1.3 depicts a number of such applications in terms of gravity accuracy and spatial resolution requirements, along with an estimate of the best accuracy and resolution limit that airborne gravimetry can currently resolve.

Clearly, airborne methods already meet the requirements for local and regional geological studies and arguably those of volcanology. As discussed in Reynolds (1997), gravity data collection is currently only considered to be one of the primary methods for hydrocarbon exploration and regional geological studies and is still considered to be a secondary method for the exploration of mineral deposits. The latter fact is unlikely to change before the accuracy and resolution of airborne gravity data reliably meet the levels of better than 1-2 mGal and 1-2 km, respectively. This implies that challenges remain for airborne gravity researchers working in the field of geophysics. They include simultaneously improving the accuracy and resolution of their systems so that data can better meet the requirements of high-resolution applications such as resource exploration.

The use of airborne gravity data in geophysics is discussed further in Chapter 6. Information about the roles of gravity and other aero-geophysics data for resource exploration can be found in Sideris et al. (1992), Reynolds (1997) and Johnson (1998).

Other applications

By limiting the discussion in this section to geodesy and geophysics, a number of important applications have not been mentioned. They include the monitoring of ice dynamics, the monitoring of sea level and groundwater, understanding and predicting tectonics and studying the continental lithosphere. If the challenges described for geodesy and resource exploration can be met, then so will those of the above named applications.

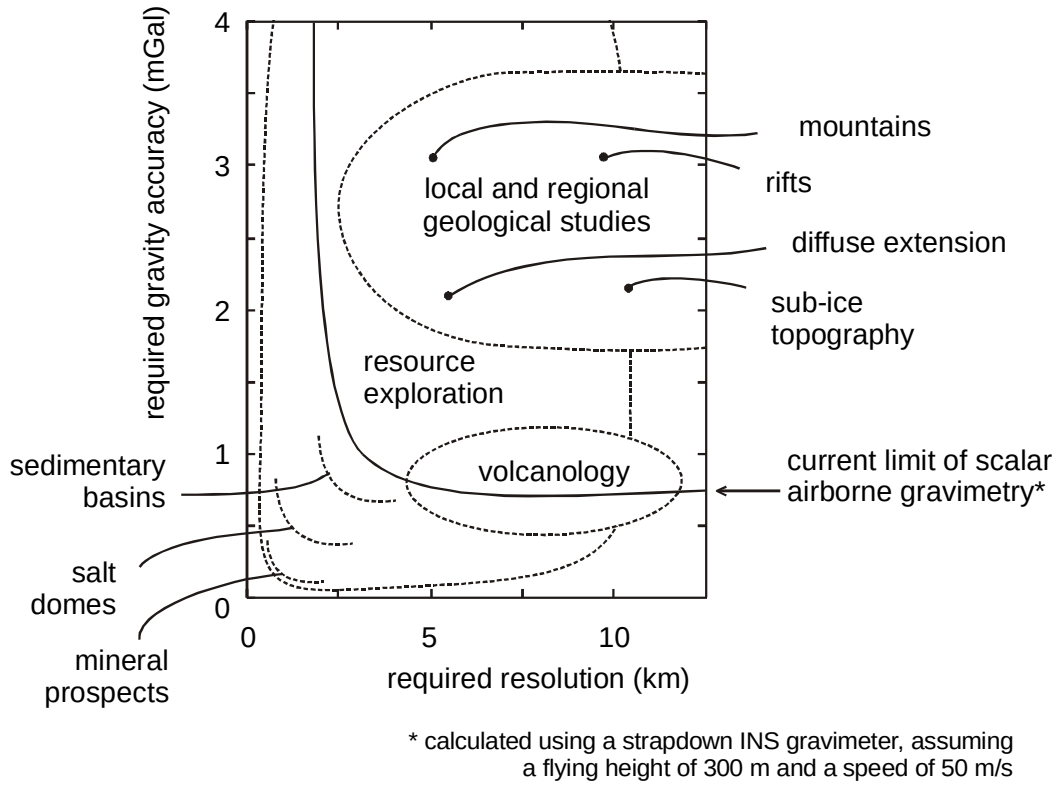

Figure 1.3 The requirements for gravity data for geophysics (After NRC (1995))

1.3 Related research

This section addresses key airborne gravity research being carried out worldwide. For brevity, remarks are concentrated on research that has been carried out since the mid 1990s (which is roughly the time when the research presented in this dissertation began). Detailed accounts of some past and current research activities can be found in Hein (1995), Schwarz and Li (1996b) and Wei (1999).

Recent activities can be classified in the following categories:

a) improvement of traditional airborne gravity systems,

b) development of new system concepts (based on inertial technology),

c) optimization of DGPS processing methods,

d) optimization of filtering and estimation methods, and

e) the application of airborne gravity data in various mapping and exploration applications.

These are each treated below.

Improving traditional scalar gravity systems

Airborne gravimetry using **damped 2-axis platform systems** is a fully operational procedure and has gained acceptance as a standard method for obtaining **low and medium-resolution** gravity information (i.e. for half-wavelengths longer than 5 km). As a result of software and hardware considerations, typically available systems of this type have improved in accuracy since the mid 1990s from 5 mGal at half-wavelengths of 10 km to roughly 2 mGal at shortest half-wavelengths of 6 km, see e.g. Brozena and Childers (2000) and Forsberg et al. (1999). An accuracy of 1 mGal at a half-wavelength resolution of 5 km is also reported in Harrison et al. (1995). The modified shipborne gravimeters that are used in these systems include LaCoste & Romberg, Bell Aerospace and Zero-Length Spring Corporation (ZLS) meters and cost upward of US$ 400,000 to purchase. Major players currently conducting research in this area include those who pioneered its use for wide-area surveys such as the group at the United States Naval Research Laboratory (NRL) and the group at the Danish National Survey and Cadastre (KMS). Extensive documentation of the activities of these groups over the last decade is available and recent reports can be found in Brozena and Childers (2000) and Forsberg et al. (1999), respectively. Other users of such systems include the Institute of Geodesy at the Swiss Federal Institute of Technology in Zurich, Switzerland, see Klingele et al. (1995), the British Antarctic Survey, see Jones (1997), and the Geodetic Survey Division of Geomatics Canada. Commercially available systems of this type include Carson Services, Inc., EDCON Aero Surveys Inc. (now offering the combined services of EDCON and Aero Surveys) and Fugro-LCT (now offering the services of the former LCT).

Developing and improving novel system concepts

New system concepts that have been under development over the last decade are all based on a combination of inertial technology with the DGPS. They are summarized below.

An **inertially stabilized platform system** was first tested by the Inertial Technology Center (ITC) in Moscow in cooperation with the University of Calgary and Canagrav Research Ltd. of Calgary (see Salychev et al. (1994) for details) and continues to be operated in airborne mode by the ITC. As discussed in Ferguson and Hammada (2000), development of another such system called AIRGrav that has been underway at Sander Geophysics Ltd. since the early 1990s was completed in 1999. Such systems usually require significant development or modification costs. Both of those named here are commercially available and have been demonstrated to yield **medium and high-resolution** estimates of the gravity field with an accuracy of 0.5-1.5 mGal.

The use of **gravimeters based on Strapdown Inertial Navigation Systems** was pioneered at the University of Calgary (see e.g. Schwarz et al. (1991)), where extensive development and testing continues. Such a system is physically very small, can be purchased off-the-shelf as a navigation system for between US$ 90,000 and 160,000 and requires no modification. Results of the first airborne test conducted in June 1995 are reported in Wei and Schwarz (1998), showing that relative gravity can be collected along a profile with an accuracy of 2-3 mGal at a half-wavelength resolution of 5 km. It was then demonstrated in Glennie and Schwarz (1999) that the same accuracy and resolution can be achieved in a 100 x 100 km area of the Canadian Rocky Mountains. Both of these results demonstrate the usefulness of the SINS for **medium-resolution** gravity estimation. Using the same system, it is shown in Bruton et al. (2000a) and herein, that the SINS approach can yield an accuracy of 1.5 mGal at a half-wavelength of 2 km and 2.5 mGal at a half-wavelength of 1.4 km, therefore demonstrating its role in **high-resolution** applications. Intermap Technologies Corporation (Intermap) of Calgary, Canada also uses the airborne gravity technology developed at the University of Calgary as a part of their commercially available airborne mapping systems. Interesting results have been obtained using a relatively low-cost SINS (US$ 60,000) by a group at the University of Porto in Portugal. See Bastos et al. (2000), where it is demonstrated that their system can be used to obtain rough estimates of the gravity field in a limited bandwidth (roughly 5-10 mGal for half-wavelengths between 10 and 100 km). Other

groups have recently begun developing airborne gravity systems based on SINSs, including KMS and the Ohio State University.

The use of a **triad of accelerometers** for determining the magnitude of the gravity vector was first published and tested in Czompo (1994) and tests have since been described in Wei and Schwarz (1998) and Glennie and Schwarz (1999). To date, these have demonstrated that the performance of a triad of accelerometers is about 20-30% worse than a comparable SINS. Efforts have been underway since 1994 to test and develop such a system at the Bavarian Academy of Sciences and Humanities in Munich, see Boedecker (1998) for a recent description.

It is discussed in Ferguson and Hammada (2000) that excellent long-term stability of an accelerometer can be achieved by precise custom temperature control. In theory, this can make systems based on the inertially stabilized platform, SINS and accelerometer triad concepts very applicable for low-resolution data collection. Although Sander Geophysics Ltd. has implemented this with great success, it has the drawback of increasing the cost and complexity of the system. Depending on the target application, this may not be deemed worthwhile, especially in the future when the low-resolution information from gravity satellite missions will be globally available.

Two comparisons between the 2-axis damped platform, SINS and 3-axis inertial platform systems have been carried out in the last two years. See Glennie et al. (1999) for an example that compared the first two of these system concepts on a common airborne platform. Except during short periods where one or more of the systems was misbehaving for a known reason, this test served to demonstrate a good agreement between the 2-axis platform and SINS systems for medium-resolution applications. Improvements on the results given in Glennie et al. (1999) are presented in Appendix E. A recent comparison was also carried out within the AGEM project with the goal of comparing all three system concepts on the same flight. Results of that comparison are not yet available, but the campaign is discussed in Appendix A.3 and initial results are given in Section 5.5.

Important developments that have not been mentioned above in this sub-section include those by Hein et al. (1990) and Segawa et al. (2000).

Vector gravimetry

Of the system concepts named above, only those that make use of inertial technology are capable of delivering all three components of the gravity vector, i.e. SINS and inertial platform systems. Horizontal components (deflections of the vertical) are much poorer in accuracy than the vertical component due to gyroscope drifts. This is discussed in Schwarz et al. (1991) and studied in detail in Wei and Schwarz (1994). Estimates of the full vector have been obtained by Jekeli and Kwon (1999) using data form the University of Calgary SINS. They propose a method of estimating the full gravity vector for flight lines that have been surveyed at least twice and demonstrate an agreement of about 8 mGal between the horizontal components and an independent reference.

Deriving acceleration from DGPS

By the mid 1990s, it was generally accepted, that when used in differential mode, the GPS has the potential to derive the acceleration of a moving platform with an accuracy of 1-2 mGal under good environmental conditions within a bandwidth corresponding to a filtering period of 90 seconds (e.g. consider investigations such as those outlined in Brozena et al. (1989), Kleusberg et al. (1990), Hehl (1990), Czompo (1991), Wei et al. (1991) and Van Dierendonck et al. (1994)). Despite these demonstrations, errors due to DGPS in the bandwidth of interest for airborne gravimetry continue to pose a major challenge to researchers attempting to simultaneously increase the accuracy and bandwidth of their systems; good conditions are by no means guaranteed in practice and special care is needed as the bandwidth is widened. These facts imply the need for a better understanding of the behavior of the DGPS error budget for acceleration determination and for models and methods that best respond to the challenges of a widened bandwidth.

Other methods of determining acceleration from GPS that have been proposed and tested over the last few years include Jekeli and Garcia (1997) who demonstrated that it is possible to derive aircraft acceleration directly from GPS phase accelerations and more recently Han et al. (2000) who proposed the use of absolute GPS positioning as an alternative method of determining acceleration (now that SA has been turned off). While both of these offer potential alternatives, they have not been shown to provide results that

are as accurate or reliable as the traditional technique under typical dynamic conditions, and therefore do not currently enjoy practical use.

Filtering and estimation methods

As the bandwidth of airborne gravity systems is increased, it is challenging to remove the sensor and dynamics-induced noise, while leaving the relatively small amplitude gravity signals intact. Extremely low signal-to-noise ratios and the dependence of noise on aircraft dynamics imply the need for sophisticated filtering and estimation techniques. A discussion of this problem and some approaches to its treatment are given in Schwarz and Li (1996b). A number of filtering methods (including both model and frequency domain approaches) were proposed in the first half of the nineties and thoroughly compared using a single data set in Hammada (1996). It was concluded there, that the deterministic frequency domain approach of the low-pass filter is the most appropriate because it makes no a-priori assumptions about the gravity field. Since then, common practice has confirmed this; band-limitation via low-pass filters is the most common approach to noise reduction in use today. The only known exception to this is the group at ITC who continue to enjoy success with model-based approaches, see Salychev and Schwarz (1995).

The subject of designing low-pass filters for airborne gravimetry has been treated recently in a number of publications including Hammada (1997), Childers et al. (1999) and Forsberg et al. (1999). The first of these demonstrates that for a given bandwidth, finite impulse response (FIR) filters designed using optimality criteria outperform those designed using other methods (including windowing). Childers et al. (1999) confirms the importance of using a low-pass filter that has good characteristics within the bandwidth of interest and demonstrates that the cut-off frequency should not be below the highest frequency of the gravity signal (but should be as close to it as possible for a given survey). It also proposes the use of a filter that operates in the frequency domain, resulting in very little data loss due to edge effects. Forsberg et al. (1999) discusses the use of a recursive Butterworth filter with good results.

A method to reduce noise in the bandwidth that remains after low-pass filtering is proposed and implemented in Bruton (1997). It resulted in improvements in the accuracy

and resolution of the University of Calgary airborne gravity system by removing GPS noise within the bandwidth of interest.

Although most airborne gravity filtering methods involve generating the difference between the GPS and specific force data streams and then low-pass filtering them, two other approaches have recently been applied with success. These are described in Bruton and Schwarz (1997) and Cunha (1998). The first proposes and demonstrates the use of an adaptive filter for the estimation of the gravity disturbance and the second includes it in the state vector of the Kalman filter. Both approaches have been shown to remove in-band noise but are not currently widely used outside of the respective research groups.

Biases in gravity field estimates can arise when the specific force sensor is subject to long-term biases. They are usually very small for damped 2-axis meters, as demonstrated in Forsberg et al. (1999) for the LaCoste & Romberg; results achieved using their meter are often as good or better without using a crossover adjustment, for example. As mentioned earlier, Ferguson and Hammada (2000) demonstrate that no special post-mission treatment of system biases is required for their AIRGrav system. However, accelerometer biases are typically a problem for the SINS and accelerometer triad approaches to airborne gravity. Glennie (1999) offers a good discussion of the problem and presents a number of attempts to estimate the effect of the accelerometer biases. This topic is dealt with in detail in Chapter 6. It is well-known that if the survey is flown such that the aircraft flies over the same points at different times, a crossover adjustment can be used to estimate these effects. Glennie and Schwarz (1997) show that this can be used to make a data set consistent (i.e. to reduce all flight lines to a plane in space).

Applications of airborne gravity data

Several national and international research projects are of interest because they have had the effect of improving airborne gravity methods by including them in broader strategies for mapping or exploration. Many of these projects have tackled the challenges of estimating the geoid from airborne gravity data; a process that includes refining the respective airborne gravity systems, representing the airborne data on a level surface through a process called downward continuation and using it in a numerical process to estimate the geoid (often by combination with other sources of gravity data).

The aerogeophysical program at NRL continues its pioneering role in the collection and application of airborne gravity data. Their concentration is on large area, medium-resolution surveying on remote regions of the Earth and they have conducted several smaller, high-resolution surveys. They reported the continuing success of large-scale survey efforts in Greenland, West Antarctica and the Arctic, in Brozena et al. (1997). The NRL airborne gravity system has recently been used in a project to determine a local geoid for coastal oceanography. Results presented in Brozena et al. (2000) describe a relative geoid with an accuracy of approximately 3-5 cm and a spatial resolution of 10 km. The accuracy of the corresponding absolute geoid is estimated to be 10 cm (J. Brozena, personal communication).

As described most recently in Timmen et al. (2000), a large-scale Danish-German-Norwegian-Portuguese cooperation called AGMASCO began in January 1996 with a planned duration of 38 months. It combined airborne gravity and altimetry systems with the goal of developing an airborne geoid mapping system for coastal oceanography. It has been demonstrated to derive relative geoids in coastal regions (Skagerrak, Fram Strait and Azores) accurate to 5 cm from airborne data with a spatial resolution of 6-7 km, Fernandes et al. (2000).

As discussed most recently in Li and Schwarz (2000), researchers at the University of Calgary have continued their efforts to determine the geoid from airborne gravity data, resulting in a relative geoid accurate to 2 cm with a maximum spatial resolution of 4-5 km in an area of the Canadian Rocky Mountains. In a partnership with the University of Calgary, a similar approach has been taken by Intermap Technologies Corporation, as discussed most recently in Wei and Tennant (2000). By directly referring their Interferometric Synthetic Aperture Radar (IFSAR)-based mapping products to a local geoid they determine, it has become part of the services they provide.

In May 1999, a national project called Airborne Gravity for Exploration and Mapping (AGEM) began in Canada. It brings together research partners in Canadian universities, government and industry to develop methodologies and products to fully exploit the potential of airborne gravity for geoid mapping and resource exploration. It includes users and developers of gravimeters based on the 2-axis damped platform, SINS and 3-

axis inertial platform concepts. A major thrust of the AGEM project is to refine Canadian airborne gravity systems to the point that they are of sufficient accuracy and resolution for geophysical exploration and high-resolution local geoid determination. In this case, the goal is to obtain high-resolution relative gravity information and has already resulted in high-resolution estimates of the gravity field (agreements with upward continued ground data of 2.5 mGal at a half-wavelength of 1.5 km when using a SINS-based gravimeter). Geoid determination using this data is underway and results are expected to be made available in 2001.

Other Airborne Gravity Research

Although not directly relevant to this dissertation, the following are included for completeness and because of their importance to the field of airborne gravity. As mentioned above, representing the gravity field on a suitable level surface (or on the surface of the Earth) using measurements taken in the air is not a trivial task. The downward continuation of airborne gravity data has been tackled recently in Forsberg and Kenyon (1995), Tscherning et al. (1997), Novak et al. (2000b), Li (2000), and Wei and Tennant (2000), among others. A detailed comparison of techniques for the downward continuation of airborne gravity data can be found in Novak et al. (2000a). Also, some very interesting work is going on the fields of absolute airborne gravimetry and airborne gravity gradiometry. Recent developments in these fields demonstrate the potential to make both of these measuring concepts more accessible and cost efficient over the next decade. For details, see Brown et al. (2000) and Bell et al. (1998), respectively.

1.4 Statement of the problem

By the mid 1990s, the use of a SINS for gravimetry was a novel technique gaining recognition as an accurate and relatively cheap means of estimating the medium-resolution components of the gravity field of the Earth. This dissertation presents and discusses contributions and improvements made to an existing airborne gravity mapping system of this type. This section outlines its objectives and lists a number of assumptions that are made throughout.

1.4.1 Objectives

The overall objective of the research presented herein is to simultaneously increase the accuracy and the bandwidth of the SINS gravimeter used by the University of Calgary. Specifically, this involves:

- improving the performance of the system for medium-resolution applications such as geoid determination (i.e. for spatial resolutions corresponding to half-wavelengths longer than 5 km and less than 150 km), and
- widening the bandwidth of the system and in turn demonstrating its suitability for high-resolution applications such as resource exploration (i.e. for spatial resolutions corresponding to half-wavelengths shorter than 5 km),

1.4.2 Assumptions

The Honeywell Laseref III (LRF-III), owned by Intermap Technologies Corporation, is used for the evaluations presented herein. It is a high quality civilian grade SINS that falls approximately in the middle of the range of such systems that could be applied to airborne gravity. This implies that a higher quality SINS should yield better results. It is also assumed that the data acquisition process for this system is external to the scope of this dissertation.

All research is carried out using scalar gravimetry.

All surveys described are carried out from conventional fixed wing aircraft (vs. other platforms such as helicopters and balloons). This restriction is strictly financial. Because of the relatively high speeds of fixed wing aircraft, this is meant to imply that significantly higher resolution results might be obtained using the SINS if data could be collected from a slower, stable platform.

Aircraft acceleration is derived from DGPS in all cases. Alternatives such as radar altimetry are not available for study at the University of Calgary.

Since there is very little practical need to obtain the results of a gravity survey in real-time, a post mission approach is taken for all computations. This is an industry standard that significantly increases the allowable complexity of an algorithm and permits the use of non-causal filtering methods.

The problems of downward continuation and geoid determination are considered to be external to the scope of this research; it is assumed that the performance of the system (and improvements made to it) can be assessed and characterized by the behavior exhibited at flying altitude.

Proposed improvements to the system will be evaluated in part by comparing the gravity disturbances it estimates to independent estimates coming from ground gravity data that has been upward continued to the flying altitude. In all cases, real data will be used. Because the accuracy of such an independently determined gravity field is usually high, relative to that coming from the airborne gravity system, this comparison will be used as a basis for characterizing the behavior of the latter as a function of time and frequency. As confirmed in Wu and Sideris (1996), it is also assumed that characterization of system errors along profiles is sufficient information to approximately characterize it as a two-dimensional function of space. When surveys are flown over the same point or along the same line at different times, repeatability at those places will also be used to assess the accuracy of the system.

Low-pass filters that are used throughout have as sharp a transition band as possible. This means that values quoted to assess performance (such as root-mean-square (RMS) and standard deviation) will represent the true accuracy of the system within the bandwidth implied by the cut-off frequency of the filter. This avoids confusion arising from the use of filters with wide transition bands that simultaneously suppress both noise and signal (and therefore take advantage of the relatively low power of the gravity signal at high frequencies to imply that a solution contains relevant information at high frequencies).

Although emphasis is on the use of the Calgary system for measuring the gravity field of the Earth, several of the contributions herein apply equally to its use as a system for navigation and direct georeferencing.

2 Airborne Gravimetry Using a SINS

A brief review of the mathematical background relevant to the task of measuring gravity using a SINS and DGPS is provided in this chapter. In addition, the major sources of error caused by each of these measuring systems are reviewed and placed into the context of the contributions made later in the dissertation. Sections 2.1 and 2.2 closely follow developments in Schwarz and Li (1996b), which should be consulted for details.

2.1 The measurement model of airborne gravimetry

In airborne gravimetry, the goal of the integrated measurement system is to determine the gravity disturbance vector, $\delta\mathbf{g}^l$, as a spatial function. The gravity disturbance vector is the difference between the actual gravity vector, $\mathbf{g}^l$, and the gravity vector given by the so-called normal model of the gravity field, $\boldsymbol{\gamma}^l$, at the same point in space: $\delta\mathbf{g}^l = \mathbf{g}^l - \boldsymbol{\gamma}^l$. The normal gravity vector is based on the gravity potential of an ellipsoid of revolution that has been chosen to best approximate the mass and rotation rate of the Earth, see Heiskanen and Moritz (1969) for details. In other words, an airborne gravity system measures the deviations of the actual gravity field from a global model of best fit. The gravity disturbance is related to the quantities derived from each of the sub-systems by Newton's equation of motion. As shown in Schwarz and Li (1996b), the gravity disturbance can be expressed as follows in the l-frame (that has x, y and z axes pointing east, north and up, respectively):

$$\delta\mathbf{g}^l = \begin{pmatrix} \delta g_e \\ \delta g_n \\ \delta g_u \end{pmatrix} = \dot{\mathbf{v}}^l - \mathbf{f}^l + \left(2\Omega_{ie}^l + \Omega_{el}^l\right)\mathbf{v}^l - \boldsymbol{\gamma}^l, \tag{2.1}$$

where $\dot{\mathbf{v}}^l$ and $\mathbf{v}^l$ are the acceleration and velocity of the aircraft derived from carrier phase DGPS, $\mathbf{f}^l$ is the specific force measured by an accurate inertial system (the sum of the gravity and the acceleration due to motion), and the third term is the Coriolis vector arising from two effects:

a) the angular rate of the e-frame with respect to an inertial frame of reference, Ω_{ie}^{l},
b) and the angular rate of the l-frame (centered at the aircraft) with respect to the e-frame, Ω_{el}^{l}.

Strictly speaking, equation 2.1 describes the model of a gravity sensing system that uses a local-level platform (in which the orientation of the accelerometer(s) is realized by an electro-mechanical feedback system). As discussed in Section 2.3, the accelerometers and gyroscopes in a SINS undergo the same motion as the body of the aircraft, implying that:

$$\mathbf{f}^{l} = \mathbf{R}_{b}^{l}\,\mathbf{f}^{b}, \tag{2.2}$$

where $\mathbf{f}^{b}$ is the vector of specific forces measured in the b-frame and $\mathbf{R}_{b}^{l}$ is the transformation matrix that relates them to the l-frame. Substituting equation 2.2 into equation 2.1 gives the following form of the measurement model for airborne **vector gravimetry** (Schwarz and Li (1996b)):

$$\delta\mathbf{g}^{l} = \dot{\mathbf{v}}^{l} - \mathbf{R}_{b}^{l}\,\mathbf{f}^{b} + \left(2\Omega_{ie}^{l} + \Omega_{el}^{l}\right)\mathbf{v}^{l} - \boldsymbol{\gamma}^{l}, \tag{2.3}$$

A subset of this is the **scalar gravimetry** case where only the third component of the above vector equation is of interest. When written out explicitly, the third component has the following form (Schwarz and Li (1996b)):

$$\delta g_{u} = \dot{v}_{u} - f_{u} - \left(\frac{v_{e}}{R_{1}+h} + 2\omega_{ie}\cos\varphi\right)v_{e} - \frac{v_{n}^{2}}{R_{2}+h} - \gamma_{u}, \tag{2.4}$$

where v_e, v_n and v_u are the east, north and upward components of the velocity vector, $\dot{v}_u$ is the time derivative of v_u, f_u is the upward component of the specific force vector, and φ, h, R_1 and R_2 are the geodetic latitude, ellipsoidal height and prime vertical and meridian radii of curvature (all with respect to an ellipsoid of revolution), respectively. The variables γ_u and ω_{ie} are the z-components of the vectors that represent the normal gravity model and the angular rotation rate of the Earth, respectively.

The terms Strapdown Inertial Vector Gravimetry (SIVG) and Strapdown Inertial Scalar Gravimetry (SISG) are employed when a SINS is used to make the measurements of specific force needed to satisfy equations 2.3 and 2.4, respectively. These approaches are discussed and compared in detail in Wei and Schwarz (1998).

Not included in the above equations is a correction that is made to take into account the fact that the GPS antenna and the SINS cannot be located at exactly the same point in space. Practically, this can be done by either correcting the time series of positions and velocities obtained from DGPS to the center of the SINS using the spatial offset and its time derivative (prior to deriving the acceleration), or by deriving a correction term to be applied after estimating the acceleration. This is called the lever-arm correction and is implemented in the acceleration domain for the system used herein.

In another approach to scalar gravimetry, the magnitude of the gravity disturbance vector is estimated more directly by deriving the magnitude of the specific force vector from a triad of accelerometers. Because the accuracy requirement for the orientation of the aircraft is relatively low in this case (it is only used to estimate the lever-arm correction), this approach is referred to as Rotation Invariant Scalar Gravimetry (RISG). It was first published in Czompo (1994) and is also described in Wei and Schwarz (1998).

2.2 The error model of airborne gravimetry

The error model for airborne gravimetry is obtained by linearizing equation 2.3. This is done in Schwarz and Li (1996b) to yield the following relationship between the errors in the gravity disturbance $d\delta\mathbf{g}^l$ and the errors arising from the measuring systems:

$$d\delta\mathbf{g}^l = \mathbf{F}^l\boldsymbol{\varepsilon}^l - \mathbf{R}_b^l\, d\mathbf{f}^b + d\dot{\mathbf{v}}^l + \left(2\Omega_{ie}^l + \Omega_{el}^l\right)d\mathbf{v}^l - \mathbf{V}^l\left(2d\boldsymbol{\omega}_{ie}^l + d\boldsymbol{\omega}_{el}^l\right) - d\boldsymbol{\gamma}^l, \quad (2.5)$$

where $\boldsymbol{\varepsilon}^l$ represents errors in attitude (i.e. estimated vehicle orientation) due to initial misalignment and gyro measurement errors, $d\mathbf{f}^b$ represents accelerometer measurement errors, $d\mathbf{v}^l$ and $d\dot{\mathbf{v}}^l$ represent errors in the aircraft velocity and acceleration determined from DGPS, $d\boldsymbol{\omega}_{ie}^l$ and $d\boldsymbol{\omega}_{el}^l$ are errors in angular velocity and $d\boldsymbol{\gamma}^l$ are errors in the computation of the normal gravity vector. The matrices $\mathbf{F}^l$ and $\mathbf{V}^l$ are the specific force and velocity vectors in their skew-symmetric forms. As shown in Schwarz and Wei (1994), if the accuracy of the position and velocity derived from DGPS can be maintained below 50 cm and 5 cm/s (1-σ) respectively (which is a condition that is evaluated in Chapters 3 and 4), the last three terms in equation 2.5 can be neglected. Together with the inclusion of a term to reflect errors that result from imperfections in the

synchronization of the data coming from the GPS and SINS sub-systems, this results in the following simplified form of the gravity disturbance error model:

$$d\delta \mathbf{g}^{l} = \begin{pmatrix} d\delta g_e \\ d\delta g_n \\ d\delta g_u \end{pmatrix} = \mathbf{F}^{l}\boldsymbol{\varepsilon}^{l} - \mathbf{R}_b^{l}\, d\mathbf{f}^{b} + d\dot{\mathbf{v}}^{l} + (\dot{\mathbf{R}}_b^{l}\mathbf{f}^{b} + \mathbf{R}_b^{l}\dot{\mathbf{f}}^{b})\, dT, \tag{2.6}$$

where dT represents the synchronization errors. For the scalar case, the third component of equation 2.6 is taken, as follows:

$$d\delta g_u = f_e \varepsilon_n - f_n \varepsilon_e - df_u + d\dot{v}_u + (\dot{\mathbf{A}}\mathbf{f}^{b} + \mathbf{A}\dot{\mathbf{f}}^{b})\, dT, \tag{2.7}$$

where as usual, the subscripts e, n and u represent the east, north and upward components of the vector component they follow and, as shown in Wei and Schwarz (1998), $\mathbf{A}$ and $\dot{\mathbf{A}}$ are row matrices of the form:

$$\mathbf{A} = [-\cos\theta\sin\varphi \quad \sin\theta \quad \cos\theta\cos\varphi] \text{ and} \tag{2.8}$$

$$\dot{\mathbf{A}} = [\,\dot{\theta}\sin\theta\sin\varphi - \dot{\varphi}\cos\theta\cos\varphi \quad \dot{\theta}\cos\theta \quad -\dot{\theta}\sin\theta\cos\varphi - \dot{\varphi}\cos\theta\sin\varphi]. \tag{2.9}$$

The variables φ and θ are the roll and pitch of the aircraft and the dots above the variables again represent the time derivatives.

The error model for the RISG approach can be found in Wei and Schwarz (1998).

2.3 The estimation process, the observables and the measuring systems

Despite the seemingly simple relationships between the gravity disturbance and the quantities derived from the sub-systems presented in the last section, their practical implementation is quite involved. Figure 2.1 shows a summary of the process of using the raw measurements made by the sensors to derive the final estimate of the gravity disturbance. As shown in Figure 2.1c, equations 2.3 and 2.4 only constitute a part of the software package called GRavity Estimates from Airborne Techniques for Geoid UNdulations (GREATGUN), that in turn only constitutes one part of the overall process. The other steps shown in Figure 2.1 are necessary to derive the required quantities from the measurements and to remove the high levels of noise. Each noise source is discussed herein, after the following brief review of the entire process.

2.3.1 The process of estimating the gravity disturbance

The first step shown in Figure 2.1a is to interpret and pre-filter the raw measurements coming from the data acquisition system. This can involve low-pass filtering the gyroscope and accelerometer measurements from the SINS (see Skaloud (1999)) and replacing the raw phase-rate measurements by more accurate measurements derived from the DGPS carrier phase (see Bruton et al. (1999) and Chapter 5).

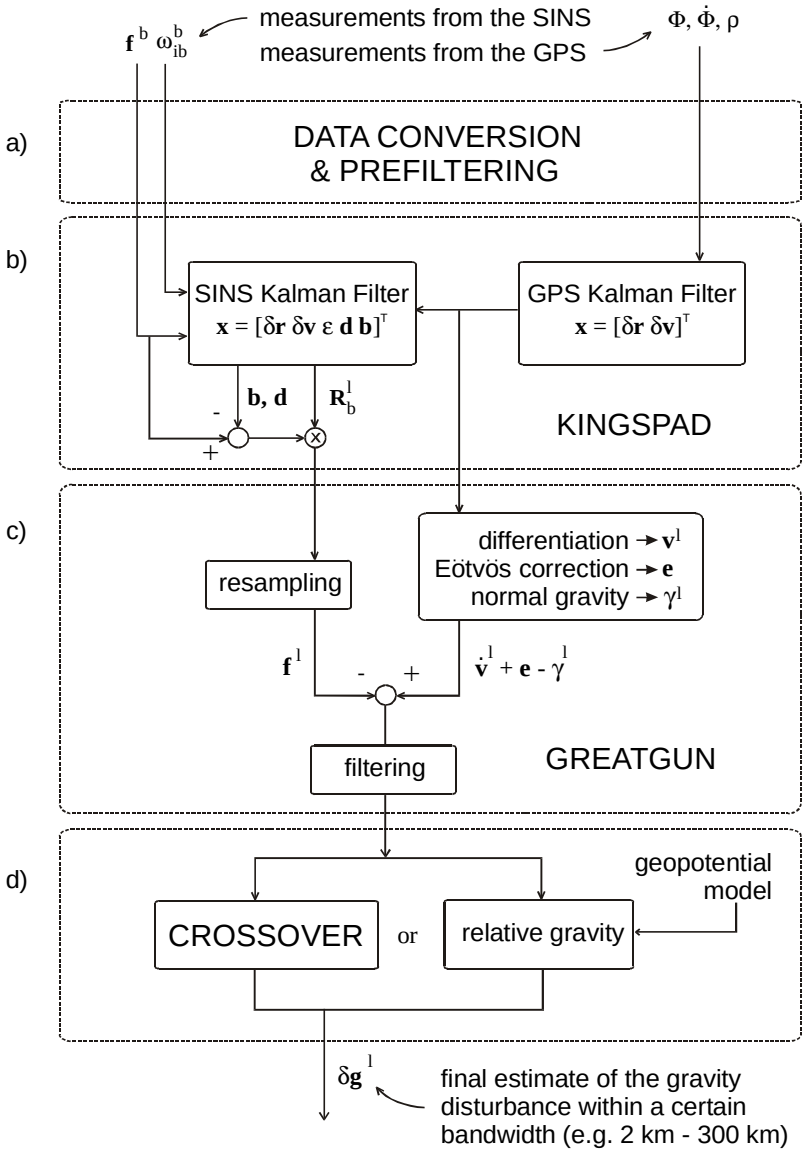

Figure 2.1 A summary of the airborne gravity data processing procedure at the U of C

The second step is the derivation of the position, **r**, velocity, **v,** and attitude matrix, $\mathbf{R}_b^l$, of the aircraft from these filtered measurements. This is carried out using a software package called KINematic Geodetic System for Position and Attitude Determination (KINGSPAD) that was written prior to the start of the research presented herein, mainly by M. Wei. As discussed in El-Sheimy and Schwarz (2000), it implements a decentralized Kalman filter to integrate the SINS and DGPS measurements. The integration process corrects for cycle slips in the GPS measurements prior to using them to derive the position and velocity of the aircraft in a double difference approach. As implied by Figure 2.1b, the SINS filter uses DGPS to estimate and apply lumped error terms representing the biases in the accelerometer measurements, **b**, and drifts in the gyroscope measurements, **d,** before estimating the attitude of the vehicle.

As shown, the GREATGUN package uses the estimated attitude of the aircraft to transform the corrected (and re-sampled) specific forces from the b-frame to the l-frame, according to equation 2.2. The position and velocity of the vehicle are then differentiated appropriately to obtain the acceleration of the vehicle. According to equation 2.3, this is used along with estimates of the Coriolis and normal gravity vectors to compensate for the motion of the vehicle and effectively isolate the gravity disturbance vector. Because the resulting signal is very noisy (see Sections 2.3.2 to 2.3.4 for details), its derivation is followed by a filtering step that removes most of the noise by low-pass filtering the data. The first version of GREATGUN consisted of 12 separate programs written by M. Wei. These 12 programs were combined into the current single software package by the author in January 1998 and underwent minor modifications by both C. Glennie and the author before December 1998. A number of more significant changes have been proposed and implemented by the author since that time (as discussed herein, in Part 3).

The third and final step is to deal with remaining low-frequency errors in the estimated gravity disturbance (that are mainly the result of residual accelerometer biases). As shown in Figure 2.1d, there are two approaches to this. In the first approach, a crossover adjustment can be used to reduce the measurements to an arbitrary plane in space. See Glennie and Schwarz (1997) and Kennedy (1999) for details. The absolute orientation of

this plane in space can then be determined if the value of the gravity disturbance is known for at least three points in the vicinity of the airborne mission. In the second approach, only relative gravity field information is derived (e.g. for spatial resolutions higher than 150 km). The latter approach is described and compared to the crossover approach in Section 6.4 of Chapter 6.

2.3.2 The SINS as a measuring system

Sections 2.3.2 and 2.3.3 are included to briefly demonstrate the important characteristics of some crucial SINS and DGPS errors affecting the estimation of the gravity disturbance, i.e. those in equations 2.6 and 2.7. As discussed above, the sensors in a SINS include two orthogonal triads; three accelerometers and three gyroscopes that both undergo the full (rotational and translational) motion of the aircraft. The accelerometers measure the specific force, $\mathbf{f}^b$, that the body of the aircraft experiences and the gyroscopes measure its angular velocity, $\boldsymbol{\omega}_{ib}^b$, both with respect to an inertial frame of reference. Each is affected by systematic and random errors, which tend to result in a combination of low and high-resolution errors, respectively, in estimates of the gravity field.

As shown in equation 2.6, the **errors in the specific force measurement** contribute directly to the error budget of the gravity disturbance estimate. They can be expressed in the b-frame as follows:

$$d\mathbf{f}^b = \mathbf{b} + \mathbf{S}_f \mathbf{f}^b + \mathbf{N}_f \mathbf{f}^b + \mathbf{n}_f , \tag{2.10}$$

where $\mathbf{b}$ is a vector representing the lumped sum of the biases influencing the accelerometer measurement, $\mathbf{S}_f$ is a diagonal matrix representing scale factor errors, $\mathbf{N}_f$ is a skew-symmetric matrix representing the non-orthogonality of the sensor axes and $\mathbf{n}_f$ is a vector representing random errors. See Titterton and Weston (1997) for details about this error equation. Generally speaking, while scale factors, non-orthogonalities and some biases can be calibrated in the laboratory, see e.g. Titterton and Weston (1997) and Shin (2000), residual biases and random errors cannot. The result is a combination of a low-frequency error (owing to the residual biases) and a broad-band error (owing to the random errors). The specifications for the accelerometers in the LRF-III SINS (that are

given in Appendix B) provide an indication of the magnitude of each of these errors for the airborne gravity system used at the University of Calgary. Consider Figure 2.2 that is included to demonstrate the errors affecting estimates of the scalar gravity disturbance by characterizing each of the errors in equation 2.7 as a function of frequency. The spectra in Figures 2-2 a, b and c have been shifted along the y-axis by an arbitrary amount so that they can be distinguished from one another. The figure is derived from results obtained in Chapters 5, 6 and 7.

It is demonstrated in Figure 2.2d that the effect that the accelerometer errors has can be separated into two regions of the spectrum, with large bias-like errors below about 0.00025 Hz and broad-band errors having a much smaller amplitude above 0.00025 Hz. Clearly, the biases present a major difficulty for estimating low-resolution components of the gravity field. They are treated in detail in Glennie (1999) and are attributed to an uncompensated temperature effect on the accelerometers. A general conclusion of that research is that the associated problems are unlikely to be resolved without the design and implementation of a system that uses higher quality, temperature compensated accelerometers. These low-frequency accelerometer errors are studied in detail in Chapter 6.

The effect that the **errors in the gyroscopic measurements of angular velocity** have on the errors in the gravity disturbance cannot be stated as simply. As discussed in Section 2.1, the major role of those measurements in an airborne gravity system is to derive the orientation of the b-frame with respect to the l-frame so that the specific force measurements can be transformed from the former to the latter. As shown in equation 2.6, the errors in the estimated gravity are a function of not only the errors, $\boldsymbol{\varepsilon}^{l}$, in this estimate of the attitude but also of the specific force experienced by the aircraft, $\mathbf{f}^{l}$. Formally, this is written as follows (see Schwarz and Li (1996b)):

$$\mathbf{F}^{l}\boldsymbol{\varepsilon}^{l} = \begin{pmatrix} f_n\varepsilon_u - f_u\varepsilon_n \\ f_u\varepsilon_e - f_e\varepsilon_u \\ f_e\varepsilon_n - f_n\varepsilon_e \end{pmatrix}, \tag{2.11}$$

where as usual, the subscripts e, n and u represent the east, north and upward components of the vector component they follow.

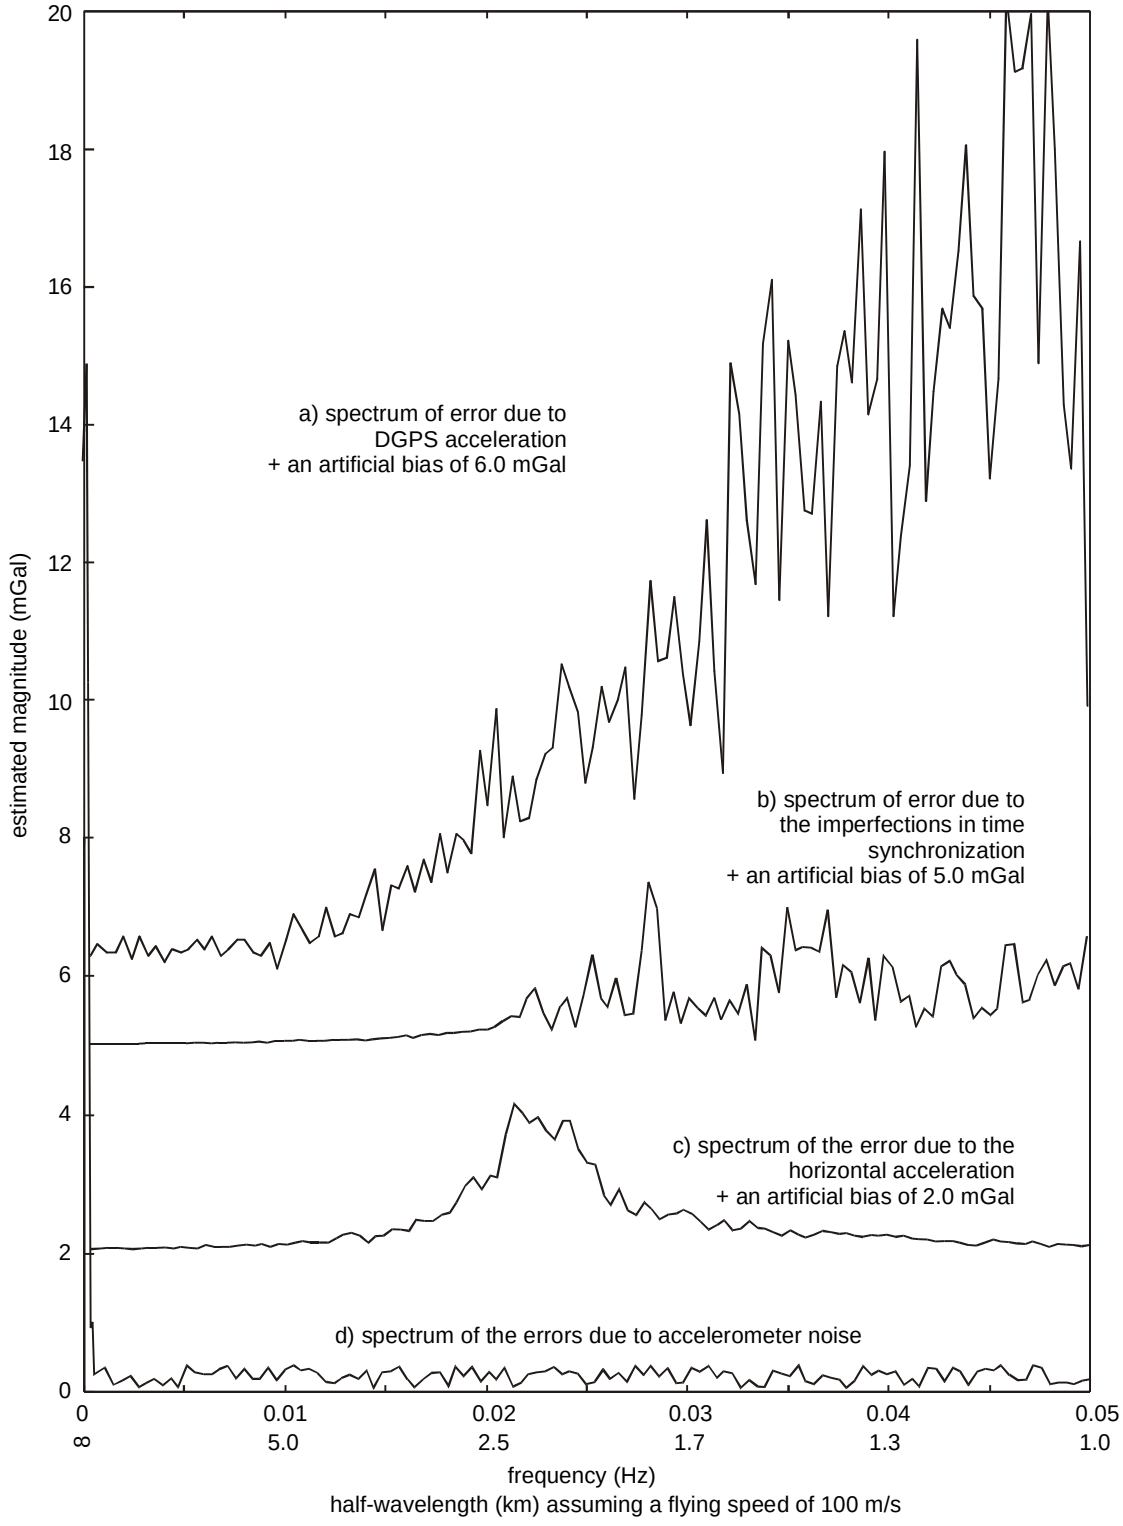

Figure 2.2 Spectra of samples of the errors affecting scalar gravimetry

Expressed in the l-frame, the errors in the estimated orientation of the aircraft can be approximated for time, t_k, by the following equation (Schwarz and Li (1996b)):

$$\varepsilon^{l}_{t_k} = \varepsilon^{l}_{t_0} + \int_{t_0}^{t_k} \mathbf{R}^{l}_{b}(t)\, d\boldsymbol{\omega}^{b}_{ib}(t)dt\,, \tag{2.12}$$

where the first term represents the errors due to the initial misalignment (including insufficiencies in the normal gravity model) and the second term represents integrated gyroscope measurement errors at time t_k. The term $d\boldsymbol{\omega}^{b}_{ib}$ represents the gyroscope measurement errors that are in turn given by:

$$d\boldsymbol{\omega}^{b}_{ib} = \mathbf{d} + \mathbf{S}_{\omega}\boldsymbol{\omega}^{b}_{ib} + \mathbf{N}_{\omega}\boldsymbol{\omega}^{b}_{ib} + \mathbf{n}_{\omega}\,, \tag{2.13}$$

where $\mathbf{d}$ is a vector representing the bias error that results in gyroscope drift, $\mathbf{S}_{\omega}$ is a diagonal matrix representing scale factor errors, $\mathbf{N}_{\omega}$ is a skew symmetric matrix representing the non-orthogonality of the sensor axes and $\mathbf{n}_{\omega}$ is a vector representing random errors. More details about equation 2.13 can be found in Titterton and Weston (1997). The error in attitude is therefore the combination of a low-frequency oscillation owing to the initial errors and a faster changing component due to the interaction of aircraft dynamics with the constant and slowly varying gyroscope sensor errors. As shown in Skaloud (1999), the magnitude of the error in attitude for a navigation grade SINS is expected to be below 50 arc seconds (with standard deviations of less than 15 arc seconds, for any given flight line). This is discussed further in Chapter 7. An estimate of the effect of the errors in attitude on the upward component of the gravity disturbance (i.e. the third component of equation 2.11) is shown in Figure 2.2c for a sample flight in which it is large owing to high horizontal accelerations. The effect is periodic in nature and can have a peak amplitude of between 2 and 10 mGal at a central frequency anywhere above 0.01 Hz (depending on the aircraft dynamics and the flying speed). For typical flying speeds between 45 and 100 m/s, this causes a major difficulty in estimating medium and high-resolution components of the gravity field. This error source is treated in detail in Chapter 7.

It should be observed that equation 2.11 explains why constant velocity conditions are favorable for airborne gravity surveying and why determination of the horizontal

components of the gravity vector is extremely challenging. If the flight azimuth and velocity are maintained nearly constant for a given survey line, then the specific force vector is approximately given by $[0,0,g]$, where g ~ -10 m/s^2. It is easily seen that under these conditions, equation 2.11 becomes $[-g\varepsilon_n, g\varepsilon_e, 0]$. That is, while maintaining a constant velocity drives the effect of the attitude error to a minimum (and towards zero for the vertical component), a portion of its effect in the horizontal components is always amplified by the magnitude of the gravity field.

The models and algorithms for strapdown inertial navigation are well-documented and can be found for example, in Britting (1971) and Wei and Schwarz (1990). The process of attitude estimation is discussed further in Chapter 7 and described in detail in Schwarz (1998) and Skaloud (1999).

2.3.3 The GPS as a measuring system

This sub-section briefly describes the processes of estimating position, velocity and acceleration from the measurements made by the GPS receivers. This background information is then used to define the relationship between errors in those measurements and errors in the gravity disturbance. This discussion of DGPS positioning is based on a similar discussion in Teunissen and Kleusberg (1998b).

The observables in the GPS include the pseudorange from each satellite in view to the user antenna and the carrier phase and phase rate (Doppler) measurements. More details about these can be found in Hoffmann-Wellenhof et al. (1994) and Langley (1998). In high accuracy DGPS, linear combinations of the carrier phase observable are formed that essentially become high accuracy range differences if the ambiguous number of cycles can be reliably estimated in a timely manner. While a number of linear combinations are possible, the single-frequency double-difference observable is introduced here because it provides a useful summary of the relevant error sources. The measurements made to each satellite from a stationary GPS receiver (located at a known point) are subtracted from those made to the same satellites by the moving receiver. This process, known as single-differencing, eliminates errors due to the satellite clocks and (depending on the distance separating the GPS antennas) reduces the errors due to atmospheric propagation and

errors in the position of the satellite. The double-difference observable is then formed by taking the difference between two such single-differences that in turn essentially eliminates the receiver clock errors. This yields the following double difference observation equation for the carrier phase case (after Teunissen and Kleusberg (1998b)):

$$\nabla\Delta\,\Phi \;=\; \nabla\Delta\,\rho + \nabla\Delta\,\delta\rho + \lambda\nabla\Delta\,N \;+\; \nabla\Delta\,T - \nabla\Delta\,I + \nabla\Delta\,m_{\Phi} + \nabla\Delta\,\varepsilon_{\Phi}\,, \qquad (2.14)$$

in metres where the double difference symbol $\nabla\Delta$ refers to differencing between satellites and receivers and where:

Φ	is the carrier phase (measurement) [m],
ρ	is the range between the satellite and the user antenna (observable) [m],
$\delta\rho$	is the ephemeris error [m],
λ	is the wavelength of the carrier [m/cycle],
N	is the ambiguous integer carrier phase cycle count [cycles],
T	is the measurement delay due to the neutral atmosphere [m],
I	is the measurement delay due to the ionosphere [m],
m_{Φ}	is the measurement delay due to carrier phase multipath [m] and
ε_{Φ}	is the noise of the carrier phase measurement [m].

Such an observation equation can be formed for each pair of satellites, resulting in *m* equations (one equation less than the number of satellites) at each epoch. For the purposes of this discussion, it will be assumed that in order to solve for the unknown baseline between the master station and the aircraft, the errors due to noise and multipath can be treated as stochastic variables and that the errors owing to the atmosphere and broadcast ephemerides can be either modeled appropriately or neglected (especially over short and medium baselines). This leaves the *m* integer ambiguity terms and the 3 components of the baseline between the master station and the aircraft as unknowns. The following linear system of equations can then be formed at each epoch using the double difference observation equations:

$$\mathbf{z} = \mathbf{A}\mathbf{a} + \mathbf{B}\mathbf{b} + \mathbf{e}, \qquad (2.15)$$

where **z** is the vector of *m* (observed minus computed) double difference carrier phases, **a** is the unknown vector of integer ambiguities, **b** is the unknown vector of baseline components, **A** and **B** are the design matrices for the ambiguity and baseline components and **e** is the measurement noise vector. Note that it is standard to use other linear combinations that incorporate pseudorange, Doppler and dual frequency carrier phase

data into this system, see e.g. Teunissen and Kleusberg (1998b). After the ambiguity terms have been correctly solved for (or their real-valued estimates sufficiently well determined), these linear combinations permit the solution of an accurate baseline. The reader is referred to Teunissen (1998) for details about ambiguity resolution. The estimate of the baseline in turn permits the estimation of the position and velocity of the aircraft relative to the known position of the master station.

When moving, the trajectory of the aircraft is often estimated in a recursive manner using a Kalman filter that plays the role of balancing new estimates of the position with estimates that it predicts based on past measurements and a model of the dynamics.

Because the models of the measurements and the vehicle dynamics are not perfect, residual errors show up in the output of the system of equations given by equation 2.15. These errors each play a direct role in the accuracy of the estimated trajectory. This issue is the topic of Chapter 3.

In airborne gravimetry, where a major role of the GPS system is to provide estimates of the acceleration of the aircraft, differentiation of the estimated position (or velocity) is required. The residual errors in equation 2.14 also play a role in the quality of the estimate of acceleration, as does the process of differentiation itself. Consider Figure 2.2a that shows the error spectrum of the acceleration of an aircraft determined from the GPS data collected in a recent airborne gravity survey. The errors are clearly small for frequencies below 0.01 Hz, but increase very significantly as a function of frequency. This represents one of the biggest challenges to widening the bandwidth of any existing airborne gravity system. Investigations of the differentiation process and the effects of each of the error sources for determining acceleration are presented in Chapters 4 and 5, respectively.

For information about the kinematic models and the Kalman filter used for trajectory determination, see Schwarz et al. (1989) and Tiberius (1998). For a discussion about the determination of acceleration from DGPS and a list of further reference material, see Schwarz and Li (1996b).

2.3.4 The effects of other sources of error on gravity disturbance estimation

Errors that have not already been treated are now considered. As shown in equation 2.6, errors due to imperfections in the synchronization of the SINS and DGPS data streams play a role in the accuracy of the estimated gravity disturbance. For the data acquisition system used at the University of Calgary, the time synchronization errors, dT, can be as high as 0.5 ms. They are due to registration errors that occur in the data acquisition system (when the data registration and time tagging process is blocked by other processes with higher priority). See Schwarz and Li (1996b) for more details about this. Clearly, these have a larger effect as the bandwidth of the system is increased, implying the need for a better data acquisition system. Consider Figure 2.2b that gives an approximate worst-case error spectrum for these errors, derived for a survey in which the rotational and translational dynamics are large.

For simplicity, dither, quantization and second-order scale factor errors have been neglected in equations 2.10 and 2.13. Their treatment can be found for example in Wei and Schwarz (1994) and Titterton and Weston (1997).

Beyond the inclusion of equations 2.10 and 2.13, demonstration of the dependence and magnitude of the uncalibrated scale factor and non-orthogonality errors on aircraft dynamics has also been neglected.

2.4 Contributions of this dissertation

In the last three sections, it has been described how the raw SINS and DGPS measurements are used to estimate the gravity disturbance and real data is used to demonstrate how the major errors in those measurements affect the latter as a function of frequency. It is shown in the above that the resolution of the system is primarily limited by accelerometer biases on the low-resolution end of the spectrum and by a combination of attitude and DGPS errors (and to some extent synchronization errors) at the high-resolution end of the spectrum. It is also shown that the current accuracy of the system for medium resolutions is a function of both the accelerometer and DGPS errors.

The above is summarized in Figure 2.3 using an approximate spatial scale corresponding to an assumed flying speed of 100 m/s. Two observations emerge from Figure 2.3. First,

the relationship between the spatial requirements of gravity data and the challenges described in this section are evident. Second, the research presented herein may be placed into context as follows. Recall from Chapter 1 that the overall objective of the research presented herein is to simultaneously increase the accuracy and the bandwidth of the SINS gravimeter used by the University of Calgary. Because they are common to airborne gravity systems of all types, the errors owing to the DGPS error budget are investigated in Part 2 of this dissertation. By means of detailed analyses of the errors in position, velocity and acceleration, it will be shown that the errors in the estimate of the required gravity disturbance due to the GPS can be significantly reduced, both within the current bandwidth and especially for higher frequencies.

Errors due to the SINS are treated in Part 3. The low-frequency errors arising from the accelerometers are investigated in Chapter 6. The dynamics of survey aircraft are studied in detail in Chapter 7, both from a theoretical point of view and empirically by using data collected on two survey aircraft. This leads to the proposal of a method for estimating and removing most of the effect of attitude induced errors, effectively removing one of the major barriers to high-frequency gravity field estimation.

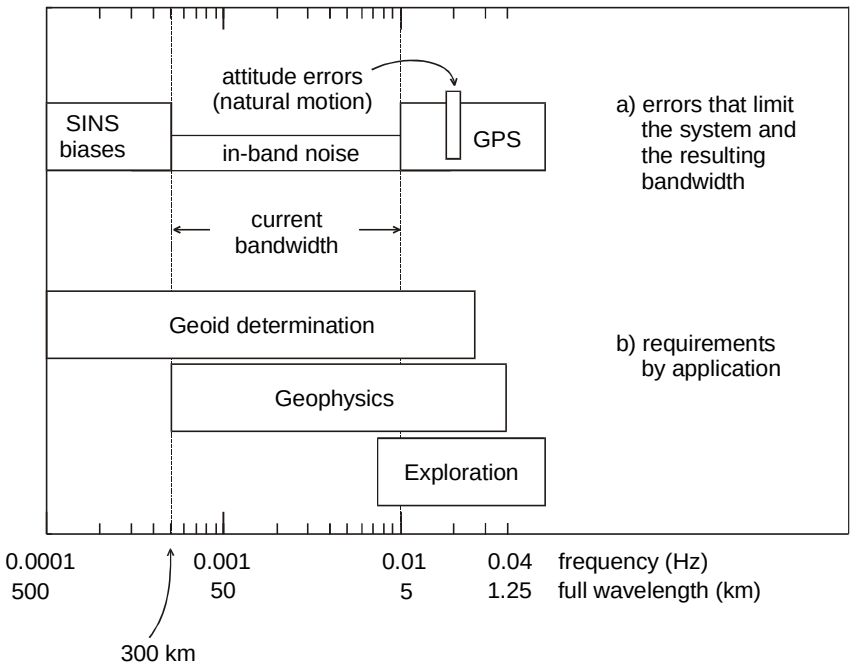

Figure 2.3 The applications of airborne gravity and the research challenges

PART 2: DGPS FOR AIRBORNE GRAVITY MAPPING

The objectives in Part 2 are to evaluate the capabilities of carrier phase DGPS for airborne mapping applications and to recommend methods for its use, with emphasis on airborne gravimetry.

In Chapter 3, an investigation is reported on the limits of the accuracy with which DGPS can be used for positioning an aircraft in flight. This research is carried out entirely from the point of view of a user who is working in the field of airborne mapping, i.e. existing commercial software has been used. This work contributes to research in the field by using real data from several well-chosen data sets, along with various data processing strategies to isolate and quantify each relevant component of the DGPS error budget. This part of the research is based on the contributions made by the author to Bruton et al. (2000b).

In Chapter 4, the estimation of high precision velocity and acceleration are considered. This is done by means of a fundamental consideration of the concept of differentiation. A number of alternative methods for the differentiation of GPS data are reviewed for deriving a Doppler measurement from the carrier phase measurements and for differentiating positions that are estimated by DGPS software. Using this approach, real data is used from a number of experiments that represent varying dynamics (e.g. static, low and high frequency dynamics). This research is based on Bruton et al. (1999).

In Chapter 5, a detailed study is presented of the DGPS error budget for acceleration determination, within the bandwidth that is relevant to airborne gravimetry. Each of the residual error sources is categorically isolated and characterized as a function of frequency. Using this, data processing strategies are recommended for both of the major application areas (i.e. geoid determination and resource exploration). It is demonstrated that very accurate estimates of the gravity field can be obtained for high-frequency applications such as resource exploration. The research presented is based entirely on the contributions made by the author to Bruton et al. (2000a).

3 On the Positioning Accuracy of Kinematic Carrier Phase DGPS

In static applications, carrier phase DGPS can be used to determine the position of an object to the centimetre and even millimetre level with respect to a fixed master station, depending on the extent of the data processing that is done and the occupation time. However, this level of accuracy cannot be expected for a moving vehicle, especially under typical dynamics and environmental conditions and as the distance separating the vehicle and the master station increases. Errors in the estimated absolute position of the survey aircraft obviously play a critical role in the performance of airborne mapping systems. The goal of this chapter is to comment on the accuracy with which carrier phase DGPS positioning of an airborne platform can be achieved. Although the emphasis is on meeting the somewhat more forgiving positioning requirements of airborne gravimetry, the findings are also related to the requirements of other airborne mapping applications.

3.1 Background

An airborne gravity system depends on the estimated position of the aircraft in several ways. The most obvious dependence arises from the process of determining acceleration via discrete-time differentiation of the relative position. An airborne gravity system also depends on position to compute the normal gravity vector and the relationship between the l-frame and e-frame. The latter computations depend on the estimated absolute position. Recall from Section 2.2 that if the errors in position can be kept below 50 cm (1-σ), they will cause negligible degradation of the performance of an airborne gravity system. Although at first thought, meeting this accuracy requirement may seem like a trivial task when using DGPS, it cannot always be guaranteed in practice.

3.1.1 A comparative study of current DGPS software

Two studies that were recently carried out within the AGEM project of the GEOIDE Network are relevant to this topic. The first is discussed in Bruton et al. (2000b) and shows that an accuracy of 50 cm is not a certainty. The study involved the comparison of

eight industrially relevant DGPS software packages that make up a representative cross-section of the DGPS engines that are currently available in industry and in academia. Each participant was given exactly the same airborne data sets for processing and the only instruction was to provide the best possible solution. For this reason, the comparison of results made therein represents a realistic evaluation of the capabilities and consistency of currently available software. Because the accuracy of this type of kinematic positioning is often quoted to be at or below the decimetre level, it was somewhat surprising to find that the solutions calculated using each of the packages differed from each other by as much as 2 metres. Unfortunately, no independent truth was available to evaluate the solutions on an absolute scale, meaning that the reasons for the large differences are not known in all cases. It is concluded that the results of that initial study do not establish the accuracy of airborne DGPS, but instead outline how inaccurate it can be. Errors in position at the metre-level cause enormous problems for any high performance airborne mapping or airborne gravity system. This motivates the need for the further study reported in this chapter.

3.1.2 A study of the accuracy of airborne DGPS positioning

The work reported in this chapter responds to the above need by addressing the accuracy of kinematic carrier phase DGPS in an airborne environment. It describes a new study that was carried out by the author within the AGEM project. By using a single software package and various processing strategies, each of the DGPS error sources that were discussed in Section 2.3.3 are isolated and roughly quantified using examples drawn from available data sets. Included are the effects of receiver noise and multipath (on the ground and in the aircraft environment), as well as the influence of the unmodeled portions of the ionosphere, troposphere and broadcast ephemerides. The effects of using dual frequency measurements for ambiguity resolution and ionospheric-free processing are also discussed and their importance is outlined in light of various realistic airborne conditions. Although brief definitions of the error sources are given whenever possible, more comprehensive information about the GPS and its error sources can be found in related publications. Many details can be found in Parkinson and Spilker (1996) and a comprehensive look at the use of the system in geodesy can be found in Teunissen and

Kleusberg (1998a). A good study of the characteristics of DGPS errors is described in Raquet (1998) and other recent analyses that treat various aspects of the relative and absolute accuracy of airborne DGPS positioning can be found for example in Cannon et. al (1992), Shi (1994), Tiemeyer et al. (1994) and Han et al. (1998).

3.2 On the DGPS error budget for airborne positioning

Because all of the analyses contained herein are empirical in nature and therefore depend on the data that is used, the data sets have been carefully chosen. They permit the isolation and emphasis of the different error sources, paying particular attention to the atmospheric effects. The first data set was collected on September 10 1996, during an airborne gravity campaign that was carried out by the University of Calgary. It is described in Appendix A.1 and referred to there as the Kananaskis field campaign. It is used here because it was flown at night and at a high altitude, implying that the effect of the ionosphere on the data is small and the effect of the differential troposphere is relatively large. The time of the flight was between 00:00 and 06:00 local time and the average ellipsoidal flying height was 4357 m. The Trimble 4000 SSI receiver on the aircraft and those at the master stations in Banff and Invermere will be employed in this chapter.

The second data set was collected on June 6th 1998 and is described in Appendix A.2. It was collected during an airborne gravity survey carried out over Greenland and is used herein to complement the Kananaskis data set because it was flown in the afternoon (between 15:00 and 18:00 local time), during a period of high ionospheric activity. The average flying height above the master station was only 300 m, implying that the effect of the differential troposphere is small. The two Trimble SSI receivers on the aircraft and the one at the Aasiaat station are used. The receivers on the aircraft each used their own antennas that were mounted on the fuselage and separated by over 3 m.

The overall methodology used in the following takes advantage of a number of operational and data processing strategies to evaluate the effect that each portion of the error budget has on the estimated position; it studies each of the errors that were introduced in equation 2.14. In order to control the experiment, the use of a single software package was necessary and Version 6.02 of GrafNav (provided to the author by

Waypoint Consulting Ltd.) was selected for two reasons. The first reason is that it permits a large number of variations in the data processing strategy. (Although most of the same conclusions were reached using KINGSPAD, the latter is not currently capable of carrying out an ionospheric-free linear combination and cannot use precise ephemerides, two components of the analyses that are useful for demonstrating certain points). The second reason is that GrafNav is a commercial DGPS software package that enjoys wide use and respect in the fields of airborne navigation and mapping.

The specific methodologies used for the evaluation and characterization of each error source are introduced in the following, as necessary for an understanding of the experiments.

3.2.1 Errors due to the broadcast ephemerides

Because positions calculated using GPS are based on the range measured between the user and each satellite, the assumed location of the satellites is of fundamental importance. The ephemeris information that is broadcast by the satellites themselves is currently the only means of obtaining this information in real time. On the other hand, more precise orbital information is available from a number of organizations between one day and four weeks after a survey. Errors in the positions of the satellites will result in an error, $\nabla\Delta\,\delta\rho$, in the double difference measurement that will be small for short baselines, but increasingly problematic as the baseline length increases. (Note that the error sources are often referred to in this chapter by using the notation that was introduced in Section 2.3.3.)

The effect of errors in the broadcast ephemeris can be evaluated in post-mission by using it to compute a position solution and comparing the result to one obtained using precise ephemeris information. Figure 3.1 shows the magnitude of this difference over the area flown for the Kananaskis field campaign, viewed from the north-east. The surface plot is generated from the time series of position differences for each flight line. As expected, the differential error grows as a function of distance from the master station. It is clear from the figure that the errors due to the broadcast information can be significant, reaching values of over 5 cm for distances of 120 km. The time dependence of the

broadcast ephemeris error can be observed by noting the saw-tooth pattern that shows up in Figure 3.1 when going from east to west. The pattern is the result of an update of the broadcast ephemeris that was made as the aircraft flew south along a line about 40 km west of the master station. Given the current accuracy of the GPS broadcast ephemeris, the range of position errors in Figure 3.1 agrees roughly with the estimated worst-case user position error of 5-10 cm for each 10 m orbital error over a baseline of 100 km that is given in Parkinson and Enge (1996).

Because the errors in position due to broadcast ephemerides are low-frequency in nature and smaller than 10 cm for typical baselines, they are unlikely to pose a threat to the performance of an airborne gravity system. They may make up a significant portion of the error budget of an accurate airborne mapping system, however, a topic that will be discussed further in Section 3.4.

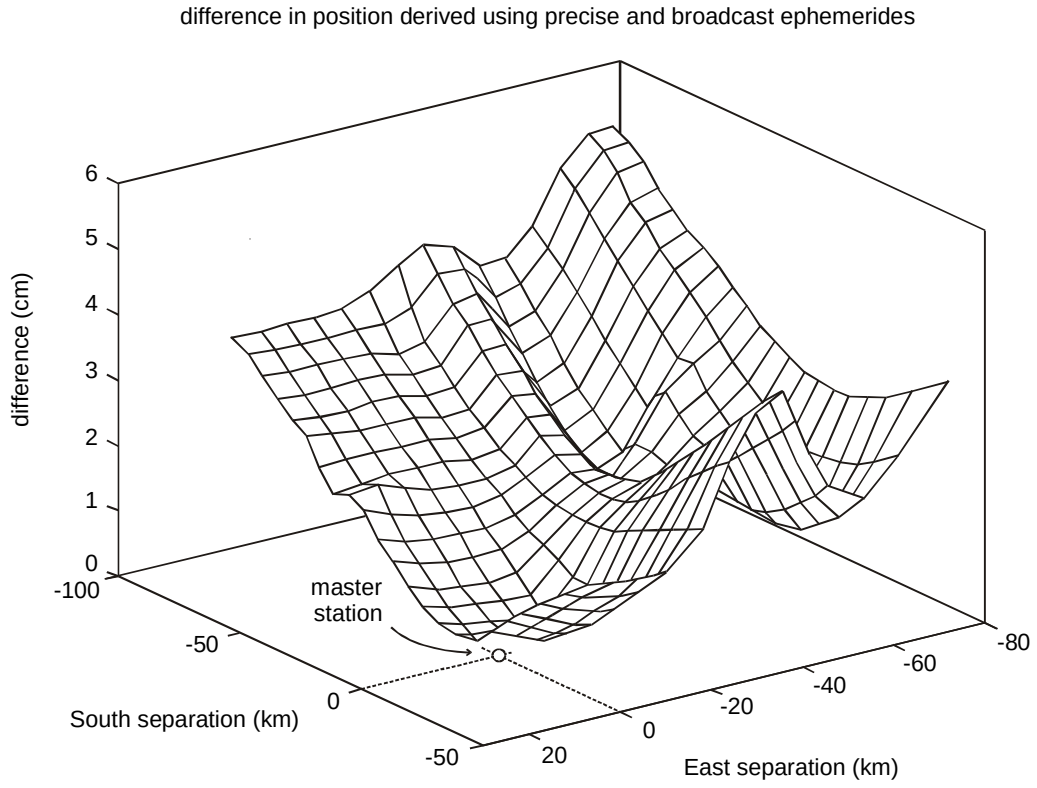

Figure 3.1 Error in position due to the broadcast ephemeris in a 100 by 150 km area

3.2.2 The importance of resolving the integer ambiguity term

The carrier phase observable is a measure of the difference between the phase of the carrier signal received from the satellite and a nominal carrier generated by the GPS receiver. These are related to each other by the time it takes for the signal to propagate from the satellite to the receiver, and the observable is therefore directly related to the range between them. Unfortunately, a GPS receiver only measures the fractional part of this phase difference, meaning that the range between the user and a given satellite cannot be determined unless the initial integer number of cycles, N, is known. Further, this ambiguity remains after differencing and must be estimated along with the other unknown parameters.

While this study puts no emphasis on how the ambiguity term, $\nabla\Delta N$, is estimated, it is highly relevant to consider the importance of using a valid estimate of this ambiguity. The goals here are to relate the observed importance of dual frequency data for ambiguity resolution and to demonstrate how much a solution based on real-valued (or float) ambiguity estimates can deviate from the corresponding fixed integer solution in typical applications. Experiences drawn from processing the Kananaskis data set provide a very useful example of this. Because of the length of the baselines flown during that campaign, it is not possible to reliably estimate the integer ambiguity terms using only measurements of the L1 carrier phase. In other words, only a float estimate of those terms is possible when using single frequency data. Integer estimates can be obtained, however, by forming the widelane observable through a linear combination of the L1 and L2 phase measurements. Teunissen (1998) can be consulted for more details about the widelane observable and its use in ambiguity resolution.

Generally speaking, a float solution should agree with the corresponding fixed solution at or below the 5-10 cm level for kinematic surveys over short and medium baselines, Lachapelle (1998). Consider Figure 3.2, however, that shows how the float solution diverges considerably more than this for the Kananaskis data, yielding errors in position that are greater than half a metre. Figure 3.2a shows the difference between the single frequency (L1) float solution and the corresponding (widelane) fixed solution and Figure 3.2b shows the horizontal distance separating the master station and the aircraft for the

same time period. It is interesting and important to note that there is little growth in the error over the first 6000 s, the time period during which the aircraft passes close to the master station for every flight line. As the shortest separation distance increases however, so does the error in position. It is likely that this occurs because, as time passes, the float estimates of the ambiguity absorb errors that are not correlated between the master and remote stations.

This is an important observation because although reliable integer ambiguity estimation can be achieved for the Kananaskis data set, conditions may not permit it in all cases. In such situations, a float solution may be the only alternative (it is seen later that this is the case for the Greenland data set), perhaps having a considerable influence on the error budget of an airborne mapping system. Even in airborne gravimetry where the requirements are often much more forgiving, errors of this magnitude might cause errors in the output that reach significant levels.

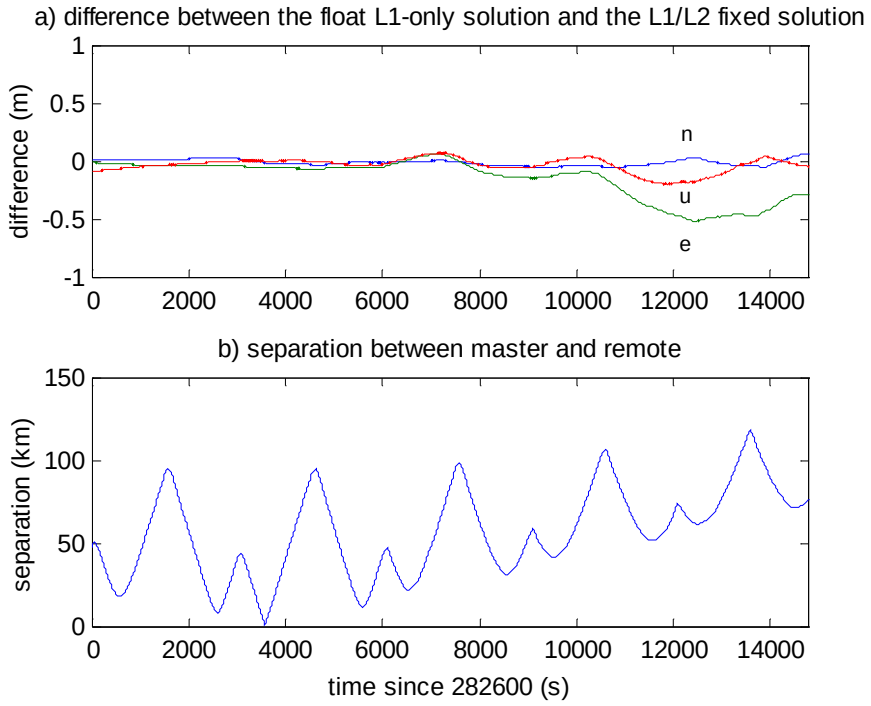

Figure 3.2 The accuracy of a float solution and the baseline length

3.2.3 Errors owing to the neutral atmosphere

It is well known that, as the GPS signals propagate through the atmosphere on their way to receivers on the Earth, their speeds are decreased by the variability of the refractive indices in the troposphere, the tropopause and the stratosphere (the portions of the atmosphere below 10 km, between 10 and 16 km and between 16 and 50 km above the surface). The combined effect that this variability has on the signals is usually referred to as the tropospheric delay. When derived using (code and) carrier phase measurements, it causes the ranges between the satellites and the user to appear between 2 and 25 m longer than the true distance, Lachapelle (1998). For the range measured to each satellite, this delay is usually modeled in terms of the wet and dry components, which account for about 20 and 80 % of the total delay, respectively. The models are applied in the zenith and projected onto the line of sight using an elevation mapping function, Lachapelle (1998). The key parameters in most models are pressure, temperature and relative humidity. Even after differencing, residual (i.e. unmodeled) portions of the delay, $\nabla\Delta T$, can cause significant errors in the estimated position that show up primarily in the height component. Although the accuracy of available models varies, the dry portion can be modeled to the centimetre or even millimetre level, while errors in modeling the wet portion can reach the decimetre level and continue to be a challenge for GPS users, see Mendes and Langley (1998) for details. Airborne mapping applications are especially susceptible to this error because of the large vertical separation between the master station and the aircraft and because it often does not show up in typical measures used for quality control (such as adjustment residuals). For example, note that large parts of the observed differences between solutions in Bruton et al. (2000b) were attributed to different treatments of the tropospheric effects. The goal of the next few paragraphs is to use a data set with a very large vertical separation to demonstrate the magnitude of the effect of the tropospheric delay on the GPS measurements and the estimated positions and to roughly quantify the accuracy with which the delay can be estimated by comparing the performance of different models. Because this is done empirically using a single data set, the result should be interpreted as no more than an example of the effect. Similar recent examples that have made efforts to highlight the importance of the troposphere can

be found in Tiemeyer et al. (1994) and Shi (1994). A thorough investigation of the absolute accuracy of many different tropospheric models can be found in Mendes (1998) and for a more rigorous discussion of the topic in view of airborne navigation, Mendes and Langley (1998) should be consulted. Comparisons made later, in Section 3.3, also treat the absolute accuracy of the tropospheric delay in more detail.

Recall that the Kananaskis data set was chosen for use in this work because it was collected at a very large flying height and is therefore subject to large differential tropospheric delays (the average vertical separation between the Banff master station and the aircraft was 3035 m during the survey). Consider Figure 3.3 that is borrowed from the more detailed discussion of this topic found in Bruton et al. (2000b). The data used to generate it was derived by T. Beran and R. Langley of the University of New Brunswick. Figure 3.3a shows the magnitude of the tropospheric delay as a function of height difference in the zenith and at an elevation of 10 degrees. Figure 3.3b shows the agreement between two models (Hopfield and Saastamoinen) for estimating that delay. Together these verify that, for a large vertical master-remote separation, the total relative delay can reach about 7 m for low elevation satellites and that, at worst, two different models agree to about 10 cm in their estimation of that delay.

This result does not necessarily quantify the absolute error in the delay estimated by models commonly used in differential positioning (both could be biased by some amount from the true delay), but it does provide a rough idea of how inaccurate at least one of them could be. For example, two packages each employing one of these models could obviously not both be correct. It is important to consider the impact of such an error on the estimated position. The vertical position error resulting from unmodeled tropospheric delays is approximately equal to the magnitude of the largest unmodeled delay in the system (and will usually correspond to periods dominated by low elevation satellites), R. Langley, personal communication. Given that, a height error of 10 cm might be expected in this case.

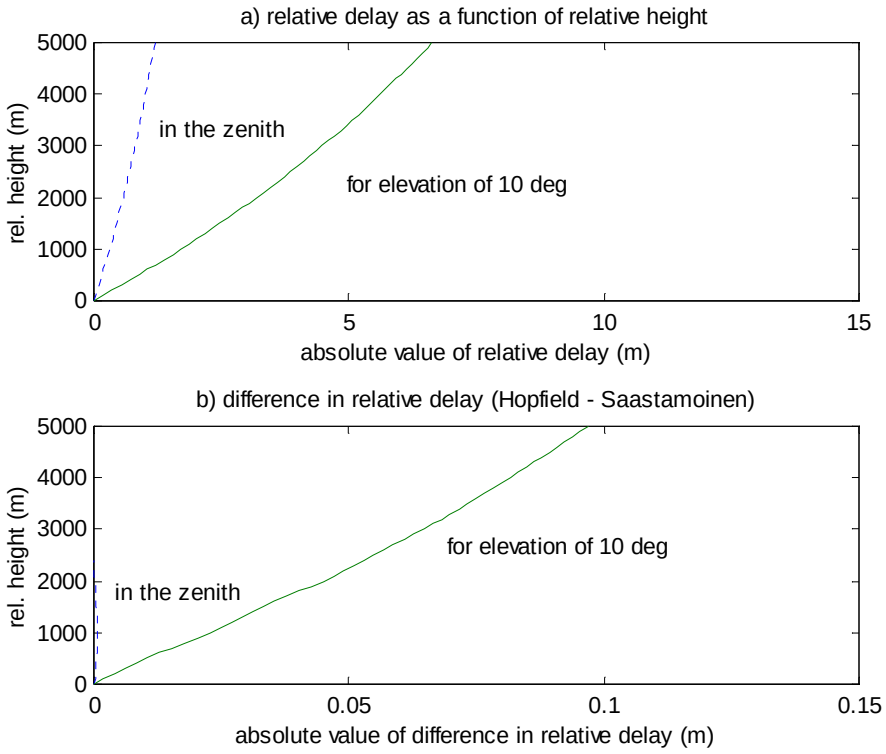

Figure 3.3 Relative delay as a function of height difference

The graphical data in Figure 3.4 is generated to show the approximate effect of the measurement domain quantities (shown in Figure 3.3) on the estimated position. Figure 3.4a shows the effect of the troposphere on the estimate of aircraft height during the period of ascent for the Kananaskis data. It is generated by comparing the position solutions obtained by processing with and without the use of a tropospheric model. It is also borrowed from Bruton et al. (2000b) and is based on data provided by A. Simsky of Sander Geophysics. It shows that the error in position, corresponding to the tropospheric delay, reaches nearly 1.5 metres for this data set. Finally, it is of interest to observe the large differences in position for solutions that have been generated using different tropospheric models for the same data. Figure 3.4b is generated using solutions derived using the two different tropospheric models available in GrafNav (Saastamoinen and Black in this case). The figure shows that the difference in height estimated using these models reaches nearly 15 cm as the aircraft ascends.

Although comparing the results obtained from these two models can be used to roughly quantify the residual error, it is important to realize that other factors, such as the meteorological parameters that feed the model, may play an important role in the absolute accuracy of such a model. While not directly indicating the accuracy of either model on its own, this approach does indicate how inaccurate at least one of them might be.

On its own, a bias-like error with a magnitude of 10-15 cm is unlikely to be a problem in airborne gravimetry. It may be problematic, however, for high-end navigation and mapping systems, in which case, it might be important to make meteorological measurements on the ground or in the air throughout the survey. Errors due to the neutral atmosphere are discussed further in Section 3.3.

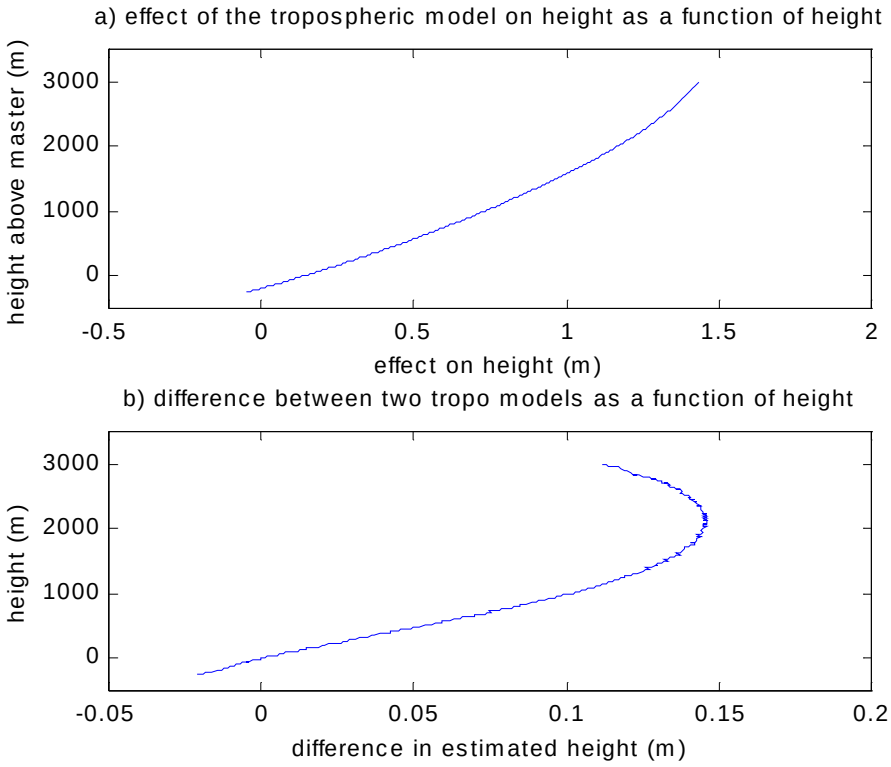

Figure 3.4 Difference in position between two tropospheric models during the ascent

3.2.4 An estimate of the effect of the ionosphere

The ionosphere (the portion of the atmosphere between 50 and about 1500 km above the surface of the Earth) has a significant effect on the propagation of the GPS signals; the

carrier phase is advanced by an amount proportional to the density of electrons and causes the measured ranges to appear shorter than they really are. Fortunately, the ionosphere is dispersive, meaning that the effect it has on a signal is a function of its frequency. In turn, this means that the effect of the ionosphere can be estimated by using a combination of the carrier phase observables on the L1 and L2 frequencies.

The goal of Section 3.2.4 is to demonstrate the potentially devastating effects of the ionosphere by employing an example where it is particularly active. Recall that the Greenland data set was selected for this study because it was known to have been collected during a period of high ionospheric activity and because the effect of the troposphere is relatively small. Because this ionospheric activity only permitted the reliable estimation of the integer ambiguities for a portion of this data set, the following analysis only treats float solutions.

Figure 3.5a shows the difference between solutions obtained using both the L1 and ionospheric-free observable (the latter is one of the possible linear combinations of the L1 and L2 phase measurements). As discussed in Klobuchar (1996), the ionospheric-free combination takes advantage of the dispersive nature of the ionosphere to remove all of its first-order effects, which account for over 99.9% of the total effect. See Teunissen and Kleusberg (1998b) for details about how the ionospheric-free observable is formed. Figure 3.5a therefore represents the error in position caused by first-order effects of the ionosphere, $\nabla\Delta I$. Figure 3.5b shows the horizontal distance separating the master station and the aircraft for the same time period. The effect of the ionosphere clearly increases as the baseline length increases, reaching values greater than one metre in each component for distances greater than 100 km. It is also interesting to notice that, although the effect of the ionosphere is reduced as the aircraft returns to the master station (after about 6500 s), the position solution is heavily biased, especially in the height component. A comparison of each solution to that obtained from a static period at the end of the survey confirms that the ionospheric-free solution is the correct one.

Based on the evidence presented herein, there is no doubt that under conditions of high ionospheric activity, such as those observed during collection of the Greenland data set, a single frequency approach is not sufficiently accurate. If errors like those observed in

Figure 3.5a are not corrected for, by using a dual frequency approach, the results could be disastrous for any mapping system.

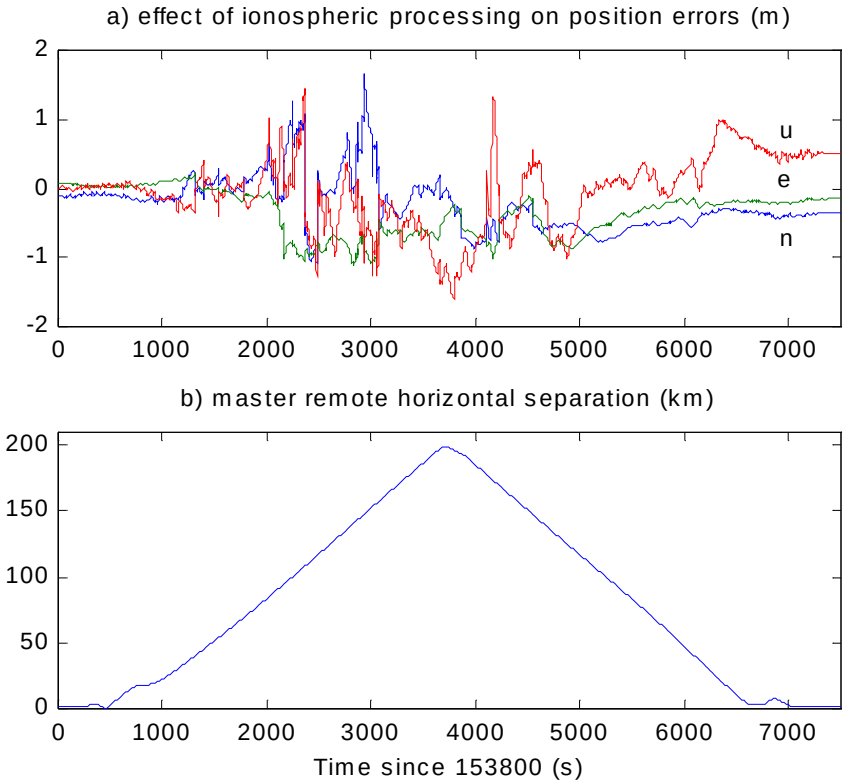

Figure 3.5 Errors caused by first-order ionospheric effects, and the baseline length

3.2.5 The effects of multipath and receiver measurement noise

The multipath effect is the result of a GPS signal arriving at an antenna from more than one direction (due to reflection) and receiver measurement noise is defined as the noise resulting from the process of taking the carrier phase measurements. Because it is extremely difficult to develop appropriate models of multipath and noise, they usually dictate limits on the accuracy of carrier phase positioning, especially in kinematic mode. Carrier phase multipath results in low and medium-frequency errors in the estimated range to a satellite. Such errors can have a maximum phase error of a quarter of a wavelength (i.e. about 5 cm for the carrier on L1). See Ray (2000) and the references made therein for good discussions of carrier phase multipath effects. GPS receiver noise

is usually a broad-band, white noise sequence. Its amplitude is typically at or below the millimeter-level for each measured range and increases for each level of differencing performed, Lachapelle (1998). It is also a function of the bandwidth of the tracking loop of the receiver, and therefore the dynamic of the vehicle.

Noise and multipath on the ground

While multipath and measurement noise each have an effect at the master station and at the airplane, it is relatively straightforward to quantify this effect for an antenna-receiver configuration on the ground. As an example, consider Figure 3.6 that was generated using the data on the L1 carrier coming from a triplet of Trimble 4000 receivers and two antennas at the University of Calgary. The antennas were arranged 754 m apart, one on the roof of the Engineering building, which is considered to be a medium multipath environment (see Ray (2000)), and one in a nearby field where the multipath is considered to be low. The signal from the antenna on the roof of the Engineering building was split to two of the receivers while the third receiver was used at the station in the field. Because the true range is known, the ambiguities are easily fixed and because residual ephemeris and atmospheric errors are negligible over such a short baseline, Figure 3.6a is a plot of the combined effects of double difference noise and multipath ($\nabla\Delta\, m_{\Phi} + \nabla\Delta\, \varepsilon_{\Phi}$). As shown, the combined magnitude of the error in position caused by the multipath and noise from the signals on L1 is at the level of 0.4 cm (1-σ) with a maximum error of about 2.0 cm. Figure 3.6b is a plot of the error in position over a zero-baseline (i.e. all errors except noise cancel completely) and shows that receiver noise on the L1 carrier only contributes an error in position at around the 1 mm-level (1-σ) with a maximum contribution of 3 mm for the same time period. Processing was done in kinematic mode in all cases. Although these analyses are empirical in nature and may vary with many factors (such as the environment around the antenna, receiver type, satellite elevation and strength of the geometry), they agree well with observations made in the more statistically meaningful studies carried out by Raquet (1998) and Bona and Tiberius (2000).

Because the ionospheric-free observable is used in many places in Chapters 3 and 5, it is important to roughly quantify the level of noise associated with it. As mentioned above, because noise should not be correlated from signal to signal or on either carrier frequency, it increases with each level of differencing done. As shown in Shi (1994), the expected noise level for the ionospheric-free combination, $\sigma_{\text{iono-free}}$, can be calculated as a function of the noise level on the L1 carrier, σ_{L1}, through the laws of error propagation. It yields the following relationship:

$$\sigma_{\text{iono-free}} = 4.09\sigma_{\text{L1}}. \tag{3.1}$$

In other words, while the ionospheric-free combination removes all first-order effects of the ionosphere, it increases the noise level by about a factor of 4 on each measured range. After differencing, the noise term, $\nabla\Delta\,\varepsilon_{\Phi}$, would therefore be increased by a factor of eight. Figure 3.6c was generated in the same way as Figure 3.6b except that the ionospheric-free observable was used instead of the L1 carrier phase observable. It shows that the same increase in noise can be expected in the estimated position.

Noise and multipath at the aircraft

Estimating the effect of multipath in the airborne environment is more difficult than in the static case because the true position of the antenna is not known. It is possible to roughly quantify it, however, by knowing the true distance between the two antennas that were mounted on the aircraft during collection of the Greenland data set. Figure 3.7 is a plot of the error in the length of that baseline due to the combined effect of receiver noise and multipath at each of the antennas on the aircraft. It was generated by taking the difference between the two L1-only position solutions derived using one master station and each of the antennas, i.e. there were no constraints placed on either solution. The ambiguities were reliably fixed for both solutions and the satellite constellation did not change during the period shown (i.e. no satellites rose or fell from view). Also, because the antennas are closely spaced, the atmospheric and orbital effects are the same, meaning again that only the effects of the terms $\nabla\Delta\,m_{\Phi}$ and $\nabla\Delta\,\varepsilon_{\Phi}$ remain. This method assumes that the multipath is uncorrelated from one antenna to the other.

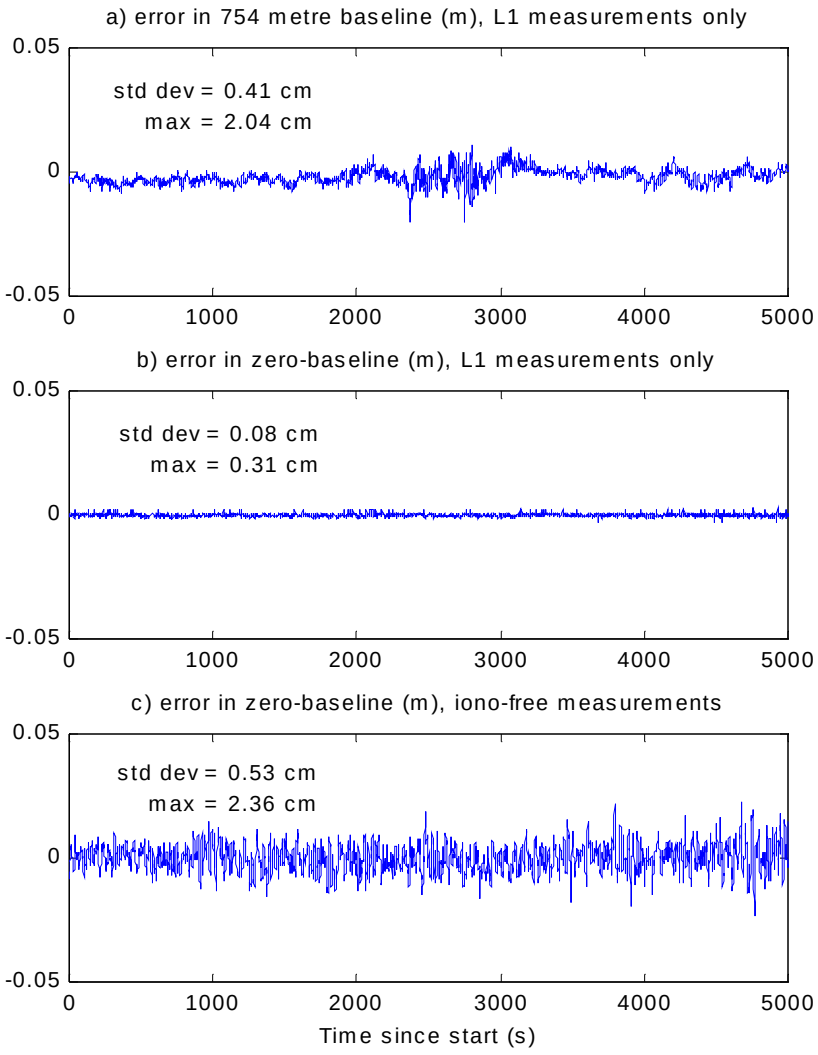

Figure 3.6 The errors caused by multipath and measurement noise on the ground

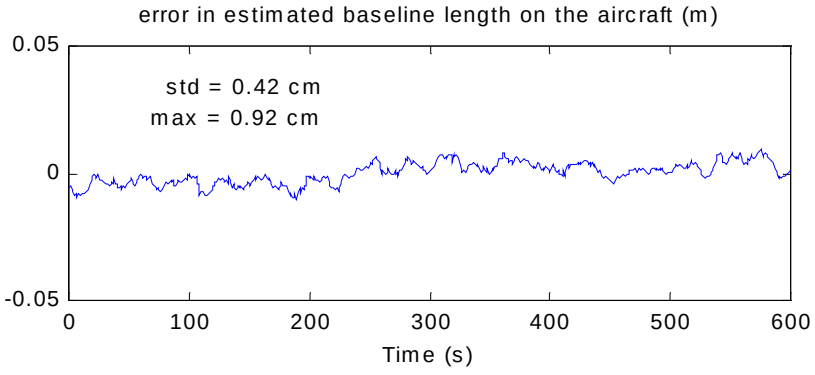

Figure 3.7 The combined effect of noise and multipath at the aircraft

As shown, the combined effect of noise and multipath for both antennas is at the level of roughly 0.4 cm (1-σ). Assuming that the effect due to multipath is equal in magnitude at both antennas but uncorrelated between them, this corresponds to a value of multipath and noise at each airborne antenna of about 0.3 cm (1-σ). This can be compared to the value of 0.4 cm (1-σ) that was observed above for multipath and noise on the ground. Visual inspection of Figure 3.6a and Figure 3.7 implies that medium to high-frequency multipath is averaged out, while much of the low-frequency content remains. Notice that the plots use different time scales. This supports the common perception that the changing orientation of the vehicle in a moving environment tends to average multipath. A complete study of this would include significantly longer data sets and a comparison of the various multipath mitigation technologies that are employed in today's receivers.

3.2.6 Errors due to changes in geometry

The reliability of a DGPS solution depends largely on the geometry implied by the satellite constellation; two periods of different geometry can yield different position solutions under otherwise equivalent conditions. Changes in geometry may result from the rising and setting of satellites, the obstruction of signals and in the case of an aircraft, from the banking that occurs during turns and other maneuvers. Even more dangerous is the situation where the elevation mask coupled with vehicle maneuvers cause drastic changes in the observed geometry. The goal of this section is to demonstrate some of these problems, again drawing on examples.

Up to this point, all computations have been performed using an elevation mask of 10^o. For the Kananaskis data set, increasing the elevation mask from 10^o to 15^o makes it impossible to reliably maintain estimates of the ambiguities through the turns, meaning that a solution that uses float estimates becomes the only option for at least some periods of time. Figure 3.8a is a plot of the difference between solutions obtained using elevation masks of 10^o and 15^o for the first data set (i.e. the difference between the solution obtained using a mask of 15^o and the fixed L1/L2 solution that was discussed in Section 3.2.3). Figure 3.8b and Figure 3.8c show the observed PDOP and number of satellites in each case, for the same time period.

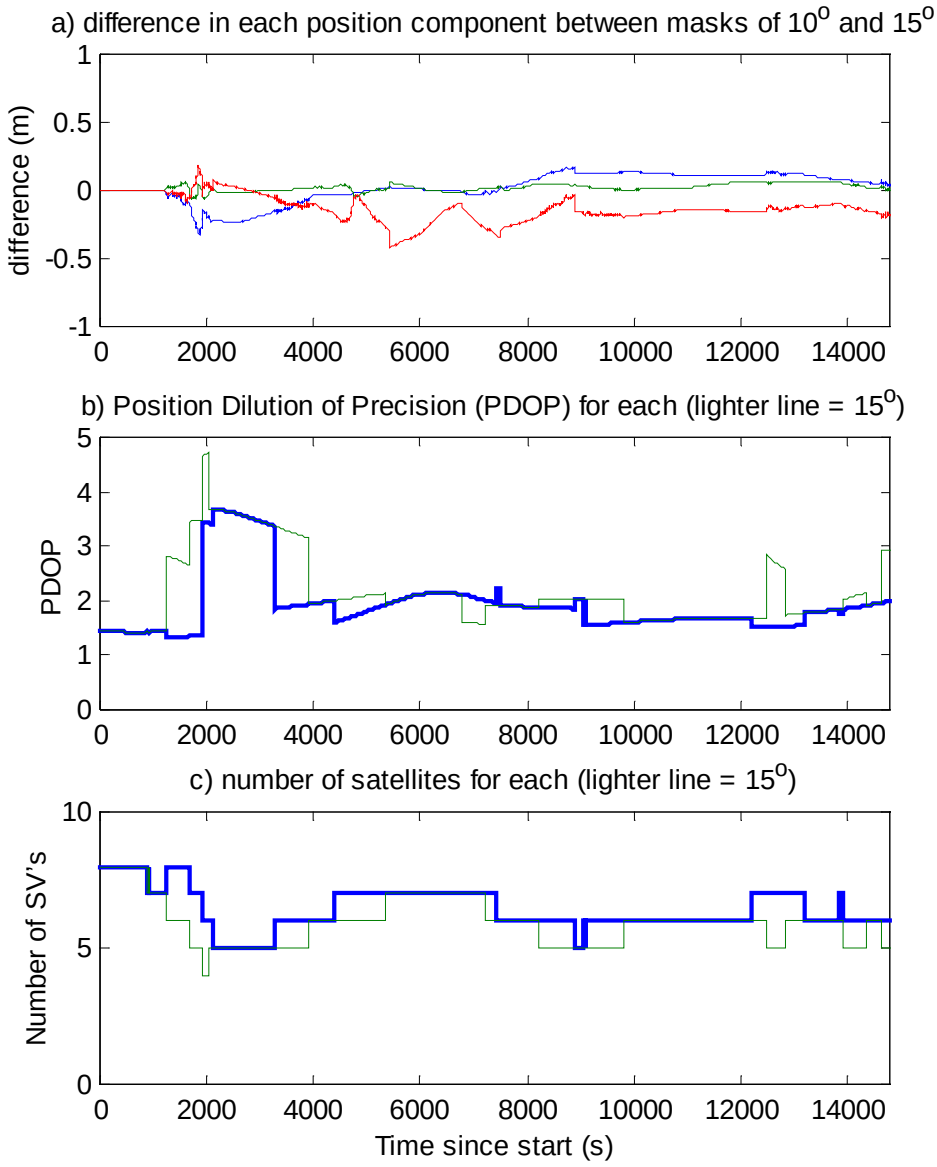

Figure 3.8 The effect of changing the elevation mask on position

Clearly, the observed satellite geometry makes a large difference following the turn that occurs at approximately 1200 s. Increasing the mask by an extra 5° forces the exclusion of satellites from the solution during that turn, bringing the total number to 4 for a short period. Reliable floating estimates of the ambiguities clearly become more difficult to achieve as the geometry worsens. Although this simulation is artificial, consider the fact

that a user would not have known any differently if those satellites had not been available in the first place due to poor health or other reasons.

Generally speaking, more satellites in more portions of the sky will imply a better geometry and a better estimate of position. Consider Figure 3.9, however, that shows an example in which this is not true. It shows the error in the same baseline that was shown in Figure 3.7 except that it was calculated by allowing one of the solutions to include an extra (low elevation) satellite for the period after 190 s. Recall that the true length of the baseline is known. Although the ambiguities remain fixed, the observed geometry improves as a result of the extra satellite. The level of multipath appears to remain roughly the same. It is interesting, however, that the estimated length of the baseline becomes biased by about 1 cm. This increases the maximum error from under 1 cm to nearly 2 cm, implying that it may not always be best to include a satellite in an effort to improve the geometry. In this case, the observed bias may come from residual tropospheric delay on the low elevation satellite which would give strong support for methods such as those outlined in Hartinger and Brunner (1999) that weight the carrier phase measurements according to the signal-to-noise ratio, the elevation of the satellite or the time since lock.

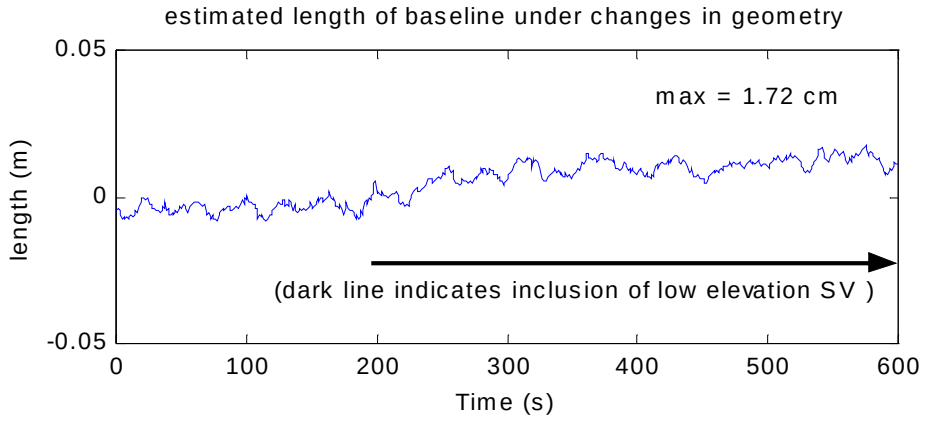

Figure 3.9 The effect of the inclusion of a low elevation satellite

3.2.7 Other sources of error

For completeness, it is important to point out that there are a variety of error sources that have not been treated thus far and that can be significant, depending on the circumstances. Among them are errors due to:

- instrumental delays in the receivers and satellites (see Teunissen and Kleusberg (1998b)),
- errors in precise orbits, especially over very long baselines (see Beutler (1998)),
- remaining tropospheric modeling errors, e.g. their dependency on instantaneous environmental conditions (see Brunner (1997)),
- ionospheric effects of second-order and greater (see Klobuchar (1996)),
- the non-white behavior of stochastically modeled parameters and their effect on the estimation process (see Wang (1999) and Tiberius et al. (1999)) and
- latency effects in real-time surveys (which are likely to be less crucial since SA was turned off).

While it can be argued that many of these are negligible under normal circumstances, their effect should always be considered for high-accuracy applications.

3.3 Checking the consistency of various solutions

There are two checks for consistency that are often carried out to verify the quality of airborne positioning results. The first is included here because it provides further insight into the quality of the tropospheric modeling being carried out and the second because it comments on the accuracy of the processing and kinematic models being used.

Figure 3.10a was derived by comparing the position solutions obtained for the Kananaskis data set using two of the available master stations (Banff and Invermere). Intuitively, it makes sense that agreement of these with each other is a necessary condition for claiming a given level of accuracy. Both are fixed integer ambiguity solutions that make use of the Saastamoinen tropospheric model and the same precise ephemeris information for the whole time period. Notice by observation of Figure 3.2b and Figure 3.10b that the agreement is below the level of 3 cm in each component when the distance of the aircraft from both master stations is below 75 km. Also notice that the

difference grows as a function of the baseline length, reaching a peak value of 18 cm in the vertical component and 6 cm in the horizontal components. This peak occurs when the distance of the aircraft from the Invermere station reaches just over 150 km and after it has not been closer than about 50 km to it for over 1.5 hours. Because the error sequence is clearly spatially correlated (and because precise ephemerides have been used), the large amplitude differences are most likely due to residual atmospheric effects. Further, because the data set was collected during the nighttime and at a very large flying height, most of the effect is probably due to unmodeled tropospheric delay (note that the Invermere master station is approximately 465 m lower than the Banff master station and 90 km away from it in the horizontal). A small component of the difference is also probably due to different multipath experienced at the master stations.

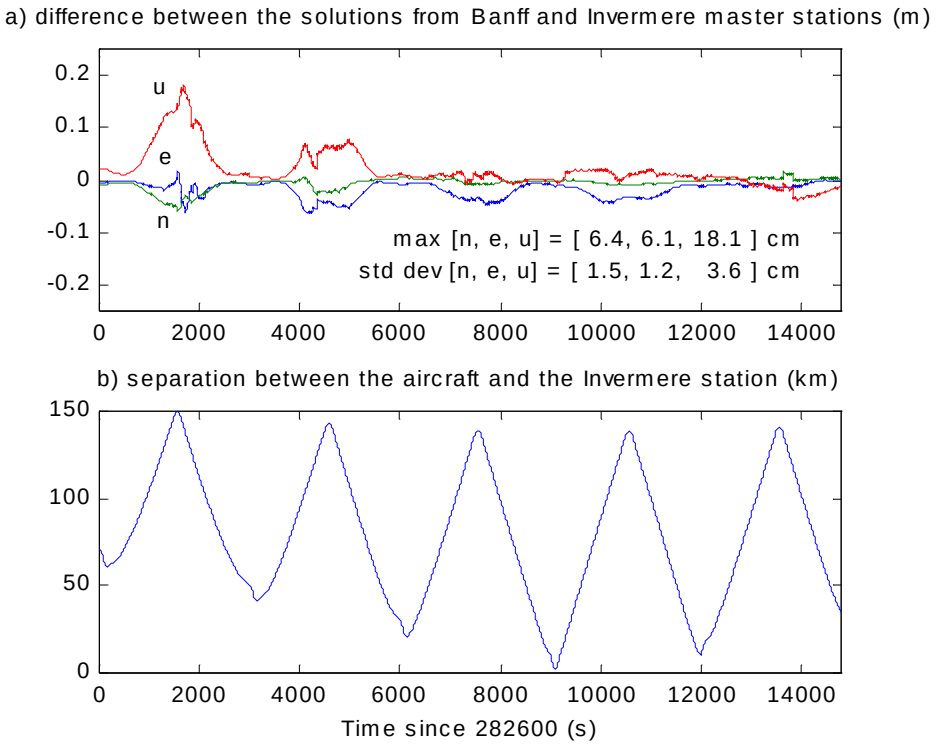

Figure 3.10 The difference between solutions from different master stations

Finally, the two solutions obtained by processing exactly the same data forward and backwards in time (using the Banff master station) are compared in Figure 3.11. It shows a high level of consistency, with a maximum difference of 6 cm in the height and

standard deviations of less than 2.5 cm in all three components. For completeness, it should be mentioned that this is less consistent than a similar result presented in Shi (1994). It is assumed that this is because there were only two instances in the latter where the satellite constellation was reported to change (i.e. satellites rising or falling from view). Consistency between solutions is very challenging to achieve when satellites enter or leave the solution; these are opposite occurrences in forward time than in reverse time. Compare the solid line in Figure 3.8c to Figure 3.11 to observe the correlation between these events and the differences between the forward and backward solutions.

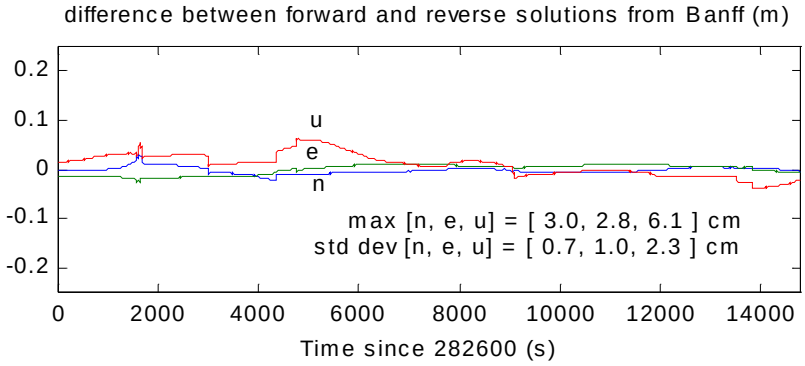

Figure 3.11 The difference between forward and reverse solutions

3.4 The effect of positioning errors on airborne mapping systems

Examples of the residual errors affecting the accuracy of airborne DGPS positioning have been isolated and presented in an effort to demonstrate the effect that each can have under realistic conditions. It was seen that an absolute accuracy of 50 cm is not always achieved. For long-duration surveys with baselines reaching lengths of 150 km, it was seen that:

- total multipath and noise seem to cause errors no larger than about 3 cm,
- the kinematic model and processing algorithm are consistent to at worst 6 cm (independent of baseline, but dependent on changes in satellite constellation),
- absolute errors in the order of 10-15 cm due to the troposphere seem possible for long baselines,

- a float solution can deviate by 50 cm from the corresponding fixed solution (also for long baselines),
- an active ionosphere can cause metre-level errors if not treated properly,
- poor geometry can result in errors at least as large as 50 cm, and
- errors in broadcast orbital information can cause errors that are greater than 5 cm.

While it can be argued that the latter three of these can be minimized (by various logistical and processing strategies), the first three represent limits to the accuracy of current airborne positioning techniques. In cases where a widelane approach is not possible (such as under a very active ionosphere), the fourth point above may also pose a limitation. Given that these represent maximum errors, however, it is reasonable to expect to obtain the position of an aircraft with a standard deviation (1-σ) of under 50 cm. This confirms that **DGPS is capable of providing the positioning accuracy required for use in airborne gravimetry.**

As for all airborne positioning applications, factors such as the selection of master stations with low multipath, keeping aircraft bank angles to a minimum during turns, using precise orbital information, flying as low as possible to minimize tropospheric effects and especially using dual frequency receivers will all help to ensure that this level of accuracy is maintained under adverse or unexpected conditions.

For completeness, consider the general positioning requirements that are given in Table 3.1 for a number of airborne georeferencing systems. They are from Skaloud (1999) and are given as function of the mapping sensor being employed. Based on the observations made in this chapter, several broad conclusions can be made regarding the ability of DGPS positioning to meet them. For large scale mapping operations (where 5-10 cm are needed), there appear to be two prerequisites: the length of baselines must remain shorter than 75 km and the aircraft must fly close to the master station on a regular basis (e.g. at least once per hour) to avoid the growth of errors that are not sufficiently spatially correlated. It is reasonable to expect DGPS to meet the accuracy requirements for the other applications listed in Table 3.1 as long as the baselines are shorter than 150 km.

Finally, despite these optimistic conclusions, the reader is reminded of the large differences between the eight DGPS packages that were mentioned earlier and reported in

Bruton et al. (2000b). Although a group of the solutions agreed with each other at the level of 10 cm, a number of them differed from that group by more than a metre. This implies that a considerable amount of caution should be exercised when developing, evaluating and using a DGPS software package for airborne applications.

For further evaluating the absolute accuracy of DGPS, the author recommends using a data set in which a precise reference is available. This might include one where either LIDAR or photogrammetry are available as an accurate and independent reference. If appropriate data sets become available, the analyses should also be extended to include various GPS receiver types.

Table 3.1 Positioning accuracy requirements for airborne mapping

Type of sensor	Required Position (cm)
Aerial camera, scale > 1:2000	5-10
Aerial camera, scale < 1:5000	75 – 100
CCD camera or scanner	25 – 100
IFSAR	20 – 200

4 Differentiation for High Precision Velocity and Acceleration Determination

In this chapter, processes of differentiating discrete-time signals are discussed, with emphasis on determining velocity and acceleration from measurements made by the GPS. There are two principal objectives:

a) to comprehensively review possible methods of carrying out the differentiation process, expressing each in a comparable and uniform way (that is, as finite impulse response (FIR) filters), characterizing each method in the frequency domain, and

b) to employ real GPS data to demonstrate the appropriate use of each method of differentiation in both the measurement and position domains.

Operations are carried out on the GPS signals at two stages in this chapter, depending on the application at hand. In the first stage, operations are carried out on the GPS measurements, prior to their use in a GPS software package (that employs a linear system like that described in Section 2.3.3). These will be referred to as operations made in the measurement domain. Primarily, the measurements being operated upon are the carrier phase measurements and the resulting phase rate measurements are often referred to in the literature as derived Doppler measurements. In the second stage, operations are carried out on the sequences of position and velocity that are available after the use of such a GPS software package. These will be referred to as operations carried out in the position and velocity domains. Likewise, the term acceleration domain will be used when discussing sequences of acceleration, i.e. after appropriate differentiation of the position or velocity sequences.

The research reported in Sections 4.1, 4.2 and 4.4 constitutes the contributions of the author to Bruton et al. (1999) and the results presented in Section 4.3 were obtained for the same publication together with the other authors. Work is presented without further reference to that publication.

4.1 Background

The differentiation of discrete-time signals has been studied on a number of occasions for velocity and acceleration determination using the GPS. In the measurement domain, investigations have been carried out to derive a phase rate measurement (Doppler) by differentiating the carrier phase measurements. The derived Doppler is used in place of the phase rate measurement that the receiver generates from the observed Doppler shift of the incoming signals. In the following, the receiver generated phase rate measurement is referred to as the raw Doppler.

In the literature, the derived Doppler has been considered advantageous, mainly because it is considered to be less noisy than the raw Doppler. Several methods of differentiation have been investigated for this purpose. Fenton & Townsend (1994) demonstrated the use of parabolic functions. Both Cannon et al. (1997) and Hebert (1997) approached the task using simulated GPS data by applying low-order Taylor series approximations of the derivative (i.e. central difference equations) and by differentiating cubic spline approximations to the data. A Kalman filtering approach was also proposed and applied by Hebert (1997). In each of the above cases, good solutions are obtained for static and low dynamic cases, but errors increase significantly when higher dynamics are included.

Several investigations have also been carried out that derive acceleration by differentiation in the position domain. That is, a GPS software package is used to derive the time series of both position and velocity which are then differentiated. Examples of such studies of acceleration determination for airborne gravimetry include Brozena et al. (1989), Kleusberg et al. (1990), Hehl (1990), Czompo (1991), Wei et al. (1991) and Van Dierendonck et al. (1994). Because of the nature of the gravity signal, emphasis in all of these studies is in the low-frequency part of the spectrum. Little work has been done that considers different methods of differentiating the signals, especially in cases that are not severely band limited.

The research carried out herein contributes to this field of study in a number of ways. First, the concept of differentiation is reviewed in Section 4.2. The ideal differentiator is introduced as a means of evaluating methods of approximating the derivative of a discrete-time signal. A number of differentiating concepts are then reviewed and

proposed for use in kinematic GPS systems. Some of these have already been applied in the GPS literature (e.g. central difference equations) and others have not. Although all of the filter concepts treated herein are widely known in the field of signal processing, it is important to note that they have not been evaluated for use with the GPS for deriving velocity and acceleration and, as a result, their evaluation is important herein. Each method of differentiation is presented in a common framework and subsequently characterized in terms of its performance in the frequency domain and with respect to the ideal differentiator. This approach provides an understanding of the performance of each method when used to differentiate GPS signals under conditions of variable or unknown dynamics.

In addition, it is possible to differentiate a signal by fitting a model to it and subsequently differentiating the model. This approach is treated by means of an example and compared to the methods mentioned above. In the example, the use of best-fitting polynomials is reformulated as a time domain convolution (and thereby analyzed within the same framework as the other differentiating methods). It is then characterized as a function of both polynomial order and window length. The ranges of applicability of polynomial fits are quantified in terms of their effective frequency range. This permits an evaluation of the effect of curve fitting on the spectrum of the resulting signal. For example, it is seen that the choice of polynomial order and window length affects the spectral signature.

To bridge the gap between theory and application, the practical implementation of each of these differentiators is discussed in Section 4.3. Each method is applied to real GPS data using a discrete time convolution. In other words, they are all considered to be non-recursive, (i.e. FIR) filters. Issues of phase response and filter order are also discussed.

Examples that use real GPS data are presented in order to demonstrate the effectiveness and limitations of these differentiating filters for determining vehicle velocity and acceleration under different dynamic conditions. Several methods are then applied to three practical cases that are characterized by increasing frequency content. Namely,

a) a static baseline,
b) an airborne gravimetry survey and
c) data collected on a precision motion table.

In the latter two cases, verification of the results is accomplished using independent references of higher quality (i.e. the known gravity field and a precise reference trajectory, respectively).

4.2 On the differentiation of uniformly-sampled discrete-time signals

For the purposes of this research, a discrete-time differentiator operates on a uniformly sampled sequence (of a continuous-time band-limited input signal) in such a way as to produce a corresponding output sequence that, after suitable band-limiting, approximates the actual continuous-time derivative of the input signal. According to Antoniou (1993), the frequency response, $H(e^{j\omega T})$, of an ideal uniformly-sampled discrete-time differentiator is given by:

$$H(e^{j\omega T}) = j\omega \text{ for } 0 \le |\omega| < \frac{\omega_s}{2}, \tag{4.1}$$

where ω is the frequency of the spectrum of the signal (in radians), $\omega_s \equiv 2\pi/T$ is the sampling frequency (in radians), and T is the corresponding sampling period (in seconds). The task of differentiating a discrete-time signal can therefore be seen as approximating this idealization and applying it to some signal of interest. The ideal differentiator will be used throughout Section 4.2 to assess the quality of each of the differentiators that are described.

4.2.1 On the design and implementation of differentiators for GPS

The differentiators discussed in this chapter are all cast as FIR filters. Practically, such a filter is applied to a discrete data set, $x(nT)$, using a convolution as follows:

$$x'(nT) = \sum_{k=0}^{N-1} h(nT)x(nT - nk), \tag{4.2}$$

where in this case, $x'(nT)$ is the derivative of the input sequence, $x(nT)$, and $h(nT)$ is the impulse response of the system, having length N. The relationship between the discrete-time unit impulse response, $h(nT)$, and the frequency response of the discrete-time differentiator, $H(e^{j\omega T})$, is given by the Inverse Fourier Transform. Practically then, the design of a digital differentiator becomes the problem of designing an impulse response, $h(nT)$, that can be applied to a data set using equation 4.2 and that has a frequency

response as close as possible to that in equation 4.1, within the frequency band of interest. For details about filter design, see Antoniou (1993).

The application of discrete-time filters in GPS systems is restricted by several important practical considerations. The first consideration relates to the linearity of the phase response of the filter, $\arg[H(e^{j\omega T})]$. It is especially important that $\arg[H(e^{j\omega T})]$ be close to a linear function within the frequency band of interest. This is especially critical for navigation systems because of their high dependence on the temporal characteristics of the signals. Careful treatment of both the magnitude response, $|H(e^{j\omega T})|$, and the phase response, $\arg[H(e^{j\omega T})]$, of the differentiator is therefore important. An FIR filter is appropriate in this respect, because it can be designed to have a time symmetric impulse response that implies that all spectral components of the input signal are subject to the same delay. The delay itself is easily corrected for in post-mission applications and may be taken into account in many real-time applications (depending on the processing requirements, the data rate and the length of the filter).

A second consideration is the integer nature of the delay and the corresponding odd length of the impulse response. As pointed out in Oppenheim and Schafer (1989), an even length impulse response has a significantly better magnitude response than one of similar order but odd length. This was confirmed recently in Skaloud and Merminod (2000), where the goal was to preserve the frequency content for high frequencies. Despite these improvements in the magnitude of the frequency response, an FIR filter with an even length impulse response has a major drawback in terms of the phase response. It causes the resulting delay to be non-integer, which in turn, can be problematic for time sensitive navigation systems. For the research presented herein, only odd length impulse responses are employed because the errors due to interpolation (from the non-integer to the required integer epochs) outweigh the advantages, especially for components of the signals in the medium and high-frequency bands.

A third consideration is the absolute length of the impulse response. A filter with a short impulse response offers the advantage of a small time delay and a filter with a longer impulse response offers the advantage that it can achieve a more accurate magnitude response. A further important consideration, relating to the length of the filter, occurs in

these systems when the states of the filter have to be reset because of unavoidable external events. In this case, errors occur due to edge effects. Cycle slips and gaps in GPS data are good examples of such unavoidable external events. These errors will affect less data if the length of the impulse response of the filter is small. A compromise must therefore be reached between minimizing the magnitude of the approximation error and the tolerable impulse response length. This will depend on the application and varies from situation to situation.

4.2.2 Low-order Taylor series approximations

Because they have been used extensively in the GPS literature, low-order Taylor series approximations of the derivative will be treated first. Consider the first-order central difference approximation to the derivative that is given by:

$$x'(nT) = \frac{x(nT+T) - x(nT-T)}{2T}, \tag{4.3}$$

so that the derivative of x at time nT is estimated using data at two epochs (times $nT+T$ and $nT-T$). Higher order central difference equations are often presented in the literature as an alternative to this simple first-order approximation (because it is intuitive that using more data on either side of time nT will provide a better estimate of the derivative). These higher order derivations are also based on the Taylor series. Consider for example the third-order central difference equation that has the following form:

$$x'(nT) = 45\frac{x(nT+T) - x(nT-T)}{60T} - 9\frac{x(nT+2T) - x(nT-2T)}{60T} + \frac{x(nT+3T) - x(nT-3T)}{60T} \tag{4.4}$$

The expression of these basic equations as FIR filters is relatively straightforward and useful; the impulse responses corresponding to the first and third-order differentiators given above are usually represented as vectors as follows:

$$\mathbf{h} = \frac{1}{2}\begin{bmatrix} 1 & 0 & -1 \end{bmatrix}^T, \text{ and} \tag{4.5}$$

$$\mathbf{h} = \frac{1}{60}\begin{bmatrix} 1 & -9 & 45 & 0 & -45 & 9 & -1 \end{bmatrix}^T. \tag{4.6}$$

Note that the terms *first-order* and *third-order* are used herein to refer to the filters given by equations 4.3 and 4.4, despite the fact that strictly speaking, the impulse responses in equations 4.5 and 4.6 define second and sixth order discrete-time filters. This is to remain consistent with their use in GPS literature.

Such filters can be implemented using equation 4.2 and are obviously non-causal. Expressing them in this way is useful because it permits their comparison to the ideal case. Consider Figure 4.1, for example, showing the magnitude of the frequency response for each case, derived using a normalized sampling frequency given by ω_s = 1 Hz. The ideal differentiator is also shown by the dotted line. It is clear that the first and third-order filters are only good approximations of the ideal differentiator at low frequencies. For example, they begin to deviate considerably from ideal at about 0.05 Hz and 0.15 Hz, respectively.

The above observation has several implications insofar as GPS applications are concerned. First, if the signal being differentiated is primarily a very low-frequency signal, then either of these approximations may suffice. This is the case for data obtained during a static survey and might be the case for data collected from a very slow moving platform. Second, if the signal is primarily low-frequency in nature, and there exists mainly noise in the upper part of the spectrum, then these differentiators will also serve to suppress that noise; a very encouraging result at first glance. However, the proper use of these differentiators depends on the signal to be differentiated. Caution must be exercised if the signal has medium to high-frequency components (relative to the sampling frequency) because this type of differentiator will result in spectral distortion at those frequencies, and ultimately in a derivative of lower accuracy. Investigations in which low-order Taylor series approximations have been used in the GPS measurement domain include both Cannon et al. (1997) and Hebert (1997). Results presented in both of these publications support the observations made above; that is, the approximations worked well for static and very low dynamic data sets, but give poor results for cases involving higher dynamics.

Note that by increasing the order of the Taylor series approximation, better approximates of the ideal differentiator can be obtained (especially at low frequencies). This will be examined further in the discussion of numerical designs in Section 4.2.5.

4.2.3 Designing differentiating filters using the ideal impulse response

Because the frequency response of the ideal differentiator is known, its ideal impulse response can be calculated analytically using the Inverse Fourier Transform. Practically, a differentiator can be designed by sampling, truncating and delaying the resulting infinite length impulse response. For a 24th-order filter, this yields the differentiator shown by the solid line in Figure 4.2, i.e. the filter has an impulse response that is 25 samples long. The upper graph shows the magnitude of the frequency response of the filter and that of the ideal response. The lower graph shows the error in the magnitude (i.e. the difference between the latter two curves).

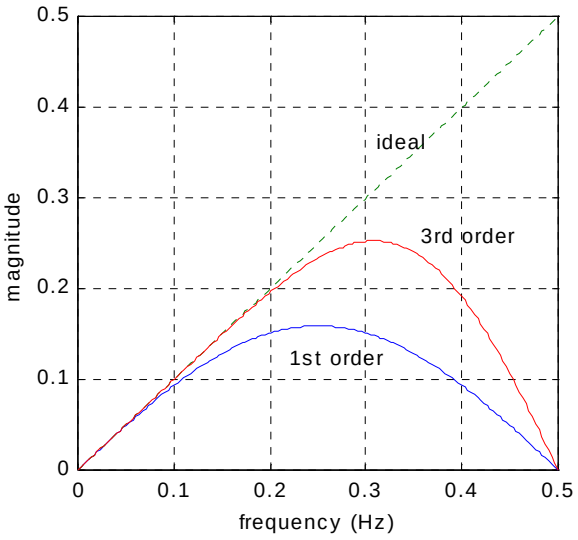

Figure 4.1 Magnitude of the frequency response of the ideal differentiator and two low-order Taylor series approximations

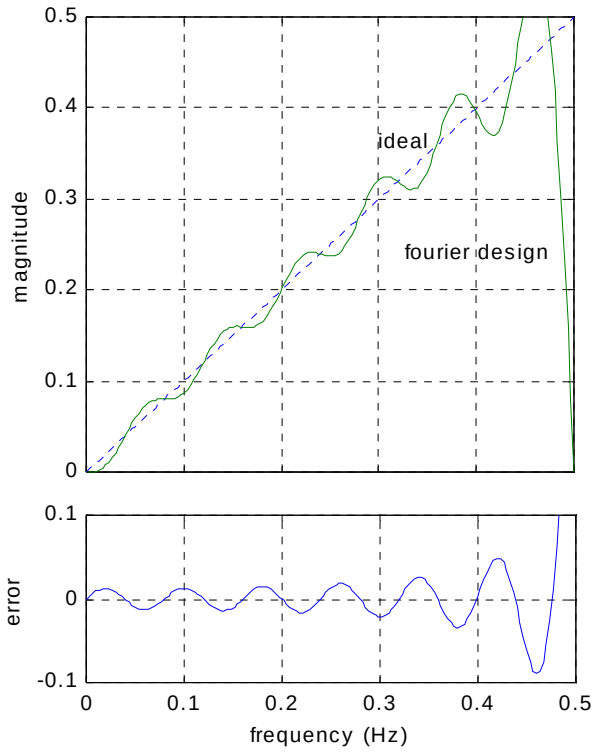

Figure 4.2 A 24th-order differentiator designed using the ideal impulse response

While this filter does not compare well in the low-frequencies with the Taylor series filters discussed in the last section, it does offer a better approximation for high frequencies. The oscillatory nature of the deviations in the response that are shown in Figure 4.2 is commonly referred to as the Gibbs' phenomenon and is the result of the finite length of the impulse response. There are several ways to improve on this design. The first and most obvious way is to increase the length of the filter, but this is subject to the disadvantages mentioned in Section 4.2.1, especially for navigation systems. The second approach is to attenuate the Gibbs' oscillations and avoids an increase in filter order. Essentially, this approach replaces the abrupt truncation of the impulse response by a gradual tapering that is defined by a window function, see Antoniou (1993). There are a number of possible window functions that can be employed, including the Blackman function that is illustrated in Figure 4.3b. Figure 4.3a and Figure 4.3c show the impulse response corresponding to the 24^{th}-order filter, before and after application of the Blackman window function (note that it was designed to have a cut-off frequency of ω_c = 0.4 Hz). Clearly, the window function reduces the magnitude of the impulse response on either side of its centre. Figure 4.4 shows the corresponding effect on the magnitude of the frequency response of the filter. As in Figure 4.2, the upper graph shows the amplitude response of the filter, and the lower graph shows the error relative to the ideal impulse response. The advantage of using the window function is evident as a result of the reduction in the magnitude of the Gibbs' oscillations. By using only symmetric impulse responses and symmetric window functions, the symmetry and therefore the linear phase properties of the filter are preserved. Examples of other window functions are given in Antoniou (1993) and Orfanidis (1996). Examples where window functions have been used in the design of low-pass filters for GPS related research include Hammada (1996) and Skaloud and Schwarz (1998).

Upon inspection of the magnitude response in Figure 4.4, it is clear that a differentiator that is based on the ideal impulse response will offer a good approximation to the ideal differentiator when a window function is used. It is also clear that such a filter has a wider range of applications than the low-order Taylor series approximations that have been considered in Section 4.2.2, because it offers a better response over the whole

frequency band. This has several implications. If the signal being differentiated has components at frequencies across the spectrum, then the output of the system will be a better approximation of the derivative. On the other hand, if noise only exists across a part of the spectrum, then the system will tend to amplify that noise. This is especially true if the signal of interest contains only low-frequency information.

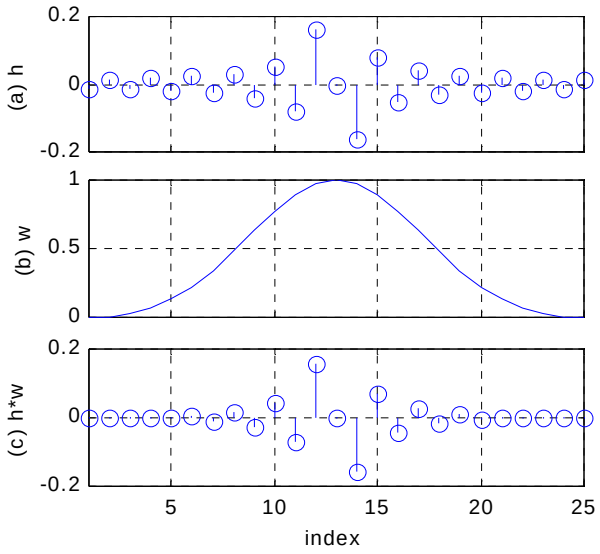

Figure 4.3 a) impulse response of the 24th-order differentiator shown in **Figure 4.2**, b) the Blackman window function and c) the windowed impulse response

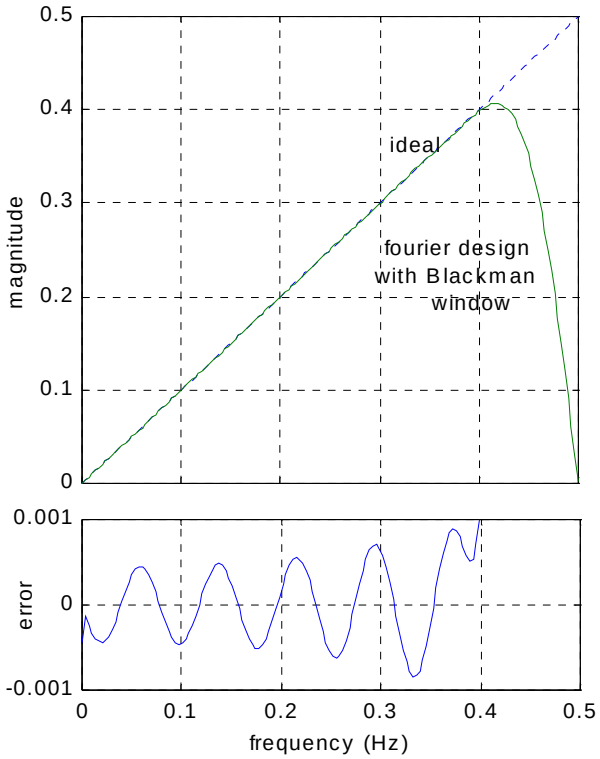

Figure 4.4 Reduction of Gibbs' phenomenon using window functions

4.2.4 Optimal designs

Although no discrete-time system can be designed to perfectly match the ideal differentiator, FIR filters can be designed that approximate it very well for all frequencies of interest. In optimal design, the impulse response of the filter is obtained by minimizing the error in the frequency response (with respect to the desired frequency response) by means of conventional numerical optimization techniques. In this section, the Remez Exchange Algorithm is used as a demonstration of this. Consider Figure 4.5 that shows the amplitude response of a 24th-order differentiator designed using the Remez Exchange

Algorithm that is provided in software accompanying Antoniou (1993). The filter was designed to have a transition band from 0.4 to 0.5 Hz, thereby ensuring that no signal content is transmitted above the Nyquist frequency, $\omega_s/2$ = 0.5 Hz, and therefore that no aliasing of the signal can occur. Clearly, this filter offers a better approximation to the ideal case over the whole frequency band than any that has been presented previously. In turn, this provides a good solution for conditions of unknown or high dynamics.

4.2.5 Numerical designs

As mentioned in Section 4.2.2, better approximations of the ideal differentiator can be obtained from the Taylor Series if higher order expansions are used. This is a subset of the filter design problem referred to as numerical design. Consider a filter design based conceptually on a central difference formula derived from a 24^{th}-order Taylor series expansion of the derivative. Figure 4.6 shows the magnitude of the frequency response of such a filter and its deviation from the ideal differentiator. A higher order Taylor series approximation clearly better represents the ideal case over a wide bandwidth.

It is interesting to note that although the amplitude response of this filter over the whole spectrum is not as good as the one designed using an optimal approach, it is very good in the low-frequency band. It will be shown in Section 4.3 that this makes it especially suitable for cases in which the spectrum of the signal is primarily at low-frequencies (e.g. static surveys and airborne gravity).

4.2.6 Curve fitting for differentiation

An intuitive approach to estimating the derivative of a noisy discrete-time signal is to fit a curve to the data and then differentiate the curve. For purposes of illustration, consider using a least squares approach to determine a polynomial that best fits the data. The instantaneous derivative of the data is obtained by differentiating the polynomial. The curve fitting approach essentially achieves a band limiting by means of polynomial interpolation. As well as providing an estimate of the derivative, this will offer a certain level of smoothing that may or may not be desirable (depending on the situation). The amount of smoothing will be a function of the order of the polynomial and the length of the sliding window in which the data is being fit.

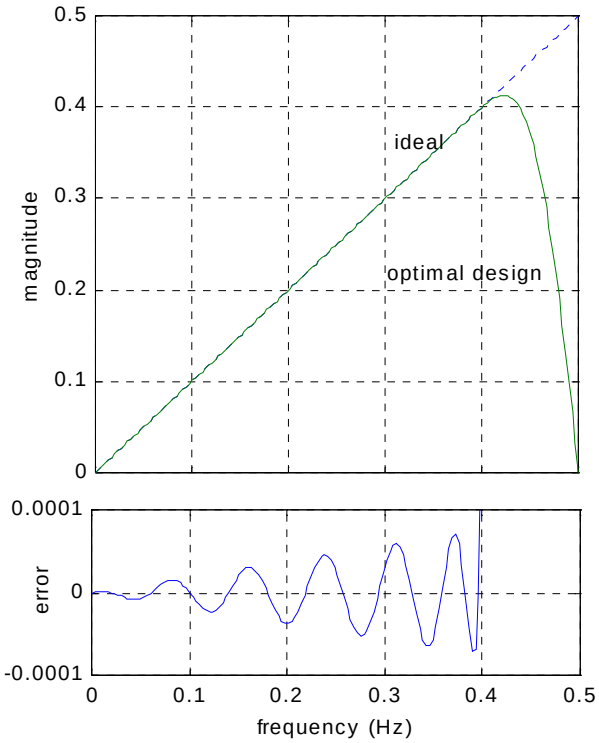

Figure 4.5 Magnitude response of the 24th-order optimal differentiator

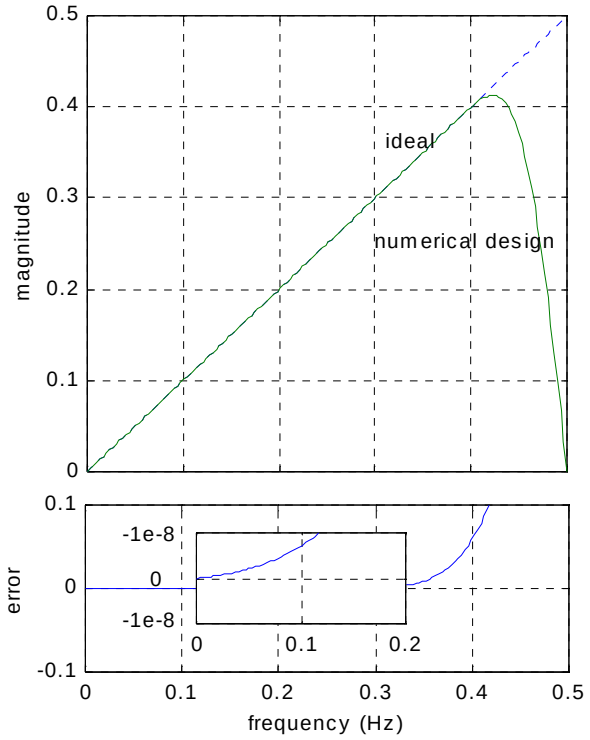

Figure 4.6 Magnitude response of the 24th-order numerical differentiator

Let the discrete-time signal at time instant t_k be given by f_k and let the sliding window of length (W+1) be centered around time t_k and index it such that the index i varies from -W/2 to W/2. Within this window, the best-fit polynomial of order M therefore has the form:

$$P(i) = a_0 + a_1 i + a_2 i^2 + \ldots + a_M i^M \tag{4.7}$$

and its first derivative is given by

$$\frac{dP(i)}{dt} = a_1 + 2a_2 i + \ldots + M a_M i^{M-1}. \tag{4.8}$$

The polynomial coefficients in the above equations can be solved for using the following standard equation for a least squares solution (without weights)

$$\mathbf{a} = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbf{f} \tag{4.9}$$

where $\mathbf{a} = [a_0 \quad a_1 \quad \ldots \quad a_M]^T$ and the i^{th} row of $\mathbf{A}$ and $\mathbf{f}$ are given by

$$\left.\begin{aligned}\mathbf{A}_i &= \left[\frac{dP}{da_0} \quad \frac{dP}{da_1} \quad \ldots \quad \frac{dP}{da_M}\right]_i \\ \mathbf{f}_i &= \left[f_{k-i}\right]\end{aligned}\right|_{i=-\frac{W}{2},\frac{W}{2}} \tag{4.10}$$

Performing this type of adjustment for every data point is computationally very burdensome and therefore two major simplifications are made. First, because the term $(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T$ in equation 4.9 is independent of the signal *f*, it only needs to be calculated once. Second, by setting i=0, the expression for the value of the derivative in equation 4.8 at the center point of the window ($t=t_k$) is given by:

$$\frac{dP(0)}{dt} = a_1 . \tag{4.11}$$

Thus, only one coefficient is needed at each epoch for each derivative to be calculated. This means that for the first derivative, only the second row of the term $(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T$ is needed. The solution for the value of the function therefore reduces to a convolution. If the second row of the above matrix is given by **h**, then the first derivative of the function is given by

$$\frac{df_k}{dt} \approx \frac{dP(0)}{dt} = \sum_{i=-W/2}^{W/2} h_i f_{k+i} . \tag{4.12}$$

This approach to fitting polynomials to a data set using least squares is referred to as Savitzky-Golay smoothing. See for example Press et al. (1992).

These curve fitting filters have symmetric impulse responses and will therefore have linear phase transfer functions leading to a constant time delay of all spectral components of the signal. When formulated as in equation 4.12, least squares fitting of polynomials to a discrete time signal can be seen as applying an FIR filter and the least squares adjustment becomes a filter design problem. This approach allows one to quantify the amount of smoothing that a best fitting polynomial offers as a function of both the order of the polynomial M and the size of the sliding window (W+1).

For completeness, it is important to realize how this type of filtering differs from optimal FIR filtering (that was discussed in Section 4.2.4). Although both methods might use a least squares fit, they are fundamentally different. A filter designed using a curve fitting approach uses least squares to fit a curve to the data being operated on. It will have a

corresponding frequency response that is determined by the choice of the fitting function. On the other hand, optimal design techniques use a least squares fit to minimize the errors in the frequency response with respect to some desired response.

Figure 4.7 shows the amplitude response of five polynomial fits. They are a set of 5^{th}-order polynomials (M=5) designed with window lengths W = 101, 49, 25, 15 and 7. Also shown for comparison is the amplitude response of the ideal differentiator. Notice first that each offers a reasonably good approximation of the ideal case at low frequencies. The frequency range over which they do a good job of approximating the ideal differentiator is clearly a function of the length of the window. Intuitively this makes sense: if the sliding window is short, then high-frequency information will be identified, whereas if it is large, only the low-frequency information will be preserved.

This frequency domain analysis allows one to quantify the amount of differentiation and smoothing resulting from a given polynomial. For example, consider using a polynomial with M=5 and W=15, being used to differentiate a signal. Figure 4.7 shows that a good estimate of the derivative will be obtained for spectral information up to approximately 0.08 Hz. Above that frequency, the spectral content is attenuated. The first zero of the magnitude of the frequency response of this filter occurs at 0.19 Hz and there are four more rises in the magnitude of the response as frequency increases. This characteristic can be disastrous, depending on the signal being differentiated. The filter should only be applied if one is interested in the derivative of the spectral components of the signal in the frequency range below 0.08 Hz. If there is any spectral information above 0.08 Hz in the signal, then the output of the filter will yield a very poor and perhaps unacceptable estimate of the derivative.

The behavior of a polynomial fitting filter is also a function of the order of the polynomial. Figure 4.8 shows a plot of the *maximum undistorted frequency* (MUF) as a function of window length for polynomial orders M = 2, 3, 5, 10 and 20. (MUF is defined as the frequency above which the filter no longer approximates an ideal differentiator with better than 0.5% of the maximum frequency content of the signal.) This figure can be used to determine the range of parameters that are suitable for a given application. Obviously, a polynomial fitting filter can only be used to estimate the derivative in cases

where the signals are band-limited and low-frequency in nature. The range of applicability is a function of the parameters *M* and *W*, as given by Figure 4.8.

Because the curve fitting problem has been treated by using polynomials as an example, some comments about other models are in order. Whether the curve being fit is a polynomial, parabola or cubic spline, its frequency response will be similar in nature to those in Figure 4.7; that is, it will be suitable for band-limited low-frequency signals and it will only approximate the ideal differentiator up to a certain frequency. In turn, this means that such filters can only be used for static or low dynamic conditions. This has been confirmed in the research reported in several publications. For example, Fenton and Townsend (1994) used a parabolic fit within the receiver to estimate the derivative. They found that velocity results were improved in static cases. Both Cannon et al. (1997) and Hebert (1997) approached this task in post mission using simulated GPS data by employing cubic spline fits for differentiation. In each of these cases, good solutions were observed in static and low dynamic cases, but errors increased with higher dynamics.

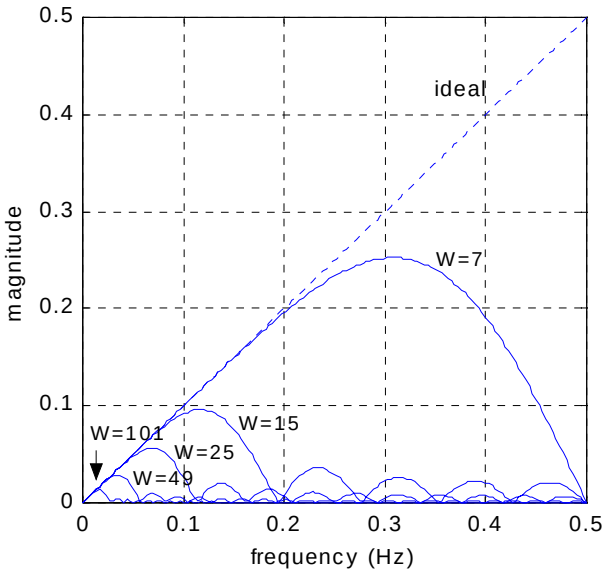

Figure 4.7 Magnitude responses of 5th-order polynomial fits as a function of filter length

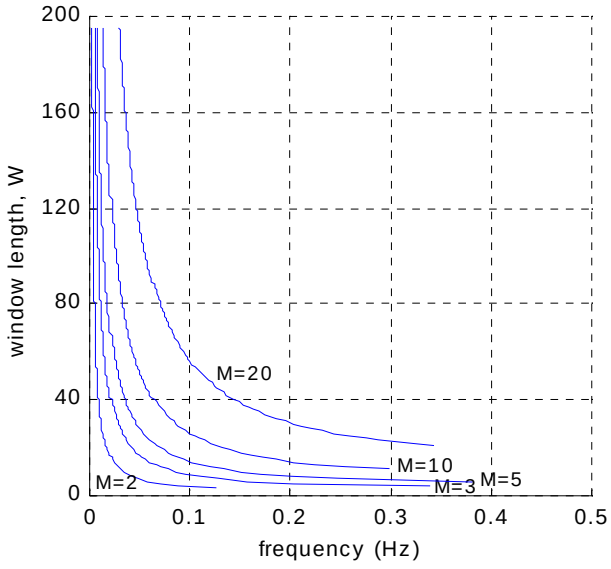

Figure 4.8 MUF as a function of window length for various polynomial orders

4.2.7 Comments on other methods of differentiation

By restricting this discussion to the use of FIR filters, several other approaches have not been mentioned. For example, it is well known that a Kalman filter smoother results in a smooth, band-limited solution. This result was confirmed in Hebert (1997). For this reason, and because it has been shown to offer little or no improvement over the low-order Taylor series approximations, Kalman filters are not treated in this dissertation. Also, infinite impulse response (IIR) filters have not been treated herein. This is for several reasons. Although they operate recursively over a smaller filter mask, the required impulse responses are just as long and their transient effects are of roughly the same length as for the equivalent FIR filter. Additionally, designing IIR filters with good phase characteristics is less straightforward than for FIR filters and one must often perform post-mission reverse time filtering in order to achieve a zero-phase response. Although not included in this initial analysis, it may be worthwhile to consider IIR filters as part of future work.

4.3 Implementation and testing using real GPS data

The goal of this section is to implement the differentiating filters that were introduced in Section 4.2 and to demonstrate their appropriate use under a variety of dynamic GPS conditions. This will be done in both the measurement domain and the position domain. In each, appropriate differentiation methods will be applied to cases of increasing dynamics, namely, a static DGPS baseline, an airborne gravity survey and DGPS data collected on a precision motion table undergoing known dynamics. In the latter two cases, results are verified using independent references (the known gravity field and the precise reference trajectory provided by the motion table).

As outlined in the last section, all of the differentiators that are treated herein are FIR filters having linear phase transfer functions, meaning that they can be applied to any GPS data set by employing a discrete time convolution. Two software packages have been developed at the University of Calgary for this purpose. One is a pre-processor for differentiation of carrier phase measurements and one is a post-processor for differentiation of position and velocity. For the measurement domain investigations that

are carried out in the next section, the pre-processor called DERived DOPpler (DERDOP) was written to apply any impulse response to the phase measurements according to equation 4.2. It derives the phase rate (Doppler) by differentiating the carrier phase data. A phase velocity trend method is used to detect cycle slips using the raw Doppler observable and the filter is reset if gaps in the data or cycle slips are detected. In cases when the filter is reset, a central difference approximation is used to estimate the phase rate during the period in which the filter is reset (i.e. corresponding to half of the order of the filter). The carrier phase data coming from each satellite are treated independently, meaning that a slip on one does not reset the filters for the other satellites. The original version of the DERDOP software (that implements the ideas presented in the last section) was written for Bruton et al. (1999) by C. Glennie.

The software package KINGSPAD is used to estimate position and velocity from the measurements whenever required. Because the baselines are short in all cases, the solutions are based on the L1 carrier phase measurements.

For the position and velocity domain investigations that are carried out in Section 4.3.2, a post-processor called GPSACC is used. It is part of the GREATGUN package and also applies any impulse response according to equation 4.2. Any gaps in the time tagged position and velocity solutions are filled in using a linear interpolation prior to filtering.

In the following numerical studies, curve fits and differentiators based on the ideal impulse response are not used. This is because the range of applicability of the curve fits is limited and similar conclusions are reached when using the differentiators based on the ideal impulse response as when using optimal designs.

4.3.1 Analyses based differentiation in the measurement domain

Three data sets are analyzed in this section to assess the performance of the different methods of differentiation in the measurement domain (i.e. performance of the derived Doppler). They were chosen to represent varying levels of dynamics. All three data sets were collected using a 1 Hz sampling frequency.

No dynamics case: static data analyses

For the static case, approximately 40 minutes of data was collected on a baseline of approximately 6 metres with an Ashtech Z-XII and a Trimble 4000SSE receiver. After a

static initialization period of 100 seconds, the data was processed in kinematic mode. The errors in velocity obtained by deriving a Doppler using the first and third-order Taylor series approximations and 48th-order differentiators designed using numerical and optimal approaches are displayed in Table 4.1. Also shown are the results obtained if the receiver-generated raw Doppler is used. Table 4.1 shows that each method of phase differentiation provides significantly better results than the raw Doppler in both RMS and mean values for all three components of the velocity. The best RMS improvement is realized using the first-order Taylor series approximation, followed by the third-order and the numerical filters. The optimal differentiator gives the poorest results, but this is expected. Because the data set is static, there is no velocity information in the frequency bands where the Taylor series approximations begin to deviate from the ideal differentiator. Therefore, instead of causing differentiation errors in this band, the two approximations actually behave similar to low-pass filters and dampen the higher frequency noise. However, since the numerical and optimal differentiators are still close to ideal in the higher band, they tend to amplify the noise. This effect is shown in Figure 4.9 where the frequency spectra of the residuals for the first-order and optimal differentiators are displayed for one of the components.

It is also important to notice that the numerical differentiator offers a better solution than the optimal differentiator for this static case. This can be explained by looking at Figure 4.5 and Figure 4.6. The error in the response of the numerical filter at very low-frequencies is of the order of $1x10^{-8}$, while it is of the order of $1x10^{-5}$ for the optimal filter.

Table 4.1 GPS velocity errors for static data

Doppler type	**RMS (mm/s)**			**Mean (mm/s)**		
	north	**east**	**up**	**north**	**east**	**up**
Raw	4.00	3.11	7.04	0.20	-1.32	1.75
1st-order	0.65	0.48	1.00	0.20	-0.05	0.18
3rd-order	0.93	0.70	1.54	0.21	-0.07	0.17
numerical (48)	1.24	0.96	2.16	0.20	-0.05	0.17
optimal (48)	1.33	1.04	2.32	0.20	-0.05	0.18

Note that although the exact numbers in Table 4.1 would be different if other types of receivers had been used, it is clear that the relative improvement offered by the various methods of differentiation would not change. It should also be noted that if it is known a-priori that the antenna is static, it is also known that the velocities should be zero.

Low dynamics case: airborne gravimetry

Airborne gravimetry provides a very illustrative example for several reasons. The first reason is that, as noted in Chapter 2, errors due to GPS still make up a significant portion of the error budget of an airborne gravity system, especially as its bandwidth is widened. This is especially true if the acceleration is derived from estimates of the velocity rather than from estimates of position. The second reason is that, although the dynamics experienced by the survey aircraft during a flight line are obviously higher than in the static case, they are nevertheless relatively low due to the benign flight conditions, so that it becomes a good intermediate example.

To analyze the improvement in airborne gravity results when the acceleration is estimated from GPS velocity (that is in turn based on the derived Doppler data), four flight lines of approximately 1200 seconds duration from an airborne gravity test are analyzed. The test took place on September 10 1996, over the Rocky Mountains in Alberta, Canada. More details can be found in Appendix A.1, where it is referred to as the Kananaskis data set.

The Trimble 4000SSI receivers in Banff and on the aircraft are used as the master and remote stations, respectively. The onboard SINS/DGPS combination is used to derive gravity disturbances (according to the procedure that is outlined in Chapter 2), and the estimates are then compared to reference gravity disturbances that were upward continued from ground gravity measurements. The accuracy of the estimates is therefore a function of both DGPS and SINS errors. Since the SINS errors are dominant mainly at lower frequencies, the accuracy of the estimated disturbances is largely a function of DGPS noise.

Table 4.2 gives the errors in gravity disturbance for different filtering periods. It shows the average of the results obtained for the four flight lines closest to the master station (see Bruton et al. (1999) for details). The optimal and numerical differentiators are each 48^{th}-order filters. The 60 and 90 second filtering times are typical periods used in

airborne gravity and the others are included to demonstrate the performance for higher frequencies.

An analysis of the results in Table 4.2 shows that the first-order filter now performs worse for the shorter period than the raw Doppler observations, which is in opposition to the observations made above for the static data. This change in performance is due to the introduction of vehicle dynamics. Consider the average spectrum of the upward component of the velocity for the airborne test that is shown in Figure 4.10. There is clearly velocity information up to approximately 0.2 Hz. The remaining spectral content above 0.2 Hz is due mostly to measurement noise. Recall Figure 4.1 that shows the amplitude response of the first and third-order differentiators. The first-order response deviates significantly from ideal in the band between 0.02 to 0.2 Hz. This deviation is in a band where vehicle motion is present and explains why poor results were obtained from the first-order differentiator. The third-order differentiator, however, is still very close to ideal up to 0.2 Hz, and therefore shows good results that are practically identical to the numerical and optimal results.

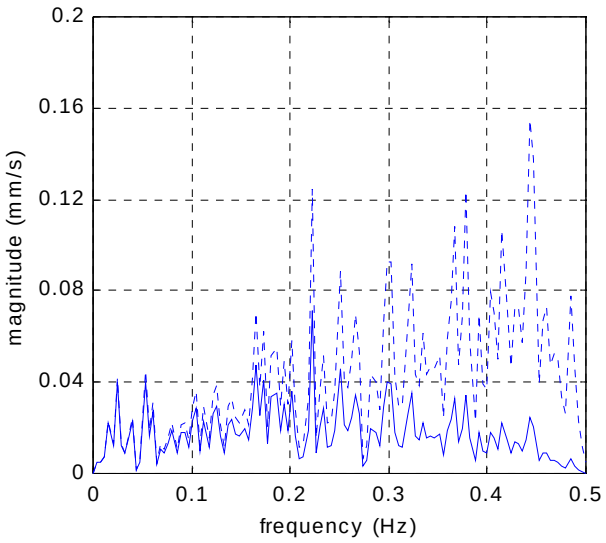

Figure 4.9 Spectrum of the residuals when using the 1st-order and optimal (dotted line) differentiators, north direction

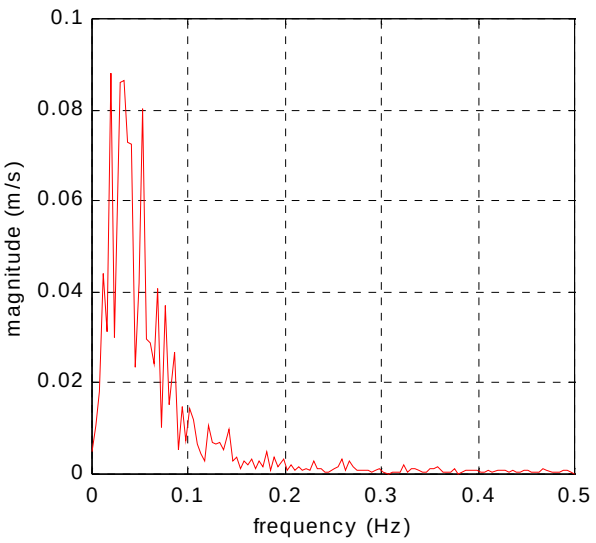

Figure 4.10 Spectrum of the upward component of the velocity of the aircraft for the airborne data

Table 4.2 Average agreement of gravity estimates with the reference (Based on derived Doppler)

Doppler type	**Standard deviation (mGal)**			
	15 s	**30 s**	**60 s**	**90 s**
raw	57.1	15.1	4.9	3.2
1st-order	144.7	20.3	3.9	2.8
3rd-order	41.9	11.0	3.6	2.8
numerical (48)	41.8	11.0	3.6	2.8
optimal (48)	41.9	11.0	3.6	2.8

High frequency dynamics case: motion table data

As a final test for comparing differentiation techniques in the measurement domain, a data set collected on an Anorad precision motion table was analyzed. An Ashtech Z-XII antenna was placed on the motion table and followed a prescribed periodic trajectory. A second Ashtech Z-XII was used and approximately 20 minutes of GPS data was collected in differential mode. The Anorad table allows the user to specify a trajectory along one axis, and provides a position and velocity reference for the motion at the sub millimetre (and millimetre/s) level. It therefore permits the generation of a data set with a dominant frequency (in this case at 0.2 Hz) and a very accurate reference trajectory. To provide a better idea of the test dynamics, a spectrum of the reference trajectory for the test is shown in Figure 4.11.

The results of the analysis are given in Table 4.3. The optimal and numerical differentiators are again 48th-order filters. The results presented show that for this higher frequency data set, the optimal and numerical differentiators perform significantly better than the two low-order Taylor series approximations. This is due to the presence of the higher frequency velocity information between 0.15 and 0.25 Hz and in the band around 0.4 Hz. In these bands, the Taylor series approximations deviate significantly from the ideal differentiator, and therefore do a poor job of estimating the higher dynamics. Because the optimal and numerical differentiators have nearly ideal responses up to 0.4 Hz, they have no problem properly differentiating the higher frequency dynamics. Figure 4.12 shows the difference in the residuals between the third-order approximation and the optimal differentiator, clearly demonstrating that the errors in the third-order derivative

case grow with increasing frequency. The optimal differentiator outperforms the numerical differentiator in this case because it is designed to have a better response over the whole spectrum. This might not be the case if the motion had more low frequency content.

Table 4.3 GPS velocity errors for the Anorad test (measurement domain)

Doppler Type	RMS (mm/s)	Mean (mm/s)
raw	261.9	-2.5
1st-order	237.1	-1.2
3rd-order	48.9	-0.5
numerical (49)	27.3	-0.5
optimal (49)	27.0	-0.5

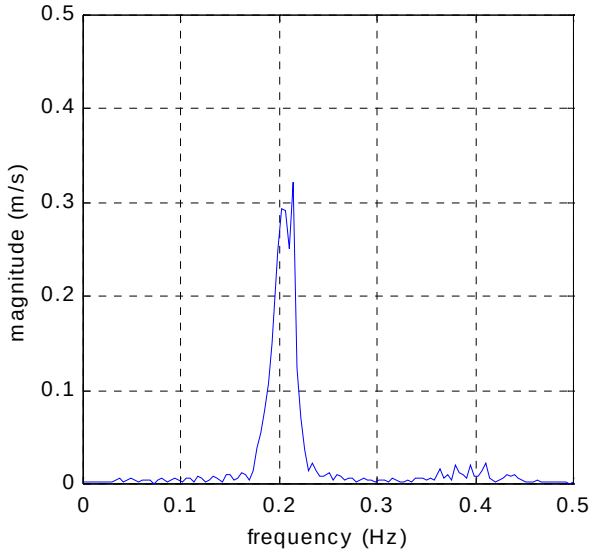

Figure 4.11 Spectrum of the reference trajectory (velocity) for the motion table test

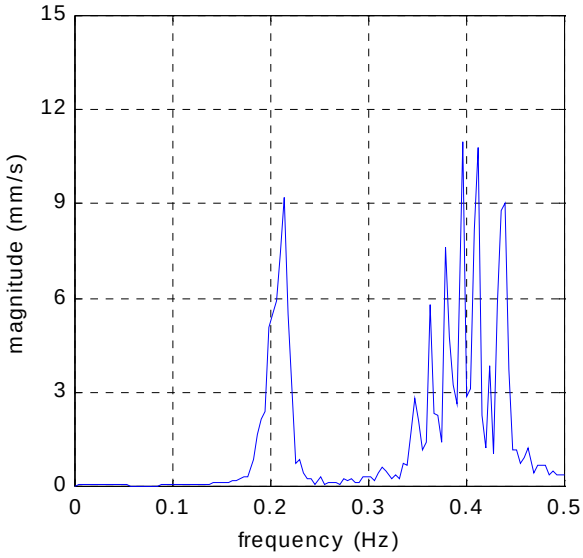

Figure 4.12 Spectrum of the difference between the 3rd-order and optimal solutions

4.3.2 Analyses based on differentiation in the position domain

Two tests were conducted to analyze the performance of the different methods of differentiating position and velocity to determine velocity and acceleration, respectively. In both, the output of KINGSPAD was used as input to the post-processor GPSACC. In

the first case, the low dynamic airborne gravity data is again studied. In the second, the higher frequency trajectory provided by the precision motion table is used again.

Low dynamics case: airborne gravimetry

In this case, the position of the aircraft is determined using KINGSPAD from carrier phase measurements and then differentiated twice using each filter to obtain acceleration. The same flight lines that were used in Section 4.3.1 are again used. The upward continued reference is again used for the evaluation. The results are shown in Table 4.4. Again, the best results were found to be at a filtering period of 90 seconds and agree with those obtained by differentiation in the measurement domain (see Table 4.2). It is important to notice that as the frequency band becomes larger (i.e. at filtering periods of 60, 30 and 15 s), the acceleration results achieved by differentiating the position data are better than those from the derived Doppler measurements. This may become increasingly important as the resolution of airborne gravity systems is increased. On the other hand, it is of interest that the best level of agreement (at 90 s) was matched using a solution that is based on the derived Doppler. This implies that acceleration can be obtained without the explicit need for ambiguity resolution, a fact that is discussed further in Chapter 5. Numerical and optimal filters of 48^{th}-order were also investigated and found to offer no improvement over the 24^{th}-order filters.

High frequency case: motion table data

In this case, the position of the antenna on the motion table is determined from carrier phase measurements using KINGSPAD and then differentiated using a number of filters to obtain velocity. This was then compared to the truth offered by the reference trajectory. The results are shown in Table 4.5 and several observations can be made. The first is that the Taylor Series approximations offer considerably worse solutions than the numerical and optimal filters. Observing the filter responses (Figure 4.1) and the spectrum of the motion (Figure 4.11), this is obviously due to the fact that they are far from ideal in the part of the spectrum corresponding to the motion of the table. Second, it should be noted that the 48^{th}-order numerical filter offers a better solution than the 24^{th}-order one. Finally, notice that in all cases, the optimal filter offers a slightly better

solution than the numerical filter. Higher order optimal filters were investigated and did not improve the solution.

It is also important to notice that the level of accuracy achieved in Section 4.3.1 with measurement domain techniques is greater than that achieved with these position domain techniques. This is due to the fact that when the dynamics are high, input to the GPS engine should be as accurate as possible (which is only the case if differentiation is done in the measurement domain).

Table 4.4 Average agreement of gravity estimates with the reference (based on differentiation in the position domain)

Doppler type	**Standard deviation (mGal)**			
	15 s	**30 s**	**60 s**	**90 s**
1st-order	273.1	36.2	4.1	2.8
3rd-order	35.2	10.1	3.4	2.8
numerical (48)	35.2	10.1	3.4	2.8
optimal (48)	35.2	10.1	3.4	2.8

Table 4.5 GPS velocity errors for the Anorad test (position domain)

Doppler Type	**RMS (mm/s)**	**Mean (mm/s)**
1st-order	240.7	0.6
3rd-order	48.4	0.6
numerical (24)	34.1	0.6
optimal (24)	31.1	0.6
numerical (48)	31.6	0.6
optimal (48)	31.0	0.5

4.4 Summary

The process of differentiation for precise velocity and acceleration determination has been addressed in this chapter. A number of differentiating filters of varying complexity were studied and characterized in the frequency domain to permit an assessment of their applicability in different dynamic scenarios. Their implementation in both the measurement and position domains has been discussed with respect to deriving velocity and acceleration for both static and kinematic cases. The performance of the

differentiators has been demonstrated using real GPS data in cases where reference trajectories were available.

To compare methods of computing a derived Doppler for GPS velocity and acceleration determination, three data sets with varying degrees of dynamics were analyzed. Based on this analysis the following conclusions are drawn:

a) Deriving the phase rate (Doppler) from the carrier phase measurements provides accurate estimates of vehicle velocity, as long as the applied differentiator is close to ideal in the frequency band where the vehicle motion is present.
b) The optimal differentiator provides the best results for high dynamics due to the fact that it most closely approximates the ideal differentiator throughout the frequency band of interest. It has been shown, using the precision motion table, that by using the correct filter, velocity estimates under medium-frequency dynamic conditions are improved by an order of magnitude to 27.0 mm/s, relative to the raw Doppler, for the receivers used in this case.
c) The Taylor series approximations provide reasonable results only for static data or for cases of low-frequency motion. This is due to their deviation from an ideal differentiator in the middle of the frequency band. Caution should be exercised when trying to apply these differentiation techniques. For the data mentioned in b) above, central difference approaches were shown to be deficient, yielding results up to 10 times worse than the proposed preferred methods.

To compare the results obtained by differentiation in the measurement domain to those in the position domain, the low and high-frequency cases were repeated for the latter. Based on this analysis the following conclusions are drawn.

d) In the low dynamic case of airborne gravity, differentiation in the position domain appears to give better results than in the measurement domain.
e) For the high frequency case of the precision motion table, differentiation in the measurement domain gives better results than in the position domain.

The results obtained using the precision motion table in this chapter demonstrate that it is possible to estimate the velocity of an antenna moving with high frequency, with an accuracy of better than 5 cm/s, as required for airborne gravimetry.

Future research on estimating accurate velocity and acceleration from GPS measurements might include:

- Using a variety of different GPS receivers in order to evaluate the performance of each. For example, it is likely that the bandwidth of the tracking loop and the method of raw Doppler calculation will vary from receiver to receiver. Herein, this was considered secondary to providing an understanding of the characteristics of various differentiators, demonstrating their use under various dynamic conditions and demonstrating the relative performance levels of each.
- Band-pass filtering the data after a careful analysis of the spectral content of the vehicle motion and prior to differentiation. For example, for the airborne gravity flight, the signal above 0.2 Hz is mostly measurement noise. Therefore, an optimal differentiator might be designed with a cut-off frequency of 0.2 Hz instead of 0.4 Hz.
- Applying a noise reduction technique such as Kalman filtering/smoothing or spline fitting, to remove noise after the phase measurements have been appropriately differentiated.
- Evaluating the appropriateness and possible advantages of using infinite impulse response (IIR) differentiating filters.

5 On the Optimal Determination of Acceleration from DGPS

In this chapter, a detailed study is presented of the DGPS error budget for acceleration determination, within the bandwidth relevant to airborne gravimetry. There are two principal objectives:

a) To use currently available methods to analyze and quantify the effect that each component of the error budget has, employing data sets that represent a wide range of operational conditions. Conceptually, this extends the analyses carried out in Chapter 3 to the estimation of acceleration.

b) To propose and demonstrate specific approaches to the optimal estimation of aircraft acceleration for the purposes of determining the geoid and for resource exploration.

5.1 Estimating the acceleration of the aircraft

It is recalled from Chapter 2 that the acceleration of the aircraft, $\dot{\mathbf{v}}$, is typically derived from DGPS by using carrier phase and phase rate measurements to estimate position and velocity and then appropriately differentiating these quantities. As shown in Chapter 3, the accuracy of DGPS positioning is mainly a function of receiver noise, multipath, sub-optimal estimates of the ambiguity terms, and (depending on baseline length) uncompensated atmospheric and orbital errors. After differentiation, each of these error sources plays a role in the error, $d\dot{\mathbf{v}}$, of the resulting acceleration. Also, it is recalled from Chapter 4 that the process of differentiation amplifies these errors as a function of increasing frequency, causing them to be larger as the bandwidth is increased.

Within the bandwidth implied by the sampling frequency, the result is a very noisy estimate of the acceleration that is usually complicated by errors due to aircraft dynamics. Consider Figure 5.1, for example, that shows how the error in estimated acceleration increases with frequency for a controlled kinematic test. See Bruton et al. (1999) for details of this test.

Although somewhat dependent on receiver type, it is observed in practice that estimates of acceleration obtained by twice differentiating time series of position data are more accurate than those obtained by differentiating time series of velocity data. This is due to the typically more noisy phase rate data obtained when using the raw Doppler measurement (see Chapter 4). This section employs the widely used approach of deriving aircraft acceleration by twice differentiating the time series of estimated positions. Because the gravity signal is relatively low-frequency in nature (usually with most of its spectral energy well below 0.05 Hz for typical flying speeds), much of the noise shown in Figure 5.1 is removed in practice by low-pass filtering. Recall from Chapter 4, however, that computing acceleration using velocity estimates obtained from a Doppler measurement (that is in turn derived from carrier phase measurements) can yield nearly the same performance as deriving acceleration from the estimated position, but without the need for ambiguity resolution.

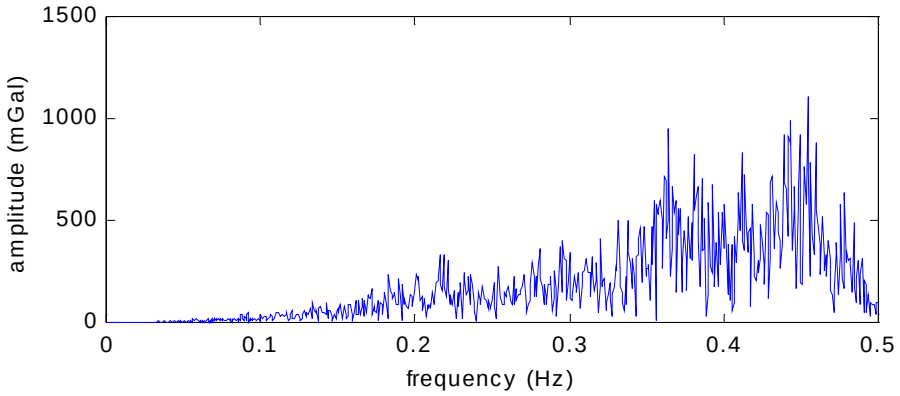

Figure 5.1 Spectrum of the error in acceleration determined for a kinematic test

This whole process was treated extensively in the late eighties and early nineties in investigations such as those outlined in Brozena et al. (1989), Kleusberg et al. (1990), Hehl (1990), Czompo (1991), Wei et al. (1991) and Van Dierendonck et al. (1994). As a result of their work, it is generally accepted that an accuracy of 1-2 mGal can be achieved under good environmental conditions, within a bandwidth corresponding to a cut-off frequency, f_c, of approximately 0.01 Hz (or for a period, $T = f_c$, of 100 seconds). However, for airborne gravimetry, useful information exists above 0.01 Hz and this

information should be extracted from the data and employed to obtain more accurate estimates of acceleration. This is made especially challenging by the fact that errors due to GPS increase in this bandwidth. As a result, errors due to GPS continue to pose a major challenge to researchers attempting to increase the bandwidth of their systems.

Recall from discussions in Appendix C.1 that the cut-off frequency is related to the maximum achievable spatial resolution via the flying speed of the aircraft. For example, a cut-off frequency of 0.01 Hz implies a shortest half-wavelength of 2.0 km when flying at 45 m/s. In turn, this implies a sample spacing of 2.0 km. In order to increase the spatial resolution of the gravity estimate along the flight path, it is therefore desirable to increase the cut-off frequency. However, this inevitably leads to lower accuracy in the estimates of the gravity vector in large part because of the error spectrum shown in Figure 5.1. This critical issue is investigated in this chapter.

5.2 Analysis of the DGPS error budget for acceleration determination

The effects that each of the residual error sources in equation 2.14 have on the determination of acceleration are presented in this section. Although they are treated in a different order than in Chapter 3, the general approach is very similar. The same DGPS processing package is used as well as a number of well-chosen data sets (including the two used in Chapter 3). Similar data processing strategies are used to isolate and quantify each portion of the error budget. For each portion of the error budget, the error in the DGPS position solution is differentiated to derive the corresponding error in acceleration. This results in a set of empirically derived error spectra that characterize the errors as a function of frequency (and therefore as a function of space for a fixed flying speed). These error spectra are generated within the bandwidth between 0.0 and 0.05 Hz, which corresponds to half-wavelengths as short as 450 m at a flying speed of 50 m/s. The results are obtained for the error in the estimated three-dimensional acceleration vector (for which the error in vertical acceleration is one component).

5.2.1 GPS receiver measurement noise and multipath in static mode

Because GPS receiver measurement noise is a fundamental limit on the accuracy of the derived acceleration, it is treated first. Although the noise level is dependent on the

tracking loop bandwidth (and therefore the receiver type and the vehicle dynamics), it can be approximately estimated using a zero baseline test in a static environment. Consider Figure 5.2 that has been generated using the same triplet of receivers used in Section 3.2.5. Figure 5.2a shows the error spectrum of the error in acceleration due to double difference carrier phase receiver noise (for the L1 and ionospheric-free measurements), $\nabla\Delta\,\varepsilon_{\Phi}$, between the two receivers sharing the rooftop antenna. For the L1 carrier, this error is almost negligible for the bandwidth shown and it is much more significant for the ionospheric-free observable, especially above 0.015 Hz. Figure 5.2b shows the magnitude of the combined error due to noise and multipath, generated by differentiating the error in the L1 double difference position solution calculated over the baseline between one of the rooftop receivers and the receiver in the field. It was generated for exactly the same time period as Figure 5.2a and shows that the double difference multipath, $\nabla\Delta\,m_{\Phi}$, on the L1 carrier considerably increases the error in the same bandwidth, especially for frequencies around 0.035 Hz (or for half-wavelengths shorter than 640 m at a flying speed of 45 m/s).

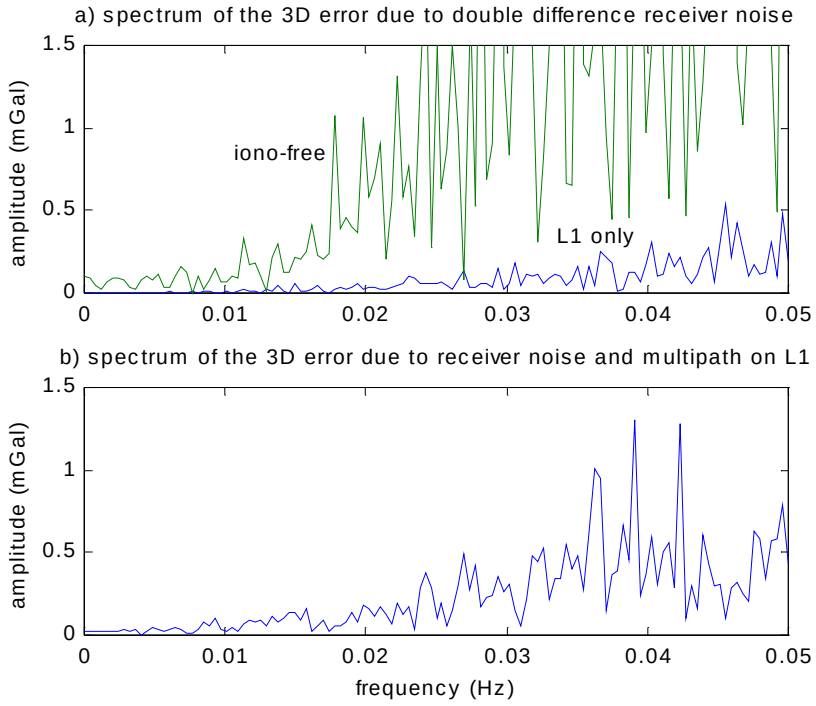

Figure 5.2 Error spectra owing to receiver noise and multipath

Although these graphs have been derived empirically using a limited data set, they are considered to be conservative estimates, given that both represent the double difference errors in three dimensions and that the multipath environment is not calm.

5.2.2 GPS receiver measurement noise and multipath in kinematic mode

The effect of noise and multipath in the airborne environment is again roughly quantified by deriving the relative acceleration between the two antennas that were mounted on the aircraft during the collection of the Greenland data set. The error can be estimated by differentiating the difference between the two position solutions derived using the same master station and each of the antennas (i.e. differentiation of the sequence that was derived in Section 3.2.5). Figure 5.3 shows the resulting spectrum. As shown, the estimated effect of the combined kinematic receiver noise and multipath at the aircraft does not differ greatly in the relevant bandwidth from that observed in the static situation on the ground that was shown in Figure 5.2b. This implies that the level of multipath in this bandwidth is similar on the aircraft to that seen in a medium multipath environment on the ground.

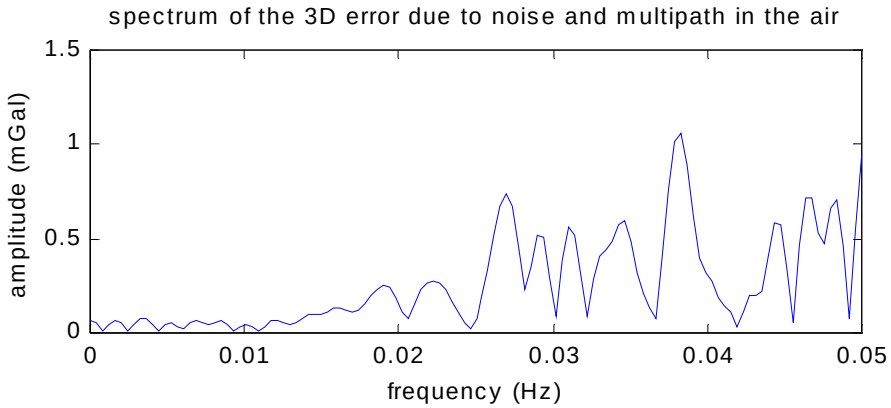

Figure 5.3 Spectrum of the error due to noise and multipath at the aircraft

5.2.3 The importance of ambiguity estimation

Although this study puts no emphasis on the process of estimating the ambiguity term, $\nabla\Delta N$, it is relevant to consider the importance of using a valid estimate of it.

Consider Figure 5.4 that shows the spectrum of the difference in acceleration (for the Kananaskis data set) between solutions derived using an L1 float solution and the (widelane) fixed solution. This shows that reliable ambiguity fixing makes very little difference over what appeared to be an equally reliable float solution.

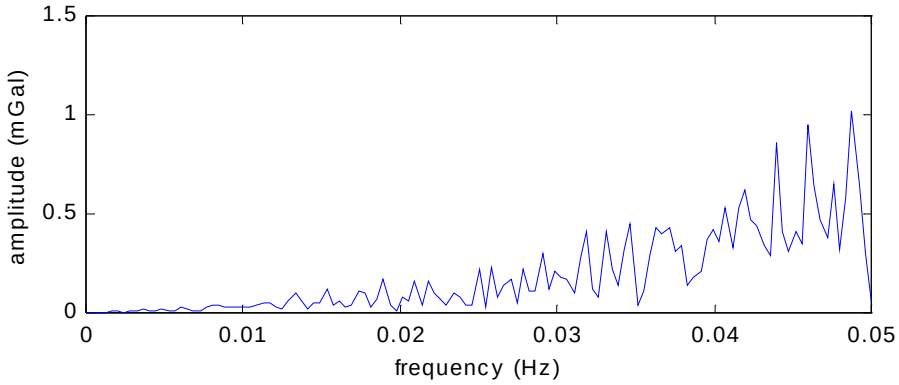

Figure 5.4 Spectrum of the error in a float solution

In order to demonstrate that the fixed solution provides a better absolute accuracy than the corresponding float solution, both were used to generate estimates of the gravity field that were in turn compared to the true field. Figure 5.5 shows the error spectra generated by that comparison. The solid and dash-dot lines in the figure show that the fixed solution is indeed slightly better than the float solution across the whole bandwidth.

The shaded region in Figure 5.5 is included to demonstrate the approximate magnitude of the component of the total error spectrum that comes from the inertial sensors and from imperfections in time synchronization (i.e. the area between the shaded region and each of the lines is due to errors in the corresponding DGPS solutions). The derivation of the shaded error spectrum is given in Section 7.3.

Note that the data used for this graph was collected on an aircraft having high dynamics within the bandwidth of interest. It is shown in Section 5.5 that using a different aircraft (exhibiting less dynamics) can make a significant difference to the results and it is shown in Section 7.4 that a large portion of the effect of the dynamics induced SINS errors can be removed by post mission modeling and signal processing. Because it is affected by biases in the inertial sensor, the portion of the bandwidth below 0.00025 Hz is not shown for any of the spectra in Figure 5.5 (this corresponds to half-wavelengths greater than 90

km for a speed of 45 m/s and is discussed in detail in Chapter 6). Note that there is also some contribution, primarily below 0.01 Hz, due to inaccuracies in the assumed gravity field.

Although it has been shown here that a minor practical advantage is gained in striving for a fixed solution, it is very important to be aware that a solution that is generated using ambiguities that are incorrectly fixed can have disastrous effects on the resulting accuracy of the acceleration estimate, especially when changes occur in the satellite constellation. Recall that the Kananaskis data set is selected because it was flown at night and under excellent conditions for DGPS processing. Although nighttime conditions are not a prerequisite for good GPS, they often imply it. Unless data is collected under conditions where one has confidence in it, then using a fixed integer ambiguity solution is not recommended for airborne gravimetry. This is discussed in more detail in Section 5.4.

Finally, note that it is shown here that a very reasonable solution can be obtained using a single frequency approach. In other words, a single frequency float solution offers an effective and relatively cheap alternative to a dual frequency approach if the requirements with respect to environmental conditions can be met.

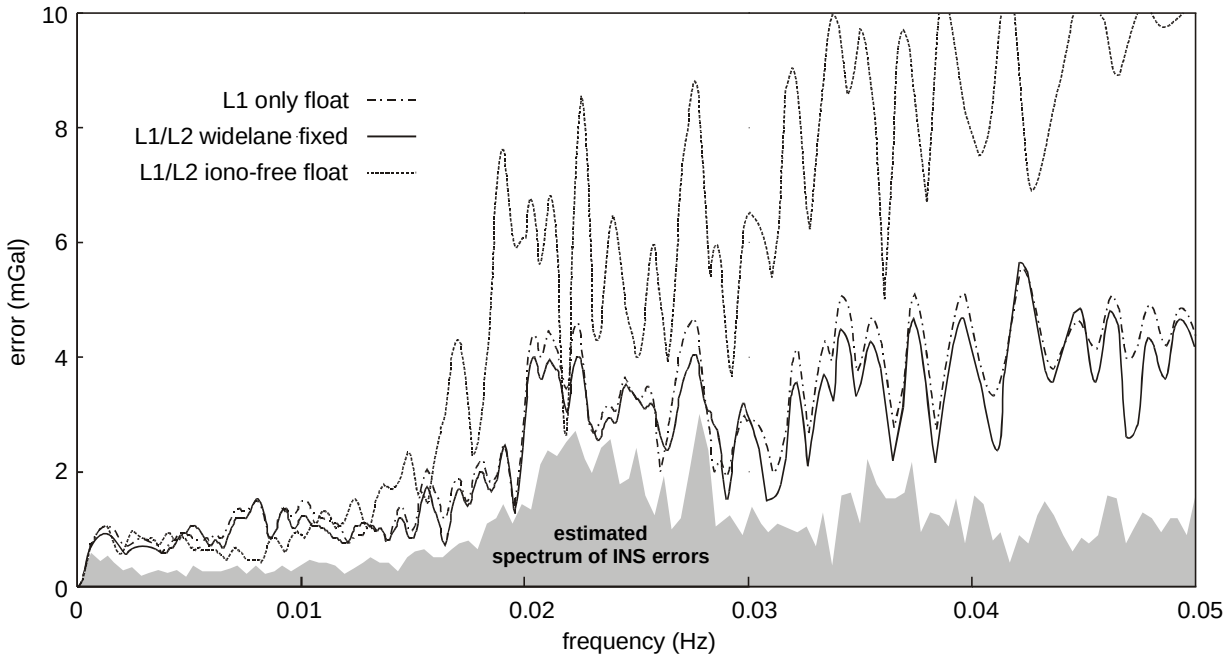

Figure 5.5 Error spectrum of the estimated gravity field for several DGPS solutions for the Kananaskis data set

5.2.4 Evaluating the effect of the ionosphere

The effect that the ionosphere has on estimates of acceleration can be evaluated by comparing the solutions that have been obtained using the ionospheric-free float to those obtained using the L1 only float solution (or a fixed solution if it can be estimated). Because the ionospheric-free combination removes all first-order effects of the ionosphere, the difference between these solutions is due to both the ionosphere and the increase in noise that results from the linear combination. The first of these effects is desirable, whereas the second is obviously not, as shown in Section 5.2.1.

Recall that the Greenland data set was selected for this study because it was collected during a period of high ionospheric activity. As mentioned in Chapter 3, one consequence of the active ionosphere, in that case, is that it was not possible to reliably resolve the integer ambiguities for more than a portion of the flight. The use of the ionospheric-free observable made it possible to obtain a reliable position solution and also results in a much more accurate solution than the L1 only float solution for acceleration determination, for the whole bandwidth. It is interesting to note that, despite the obvious advantages gained by using the ionospheric-free approach under such an active ionosphere, it can remain very difficult to obtain accurate gravity data at high resolutions. Note that this case is considered to be extreme and is made so difficult by the fact that the data was collected in the mid afternoon during a period of high ionospheric activity over long baselines and at high latitudes.

The dotted line in Figure 5.5 shows the error spectrum of the gravity solution obtained for the Kananaskis data set using the ionospheric-free float solution. It demonstrates two very important points. The first point is the very evident increase in noise over the other solutions for most of the spectrum, especially above 0.015 Hz. This confirms the preliminary observation, made in Section 5.2.1, that the increase in noise associated with the ionospheric-free combination makes that combination ill-suited for wide bandwidth applications (even in cases where the ionosphere is behaving well). The second important point is that the ionospheric-free solution has a lower error than the other solutions in the bandwidth between 0.005 and 0.01 Hz. While this may not seem very significant when displayed on the scale used in Figure 5.5, the reader is reminded that standard filtering

procedures used in geoid determination typically involve a low-pass filter with a cut-off frequency of around 0.01 Hz. It is evident that the advantage gained from using the ionospheric-free observable can be considerable in that bandwidth. This implies the following:

a) the ionospheric-free observable is best suited for use in medium-frequency applications such as geoid determination, and
b) a float (or fixed) L1-based solution is best suited for high resolution applications when the environmental conditions are good.

5.2.5 Errors due to changes in geometry

Changes in geometry due to changes in the satellite constellation that occur during a flight line, can result in very large errors in the resulting acceleration. Although such an event will often have small effects on the estimated position, it is easy to demonstrate that they can cause errors of thousands of mGal in the unfiltered acceleration data and as much as 10 mGal at a filtering period of 90 s. For the purpose of demonstration, Figure 5.6 shows an example of this (where the effect is relatively small) taken from the Kananaskis data set. It demonstrates what happens in the time domain when a satellite is introduced into the constellation at 800 s. Notice that different scales are used on the vertical axes. It is clear from this graph of the unfiltered differences that they are very high-frequency in nature and occur following the 800 s epoch. It is also demonstrated that, although low-pass filtering reduces this error, it smears the effect of introducing the satellite over a time period that roughly corresponds to the filtering period.

Figure 5.7 shows the same effect as a function of frequency. Note that the magnitude of this error can easily reach significantly larger magnitudes around the time of introduction (or removal) of the new satellite (i.e. around 800 s). There are several practical ways to alleviate the effect of such changes in geometry. If the flight lines are short and the constellation is strong enough, the user can manually force the rise and fall of satellites to occur during the turns. While this will often mean that a lower value of the Dilution of Precision (DOP) might be observed for a portion of a flight line, it ensures that rapid jumps like the one observed in Figure 5.6 will not affect the solution.

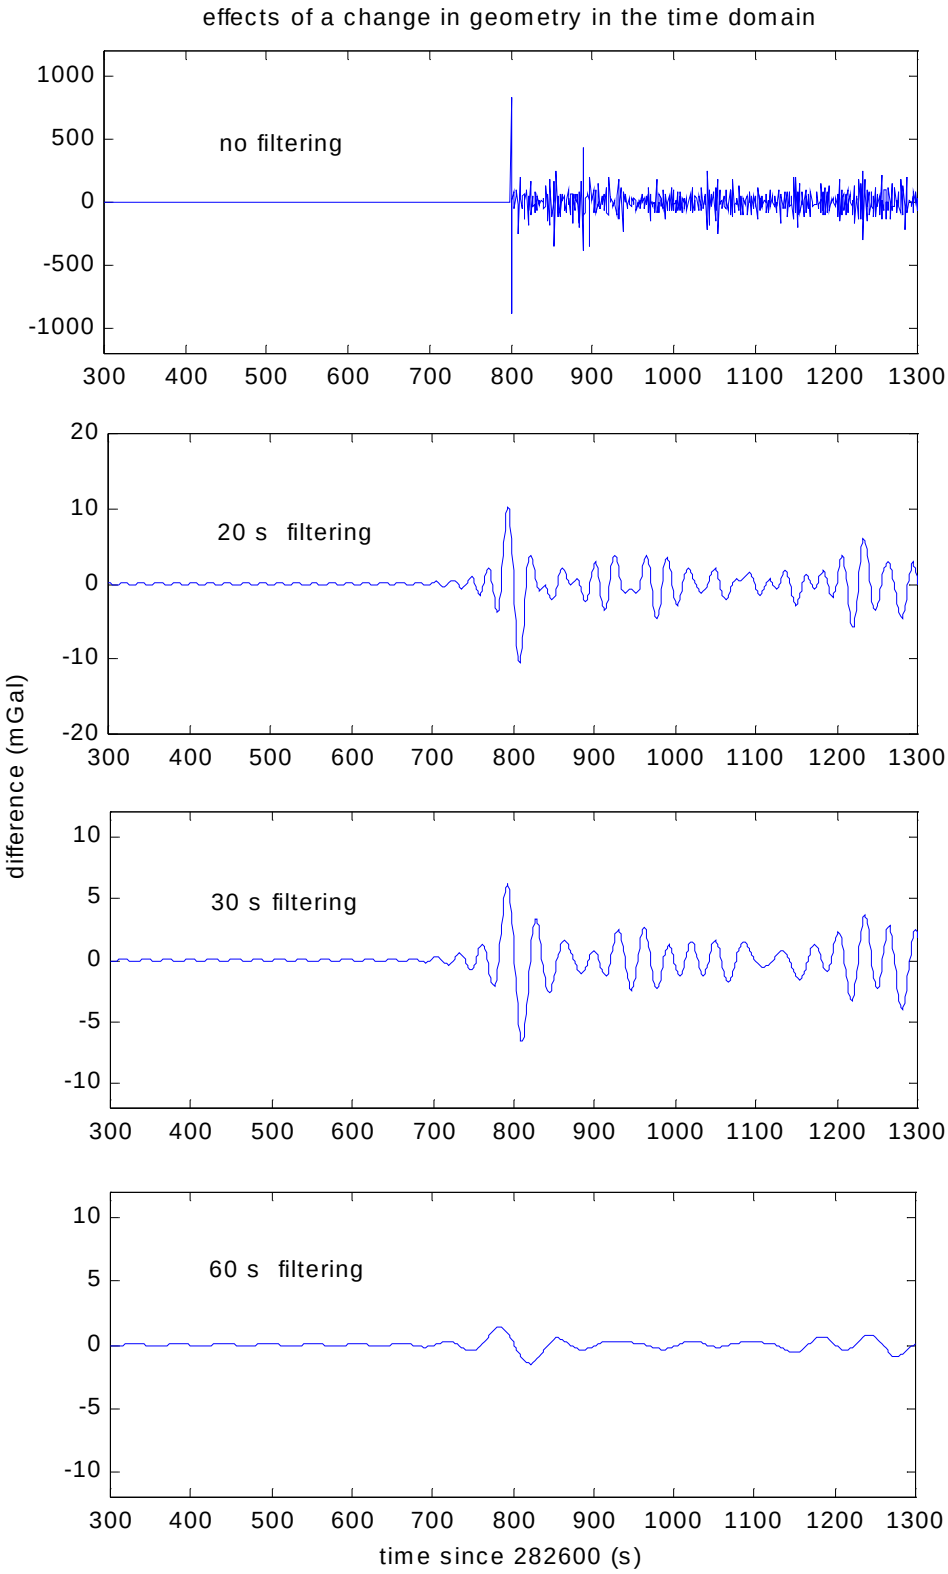

Figure 5.6 Effect of a change in geometry due to the inclusion of a new SV at 800 s

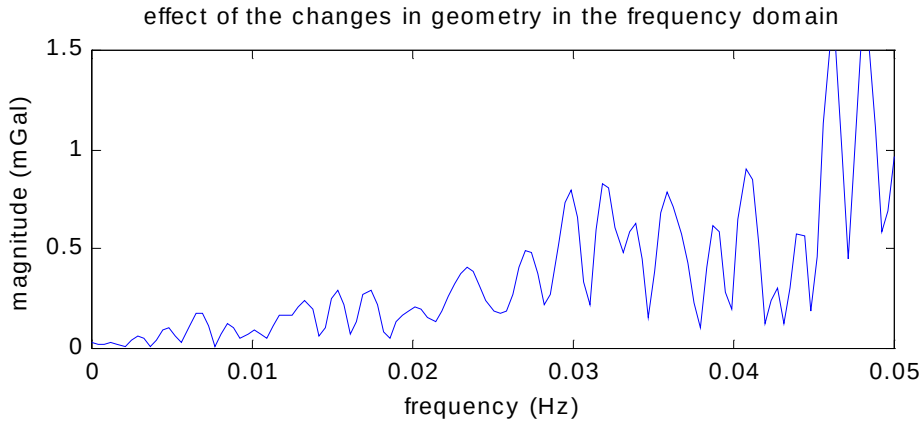

Figure 5.7 Spectrum of the effect of the same change in geometry

Another more elegant approach involves a careful introduction (and removal) of satellites from the active constellation by means of a time-varying weighting function. For example, if measurements from a new satellite are weighted as a function of the time since their appearance, a jump such as that shown in Figure 5.6 is avoided and the solution takes advantage of the improved geometry. An alternative to a weighted combination of the measurements is discussed in Brozena and Childers (2000). In that case, jumps in the position solution are estimated and removed by shifting the solution for all epochs following a jump. This results in a time series that is well-suited for differentiation. Depending on the situation, one of these types of approaches might be crucial for long flight lines or cases where the gravity data in the turns is needed.

5.2.6 Errors due to the neutral atmosphere and the broadcast ephemeris

Errors due to the broadcast ephemeris and inaccuracies of the tropospheric model do not greatly influence acceleration derived from DGPS. As noted in Chapter 3, they are both typically smooth and bias-like and are therefore removed by the process of differentiation. Figure 5.8a shows the spectrum of the error caused by errors in the broadcast ephemerides. This spectrum has been computed by comparing the solutions generated using broadcast and precise ephemerides. Figure 5.8b demonstrates the difference in acceleration between solutions obtained using two different tropospheric models (Black and Saastamoinen) under otherwise identical processing conditions. Both effects are clearly negligible.

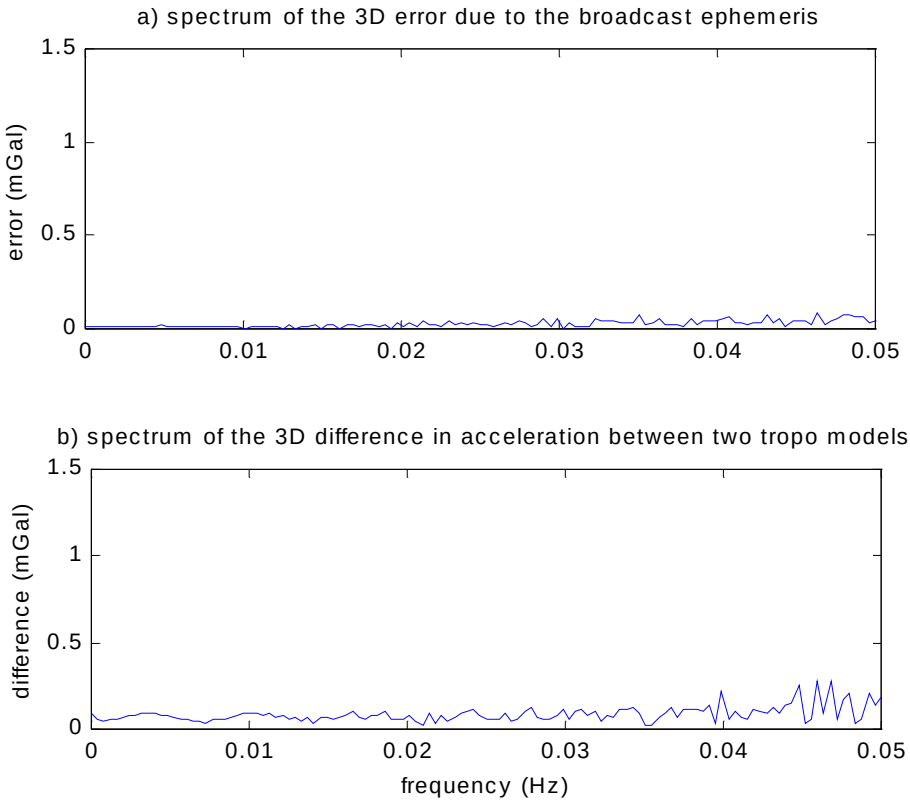

Figure 5.8 Error spectra due to orbital and tropospheric model errors

5.3 The (lack of) importance of the data sampling rate

Although the rate at which raw GPS data is sampled may affect the accuracy of the estimated position, it has very little effect on acceleration determination, especially within the bandwidth of interest for airborne gravity. This is shown in Figure 5.9 where the spectrum of the difference in acceleration between two solutions has been obtained using sampling rates of 1 Hz and 10 Hz. The figure was generated using airborne data collected by the University of Calgary and Sander Geophysics Ltd. in which a single antenna on the aircraft fed two receivers, each collecting raw GPS data at the rates mentioned above. The 1 Hz data was collected using an Ashtech Z-12 receiver and the 10 Hz data was collected using a NovAtel MiLLennium receiver. The data at the master station was collected using a NovAtel MiLLennium receiver at 10 Hz and subsequently decimated to 1 Hz, providing a means of estimating the position of the airplane using raw data at both frequencies. After computing each acceleration solution, they can be directly

compared to each other if the 1 Hz solution is up-sampled to 10 Hz. This was done using an up-sampling method that is known to preserve the frequency content of the original signal, as outlined in Orfanidis (1996). The filter used in the latter approach was designed to have a transition band between 0.35 Hz and 0.50 Hz, implying that the spectrum is almost perfectly preserved below 0.35 Hz and above 0.50 Hz.

The plot shows the spectrum of the difference in the upward component of the acceleration between these solutions (now both sampled at 10 Hz) for a period that included a flight line and the turns on either side of it. Ambiguities were reliably fixed in both cases for the whole time period. The difference is clearly very small within the bandwidth that is represented by both data sets (0.0 to ~0.35 Hz).

While some of the differences above 0.50 Hz may arise from increased noise levels in the data sampled at 10 Hz, it is also possible that they are due to aircraft motion; there is a significant amount of motion captured by the 10 Hz data that is not captured by the 1 Hz data. The inset in Figure 5.9 clearly shows that the difference between the two is negligible for the bandwidth of interest in airborne gravity. This is supported by the fact that these observations include the higher dynamic turns.

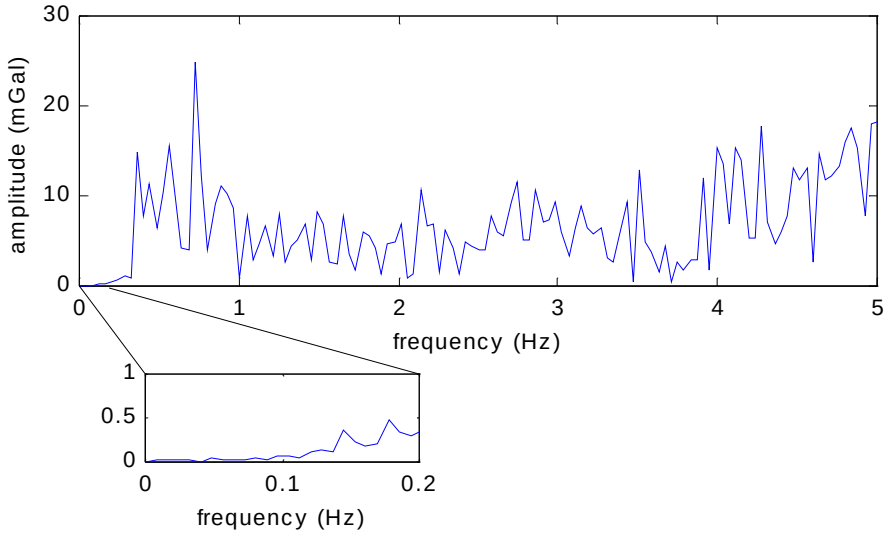

Figure 5.9 Frequency spectrum of the difference in acceleration between solutions computed from data collected at 1 Hz and 10 Hz

5.4 Recommendations for determining acceleration

Based on the findings in Sections 5.2 and 5.3, several recommendations can be made regarding the best ways to process DGPS data for acceleration determination under a variety of conditions. These include the following:

a) For applications that require medium-resolution gravity data (such as geoid determination), it is recommended that both the L1 and ionospheric-free solutions be considered. The latter can have some advantages during good GPS conditions and can be a necessary approach during poor GPS conditions (especially if the ionosphere is active).

b) For applications that require high-resolution information (such as resource exploration), an L1 solution is recommended whenever possible. It is to be understood that GPS conditions have to be good for this approach to be viable. In cases when conditions are not suitable, the approach in a) is the recommended default approach and a wide bandwidth solution should not be expected. When the GPS data is reliable, a fixed integer ambiguity solution is desirable and is often made more feasible through the use of the L2 observable in a widelane combination.

c) For all applications, it is recommended that changes in the satellite constellation be held to a minimum. If possible, satellites should only be allowed to leave or enter the solution during the turns between flight lines (i.e. periods during which gravity data is not usually being used). If the flight lines are too long, the geometry is too weak to permit this, or if the data in the turns are needed, a weighting or data shifting strategy is recommended to avoid abrupt changes in the estimated acceleration.

d) Minimization of the multipath in the environment surrounding the master station is recommended. This implies selecting a location with minimal multipath and using choke-ring antennas.

Although conditions that are conducive to the collection of very high-frequency gravity information are more likely to be achieved if baselines are kept very short and flights are flown at night, these are not necessary constraints. This is shown in the next section

where data collection conditions from a recent airborne gravity campaign are described and the accuracy of the resulting gravity estimates is shown.

5.5 How good are current DGPS techniques for determining acceleration?

This section uses results from a recent airborne gravity survey to demonstrate that an DGPS accuracy of better than of 1.5 mGal can be achieved at a spatial resolution (half-wavelength) of 2.0 km and, further, an accuracy of better than 2.5 mGal can be achieved at a spatial resolution of 1.4 km, when following the recommendations made in the last section. This data set is chosen, because, unlike the data used in Section 5.2, it was collected using an aircraft with very good dynamic characteristics and within the bandwidth of interest, implying that much of the total error spectrum is due to DGPS, especially for high frequencies.

The data was collected as part of the AGEM project of the GEOIDE Network. The campaign constitutes coverage of the Alexandria test area (near Ottawa, Canada) and had the goal of comparing all of the three currently available airborne gravity concepts on a single aircraft (an air-sea gravimeter, the LRF-III SINS and a 3-axis inertially stabilized platform). It is described in detail in Appendix A.3. Only two surveys of the campaign are described herein. The first was flown on April 19, 2000 between 16:00 and 21:00 and the second on May 4, 2000 between 09:30 and 12:30. The latter flight was flown in the morning in an effort to minimize the effect of turbulence and the atmosphere.

This campaign is very useful for demonstrating the quality of current DGPS techniques, mainly because the Cessna Grand Caravan exhibited very little horizontal motion in the bandwidth of interest for airborne gravimetry, especially for high frequencies. This is discussed in detail in Chapter 7. Therefore, the errors affecting high-frequency gravity field estimation are due mainly to the errors coming from DGPS acceleration determination and can be analyzed. Although any of the on-board gravity systems could be used (they all yield the similar conclusions with respect to the DGPS performance for high frequencies), the LRF-III system was selected for this demonstration in order to maintain consistency with other results presented in this dissertation. A detailed comparison between systems is being carried out in a separate study and the investigators of the project should be contacted for details.

Because the GPS data was reliable for both days and because the baselines reached just over 100 km, processing was done using a fixed integer widelane approach. The resulting average gravity error spectra for each day are shown in Figure 10 along with the estimated error spectrum of the inertial sensor errors (calculated for the April 19 data set using the method that is outlined in Chapter 7). This figure demonstrates that the time of day has a large influence on the level of the noise coming from DGPS, almost certainly because of the higher level of ionospheric activity in the afternoon. As shown, the level of noise due to DGPS for the data collected in the morning is only slightly above that observed for the data set collected at night (i.e. compare the solid lines in Figure 5.5 and Figure 5.10). Also notice that the noise level seen for the data set collected in the afternoon is similar to the level of noise that was introduced by doing an ionospheric-free combination for the nighttime data set shown in Figure 5.5.

It is very encouraging to observe that the total error spectrum can still be as low as 1-1.5 mGal at a cut-off frequency of 0.025 Hz. For the flying speed of 45 m/s used in this test, this implies that **gravity signals with half-wavelengths as small as approximately 1 km can be detected if their amplitude is greater than 1-1.5 mGal**.

The agreement of the solutions with the ground gravity data (at flight height) is given in Table 5.1. As usual, the standard set of filters used have transition bands that are as sharp as possible in order to very nearly represent the true accuracy in a given bandwidth.

The table demonstrates that it is possible to obtain an agreement with the reference field of 1.5 mGal at a bandwidth corresponding to a half-wavelength resolution of 2.0 km and an agreement of 2.5 mGal at a half-wavelength of 1.4 km when flying at a time of day when the ionospheric activity is low. These numbers are similar to the cumulative error spectrum. In the afternoon (on what was considered to be a day with moderate levels of ionosphere activity), it was possible to obtain accuracies of roughly 2 and 4 mGal for the same bandwidths.

In terms of achievable accuracies, the results described in Section 5.5 imply that, unless the ionosphere is very active:

a) for medium-frequency applications, such as geoid mapping, an accuracy of 1.5-2.0 mGal can be expected for the gravity disturbance, regardless of time of day, and

b) for high-frequency applications, such as resource exploration, an accuracy of 1.5 mGal can currently be achieved for the gravity disturbance, at a spatial resolution (half-wavelength) as high as 2.0 km for nighttime and morning flights.

It should be noted that the values given above, and in Table 5.1, reflect the combined accuracy of the DGPS and SINS components. In other words, the DGPS error makes up one component of the quoted values. This is reflected by the spectra shown in Figure 5.5 and Figure 5.10.

Table 5.1 Standard deviation of the agreement between estimated solutions and the reference field (mGal)

Line	Morning data set filtering periods (s)				Afternoon data set filtering periods (s)			
	30	60	90	200	30	60	90	200
T1001	10.7	2.5	1.5	1.5	10.2	4.4	3.7	3.4
T1002	7.0	2.1	1.2	1.2	10.6	3.6	2.3	2.4
T1003	7.3	2.7	1.9	1.8	27.6	4.8	2.7	1.6
T1004	11.1	3.2	1.1	1.2	22.3	3.0	1.1	0.8
T1005					15.6	3.8	2.1	1.8
T1006					23.4	4.4	1.6	1.1
average	**9.0**	**2.6**	**1.4**	**1.4**	**18.3**	**4.0**	**2.3**	**1.9**

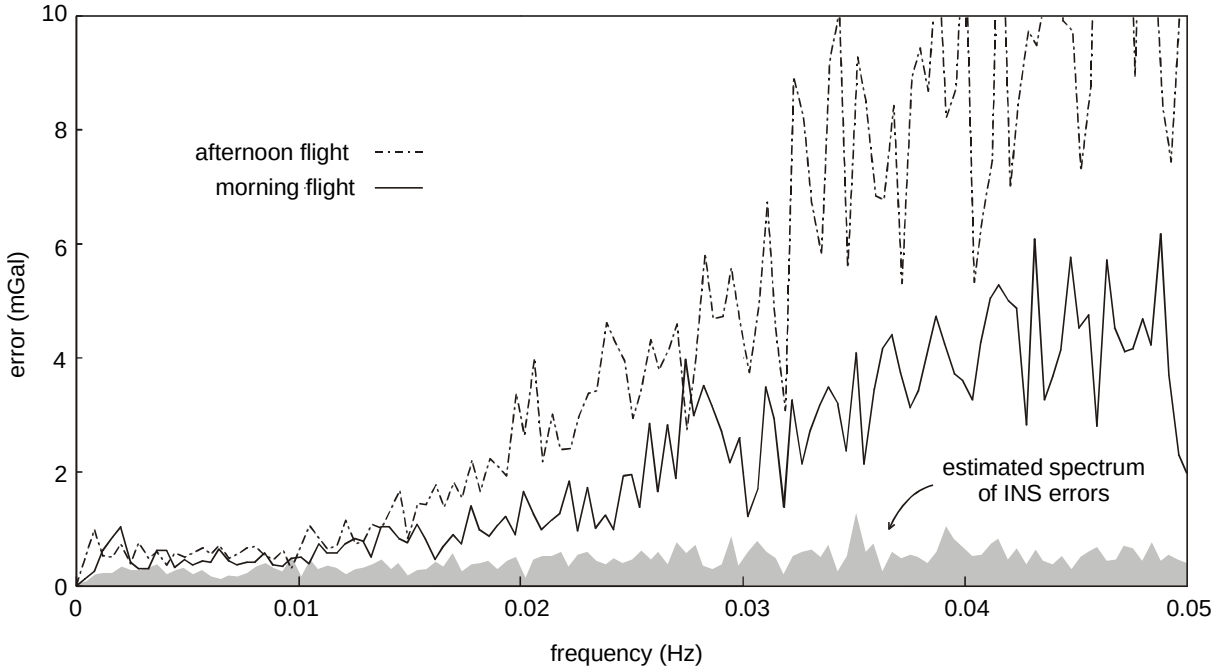

Figure 5.10 Error spectra of two estimates of the gravity field from the Alexandria campaign

PART 3: UNDERSTANDING AND HANDLING THE SINS ERRORS

The objective in the third part is to investigate the errors in estimates of the gravity field due to SINS sensor and modeling errors. The investigation has two components. The first is to thoroughly analyze the principle SINS errors that limit the bandwidth of the airborne gravity system and the second is to propose novel methods of dealing with them. The two chapters that follow concentrate on the low and high-frequency errors, respectively.

In Chapter 6, the discussion of the requirements for gravity data is carried on from the material in Chapter 1. The advantages and disadvantages of airborne gravimetry as a collection method are briefly demonstrated in order to establish the role that it plays in estimating the relative gravity field. This discussion is included in order to make the material in Part 3 more accessible to the unfamiliar reader and to clarify a number of points that have been raised recently by researchers working in the field (e.g. in meetings of the AGEM project). Original contributions are described in the balance of the chapter. Using real data, the low-frequency errors (due to residual accelerometer biases) that make it challenging to use a SINS-based gravimeter are discussed and demonstrated. A method that allows for the direct estimation of the relative gravity field from airborne data is then proposed. The performance of the system is evaluated and demonstrated in different bandwidths.

In Chapter 7, emphasis is placed on the high-frequency SINS errors that are induced by the motion of the aircraft. A detailed study is presented of the natural motion of survey aircraft in flight, both from a theoretical point of view and empirically using data from two survey aircraft. The effects that these dynamics have on estimating the high-frequency components of the gravity field are described in detail. The chapter concludes by discussing methods of reducing the errors due to aircraft dynamics, including the proposal and implementation of a novel method that removes a significant portion of them.

6 On Estimating the Relative Gravity Field Using the SINS Gravimeter

This chapter has two objectives. The first is to describe the role that airborne gravimetry plays in estimating the gravity field of the Earth and to relate it to the requirements for relative gravity information in geoid determination and resource exploration. This is done in Sections 6.1 and 6.2 through the presentation and interpretation of important background information and related research. The second objective is to describe the low-frequency errors that make it challenging to determine the relative gravity field when using a SINS-based system. This includes a demonstration of the accuracy of the system used by the University of Calgary for a variety of bandwidths and a discussion of how the relative information that is derived can be used in practice. This is done in Sections 6.3 and 6.4, which contain the contributions made by the author.

6.1 The role of airborne gravimetry as a data collection method

Although the gravity field of the Earth has power at all resolutions (even a grain of sand causes some gravitational attraction), it decreases rapidly as a function of increasing frequency. This means that the majority of the gravity signal is low-frequency in nature and has several implications. The first is that, depending on the application, sufficient information can often be drawn from an accurate band-limited estimate of the gravity field. The second is that no single measuring system can measure the gravity field accurately at all resolutions; a satellite-based system is too far from the Earth to measure high-frequency information and ground-based measurements cannot easily cover large areas for low-frequency information, to name a few examples. As mentioned in Chapter 1, there are a number of current and future measuring systems that provide information about the gravity field, and airborne gravity systems are just one example. Generally speaking, the resolution of all methods is limited on one end by factors such as sampling rate and the distance of the measuring system from the surface of the Earth and on the other end by bias-like errors and the spatial extent of the gravity survey. The following sections describe some of these in an effort to quantify the role that airborne gravimetry

plays. (Recall that the high, medium and low-resolution components of the gravity field were defined in Chapter 1 as those corresponding to half spatial wavelengths of less than 5 km, between 5 and 150 km and above 150 km, respectively. Also recall from Appendix C.1 that the terms *half-wavelength* and *spatial resolution* are used synonymously.)

6.1.1 Information about the low-resolution content of the gravity field

Currently, low-frequency information about the gravity field is available on a global scale in the form of geopotential models that have been estimated mainly by observations of the perturbations in the orbits of satellites and supplemented with data from satellite radar altimetry and shipborne surveys (over the oceans), and surface and airborne gravimetric surveys (over land). An example of such a model is the Earth Gravitational Model of 1996 (EGM96) that is a spherical harmonic expansion complete to degree and order 360, corresponding to a minimum half-wavelength of roughly 55 km. See Appendix C.3 for a brief discussion of the spatial resolution corresponding to an expansion of the spherical harmonic equations of a global geopotential model. The information in the model corresponding to half-wavelengths longer than 300 km is derived mainly from the orbit perturbation data and higher resolution information is derived from the altimetry, surface, shipborne and airborne data. Because the data available on land are often sparse or non-existent (depending on location), the accuracy of such a model for half-wavelengths shorter than about 300 km depends on data availability and on the frequency content of the gravity field in the area of interest. Errors in this bandwidth are generally more pronounced in locations where high-resolution information is not available and less pronounced in areas with smooth topography. Their average behaviour (over the globe) is formalized as a function of frequency by the lighter lines in Figure 6.1a and Figure 6.1b. They show the global error degree amplitude spectra (which can be likened to 1-σ values) for gravity anomalies derived using the EGM96 model. They were computed as the square root of the error degree variance equations given in Lemoine et al. (1998) by F. Bayoud who made them available to the author for use in this chapter. Consider, for example, that the EGM96 can only provide an estimate of the gravity anomaly with a cumulative error of less than 1 mGal for half-wavelengths longer than about 300 km. As well as accumulated errors within the bandwidths represented by the model, there is also

a considerable misrepresentation of the gravity field (called the omission error) that results from the fact that current global models are limited in resolution. Consider Figure 6.2a and Figure 6.2b, for example, that show the true gravity field in an area of the Canadian Rocky Mountains (approximately 100 by 100 km in size) and the estimate of the same field that is provided by the EGM96 model (evaluated to degree and order 360). Figure 6.2c is a plot of the difference between them and clearly demonstrates that the band-limited nature of the model can result in omitting high-resolution signals with amplitudes in the order of ± 50 mGal. More information about the EGM96 can be found in Lemoine et al. (1998) and the process of deriving global models is summarized in Rapp and Pavlis (1990).

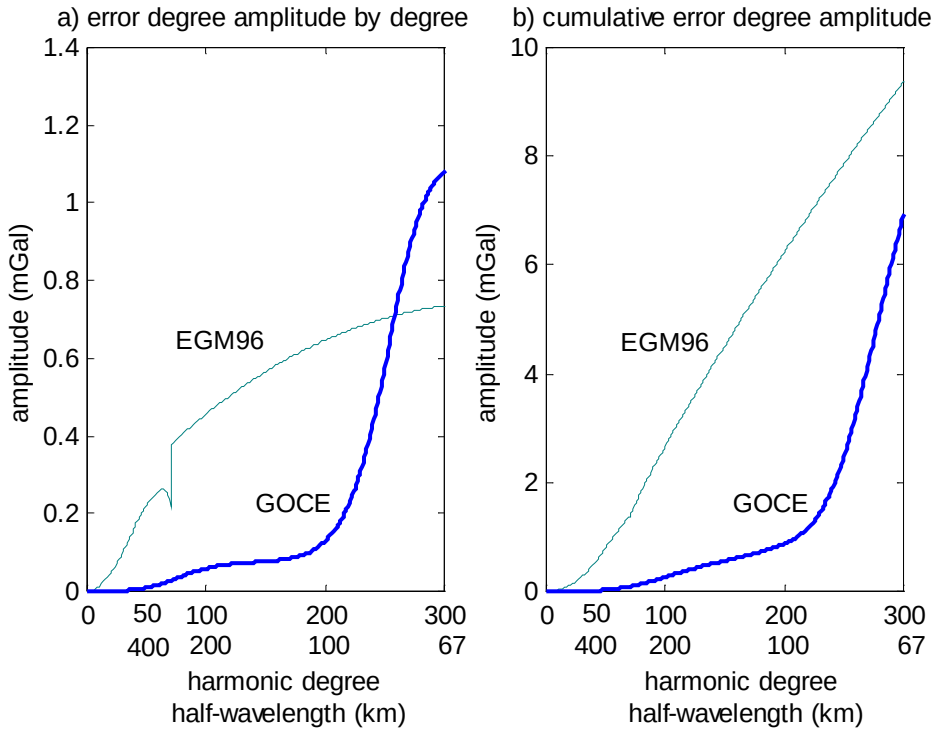

Figure 6.1 Error degree amplitude values for current and future global models

In the near future (e.g. 3 to 10 years from now), low-resolution information will come almost entirely from data provided by geophysical satellite missions that are dedicated to global gravity field estimation. These include the CHAMP, GRACE and GOCE missions that are based on hi-low satellite-to-satellite tracking, low-low satellite-to-satellite

tracking and satellite gradiometry, respectively. While these promise unprecedented improvements in the quality (and therefore useful resolution) of global models, the need for complementary medium and high-resolution information will not only remain, but increase. The reader is referred to NRC (1997) and ESA (1999) for detailed comparisons of current global models with those expected from future satellite missions. As an example, the darker lines in Figure 6.1 show how accurately the future GOCE mission is expected to provide the gravity anomaly. The degree amplitude spectra derived from the degree variance equations given in ESA (1999) are shown.

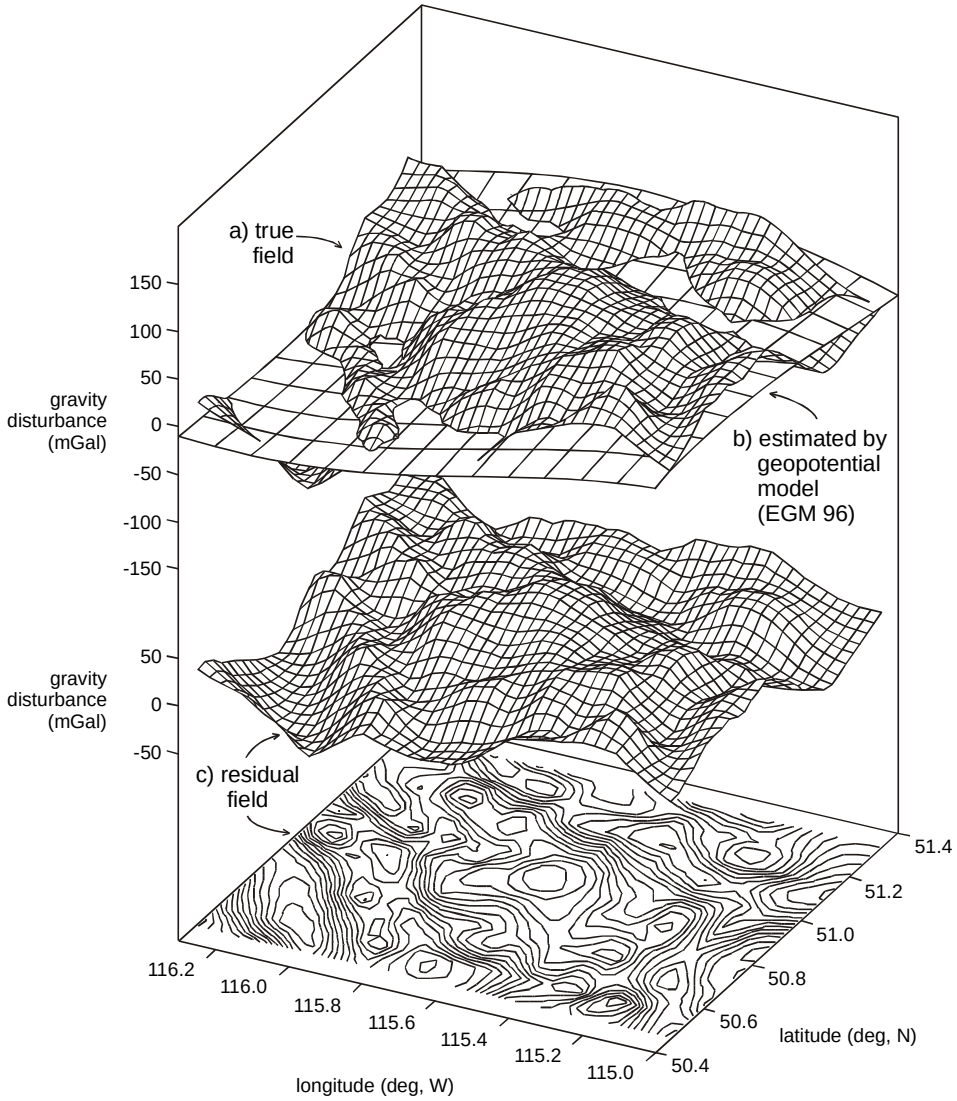

Figure 6.2 Resolution of the gravity field and of the data coming from the EGM96

It is evident from Figure 6.1 that future satellite missions will offer significant improvement over the EGM96, both in terms of accuracy and useful resolution. The models based on the GOCE data are expected to be able to provide an estimate of the gravity anomaly with a cumulative error of less than 1 mGal for half-wavelengths as short as 100 km. For higher frequencies than this, their quality degrades rapidly, however, implying that data from surface and airborne gravimetry will continue to play an important role in the estimation of the medium and high-resolution components of the gravity field. This is the subject of the next sub-section.

6.1.2 Information about the medium and high-resolution content

It was made clear above, that a role of data from surface and airborne gravimetry is to complement the information represented by (current and future) global models. This sub-section compares the two in an effort to highlight the benefits of the airborne approach.

The quality of the information that can be derived from surface data depends largely on the spacing of the available samples. For geodetic work, measurements are typically made by federal government agencies and their availability is relatively good on the continents of North America, Europe, Australia and Japan. Because of the cost and logistics of making and maintaining a database of such measurements, coverage varies in other areas of the world, from reasonable to none at all. Because of the relatively high speed, low cost and logistic simplicity of airborne surveying, it has been enjoying increased use for large-scale surveys in place of surface campaigns. For example, airborne data is currently available in selected parts of the world such as Greenland, the Arctic and Antartica (see Brozena (1992) and Brozena et al. (1997) for examples).

As an example, consider Figure 6.3a and Figure 6.3b that show the location of available surface and airborne data in the same 100 x 100 km area that was shown in Figure 6.2. The surface data shown are those available in the database of the Geodetic Survey Division of Geomatics Canada (GSD) and the airborne data was collected on September 10, 1996 by the University of Calgary (see Appendix A.1). The figure is included to demonstrate the typical characteristics of the samples available for each of these data types. Notice first, that because of ease of accessibility, much of the surface data in Figure 6.3b is collected along roadways (represented by the white dots). Alone, this

causes several problems in terms of sampling, including the obvious issue of the gaps between the roadways and the fact that such samples do not typically represent the gravity field at high elevations such as mountain tops. To compensate for these problems, the gaps were filled by helicopter-based surveys (represented by the shaded dots). Typically, a surface data set is collected in a number of campaigns that can span a number of years or decades. The data are typically maintained in a database and interpolated to a grid before use, e.g. in Canada they are also available on 2 and 5 arc minute grids. At best, it would take approximately 20 twelve-hour days to collect surface data of the density shown in Figure 6.3b, using a helicopter. This estimate is based on experience gained by the author during a recent helicopter-based field campaign for densifying the surface gravity measurements in the Kananaskis area (see Appendix A for details). On foot or even by road vehicle, the same density of points would take months to collect and may even be impossible, depending on the topography and vegetation.

Consider also the data from the airborne survey shown in Figure 6.3a. There are several important differences, including the fact that they were collected at a constant height, on a near-continuous grid and over a period of only six hours.

In both the airborne and terrestrial cases shown in Figure 6.3, the low-resolution information that can be derived from the data is limited by the geographical extent of the survey. For a survey within a given area, the maximum half-wavelength that can be resolved is given by the size of the area. For example, it is 100 km for both examples given here. On the other hand, the high-resolution information that can be derived is different for each. For the surface data, it is given by the average spacing between the samples. In this case, it varies from less than 1 km in some areas on roadways to 10 km in the areas between the roads. For the airborne case, it depends on the spacing between flight lines and on the resolution of the gravity system, i.e. the product of the aircraft velocity and the cut-off frequency of the low-pass filter used in the computations. Here, the flight lines are spaced 10 km apart, meaning that the resolution is given by a half-wavelength of 10 km in the east-west direction. If the data are filtered to 0.01 Hz and the flying speed is 100 m/s (360 km/h), then the resolution is given by a half-wavelength of 5 km in the north-south direction.

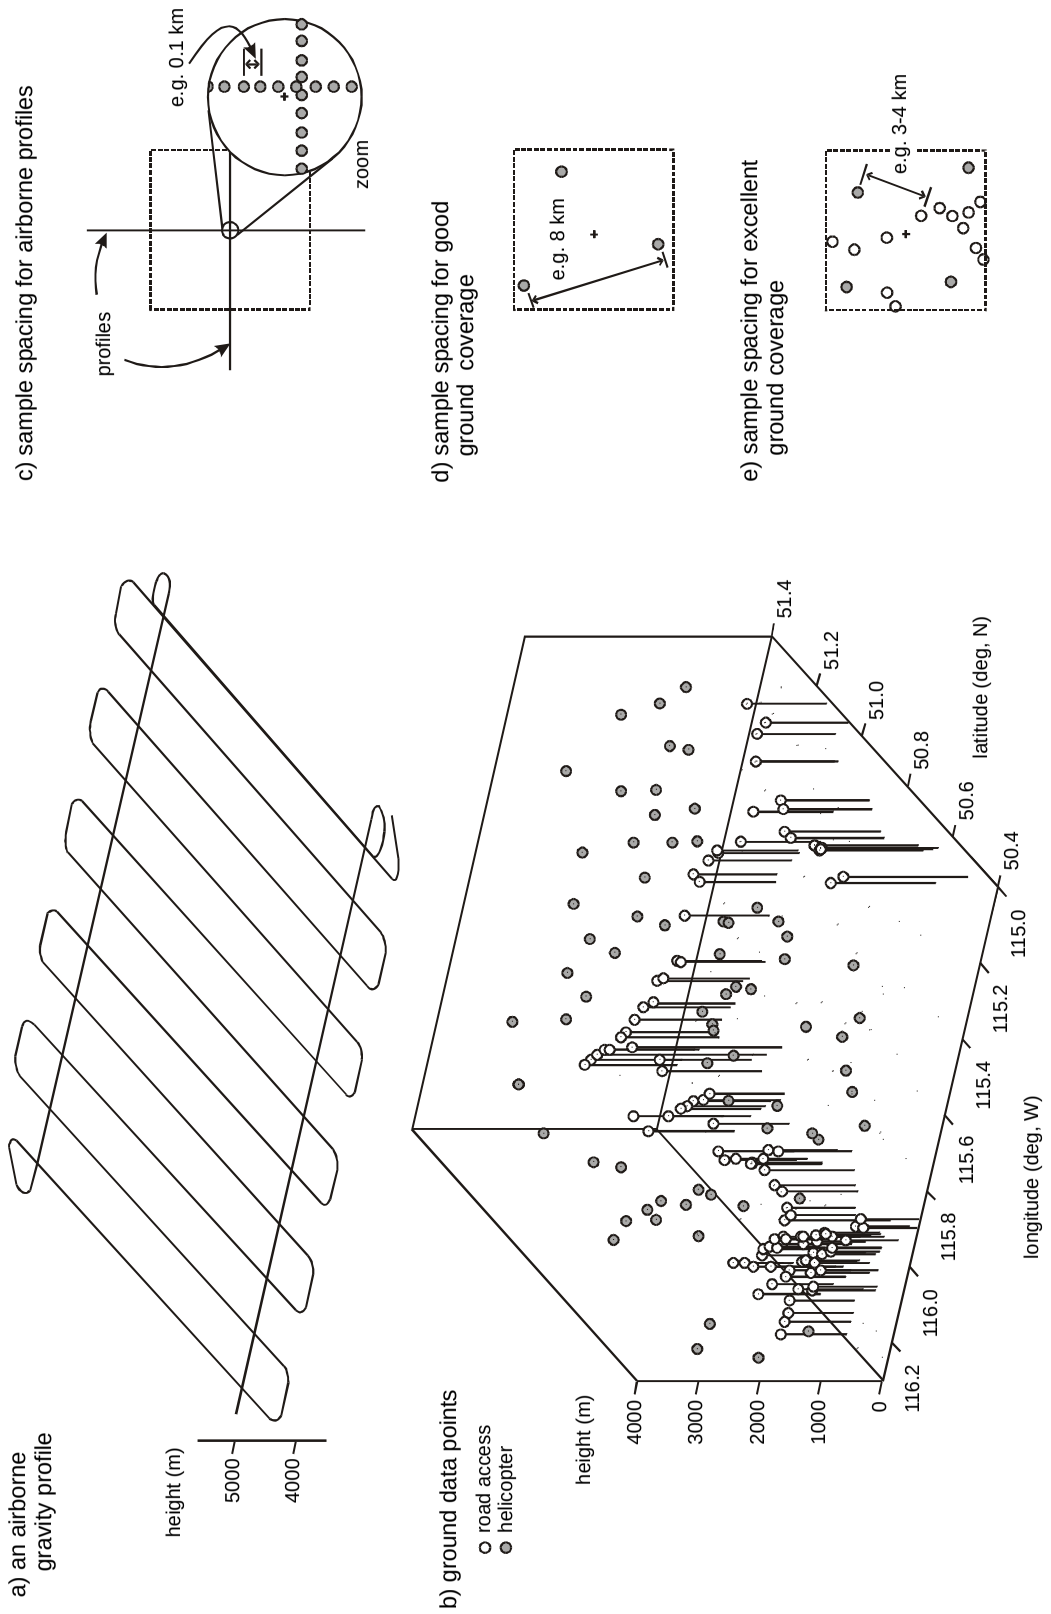

Figure 6.3 Comparing the sampling characteristics of airborne and surface gravimetry

Half-wavelength has been chosen as an expression of resolution in this chapter because it corresponds roughly to the equivalent sample spacing. For example, the maximum sample rates corresponding to the Nyquist frequency are 10 km and 5 km in the two directions discussed in this case. See Appendix C.1 for more details.

Although the factors mentioned above imply a maximum frequency content of the signal, the airborne data has a significant advantage over the surface data points because of the very high rate at which it is sampled along a profile (e.g. once every 100 m for 1 Hz data when flying at 100 m/s). This is demonstrated in Figure 6.3c for a 10 by 10 km area. As a result of this, estimates of gravity are available at practically any place along a profile without the need for interpolation. In turn, this means that if the data is needed on a grid that coincides with the flight lines, then there will be no errors due to interpolation, a possibly considerable advantage in areas with rough gravity fields. Figure 6.3d and Figure 6.3e consider the problem of gridding, for surface gravity data with varying levels of availability. The sample spacing is much poorer and the location of samples obviously implies the need to interpolate.

Before ending the discussion about sample rates, it is important to distinguish between sample rate and frequency content of the airborne data. For example, despite being highly sampled, the frequency content of the estimate is limited by the low-pass filter that is applied. Practically speaking, this means that the values at each of the points (i.e. at 0.1 km intervals) along the trajectory can be likened to mean values centered at those locations. In turn, they are mean values computed for a spatial extent that is defined by the product of the speed of the aircraft and the cut-off frequency of the low-pass filter. This is demonstrated by the solid and dotted lines in Figure 6.4. The estimate will therefore be samples of a band-limited gravity field.

Finally, it is important to realize that a potentially significant disadvantage of the airborne approach is the fact that the gravity signal is attenuated because data is collected at a higher altitude. This is demonstrated in Figure 6.5 for flying altitudes of 0.3, 1.0 and 10.0 km using the appropriate equation from Appendix C.3. It shows that the factor by which the gravity field is attenuated depends on both the altitude and resolution of the signal. Attenuation is clearly greater for higher altitudes and for shorter wavelengths, making it

more difficult to sense high-resolution information from a greater altitude. The height of the aircraft implies the need for a process by which the measurements can be represented on the geoid (or the ground). This process is known as downward continuation and the reader is referred to Chapter 1 for a list of reference material about it.

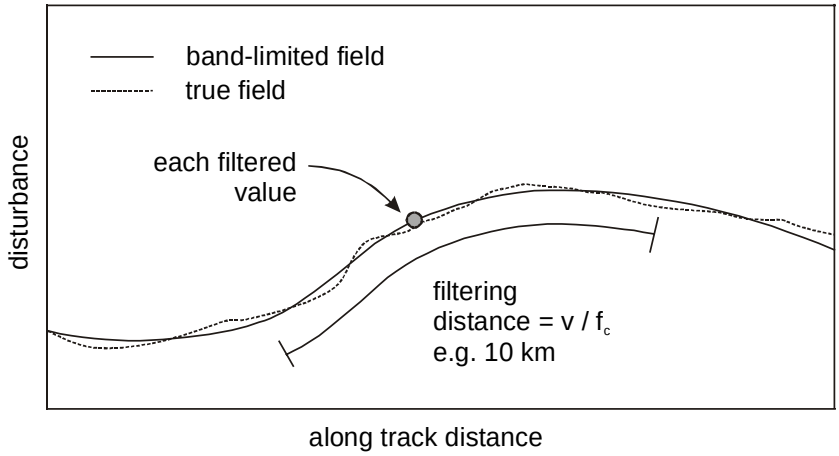

Figure 6.4 The resolution of airborne gravity data

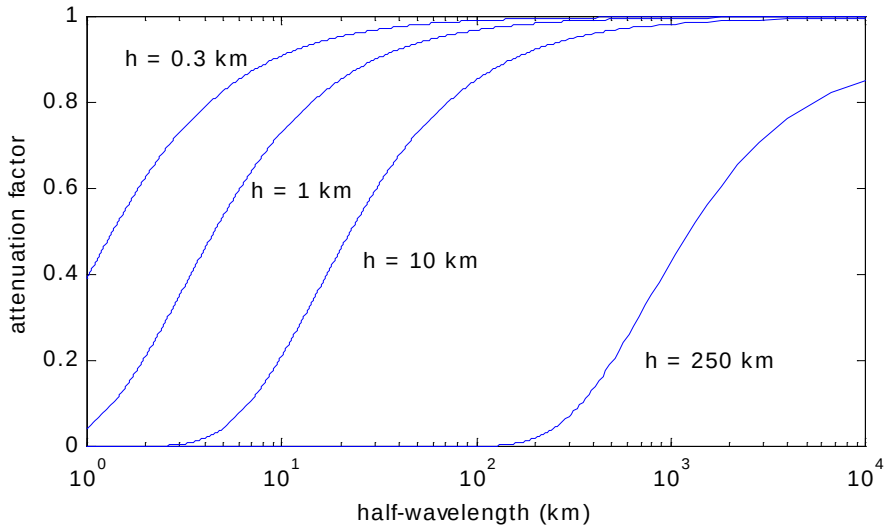

Figure 6.5 The attenuation of gravity due to altitude

Coupled with the attenuation of the gravity signal is the fact that the noise in any gravity system increases as a function of frequency (e.g. see Chapter 5 and Chapter 7). Together, these facts mean that the gravity signal becomes smaller and the noise becomes greater as the resolution of a gravity system is increased. Because downward continuation amplifies these noisy estimates, this implies a very challenging problem when trying to resolve the high-resolution components of the gravity field. In turn, researchers working to regularize the downward continuation process need to demonstrate that they are truly estimating gravity at high resolutions and not simply achieving a better fit to the gravity field by attenuating the combined signal and noise; a result that takes advantage of the fact that the signal has less power at higher resolutions.

Also shown in Figure 6.5 is the attenuation factor at an altitude of 250 km, demonstrating one of the reasons why the future satellite missions are unable to provide information for half-wavelengths below 100 km.

Note that the true field that is shown in Figure 6.2 was calculated from a combination of the surface measurements in Figure 6.3b and the surface measurements in a database that is maintained by the University of Calgary. See Argeseanu (1995) for details about this process.

6.2 Using medium and high-resolution measurements of the gravity field

Although the determination of a relative geoid and the interpretation of gravity data for the exploration of resources are beyond the scope of the research presented in this dissertation, it is useful to consider the ways in which airborne gravity measurements will be used by researchers in each. This section briefly outlines these. The goal is to emphasize the important contribution that relative measurements of the gravity field make in each case.

6.2.1 Relative geoid determination

As outlined in Sideris (1997), a local geoid is usually determined from medium and high-resolution measurements of the gravity field using the remove-restore technique. Because such measurements are usually of limited geographical extent, this approach does not attempt to compute the low-frequency components of the geoid. Instead, it uses the

measurements to compute the relative component, N_r, and obtains the low-resolution component, N_{GM}, from a global geopotential model. The total geoid undulation, N, is therefore given by the following equation:

$$N = N_{GM} + N_r . \tag{6.1}$$

For simplicity, an example is borrowed from Li (2000), in which the desired component is computed from relative gravity disturbances according to the Hotine integral as follows:

$$N_r = \frac{R}{4\pi\gamma} \iint_{\sigma} \delta g_r H(\psi) d\sigma , \tag{6.2}$$

where *R* is the radius of the Earth, $H(\psi)$ is an integration kernal known as the Hotine kernal function, σ indicates that the integration is done over the area of interest and the relative gravity disturbance is obtained as follows from the measured gravity disturbance values:

$$\delta g_r = \delta g - \delta g_{GM} . \tag{6.3}$$

In other words, the inputs to the relative geoid computation are estimates of the relative gravity disturbance that are isolated from the measured disturbances, δg, by subtracting the low-frequency component that comes from the geopotential model, δg_{GM}. As mentioned in the last section, these measurements may come from either surface or airborne gravimetry campaigns. Note that the relative gravity disturbance is the quantity shown in Figure 6.2c.

This and other methods for determining the relative geoid from airborne data are described in detail in Li (2000) and Novak and Kern (2000) where it is pointed out that the measured quantity, δg, is expressed on a level surface and corrected for effects such as the topography and the atmosphere before being used in a process such as the one described by equation 6.2. As mentioned earlier, a downward continuation step is also involved.

A current challenge being faced by researchers is the determination of the geoid with a cm-level accuracy. Recent results presented in Li (2000) demonstrate that a relative geoid accurate to 2 cm can be obtained in mountainous areas from airborne gravity data with a

half-wavelength of 4.5 km. According to Li (2000), the task of determining an absolute geoid at the cm-level seems possible by using airborne (and surface) data together with the data from future satellite missions.

6.2.2 Resource exploration

The deliverable in most geophysical applications of gravity data is the Bouguer anomaly in the form of surface or contour plots. It is an expression of the measured anomaly on the geoid (through a process that assumes that the masses that exist between the measurement point and the geoid can be represented as a slab of constant thickness and density). As discussed in Reynolds (1997), the effect of residual topography and a long-term trend are typically removed from the anomalous field so that it represents only the local changes in gravity due to features below the surface. This makes it analogous to the quantity of relative gravity disturbance that was given in the last section and represented by the contour plot in Figure 6.2c (although in that case, the effect of the topography had not yet been removed). A relative accuracy of 1 mGal at a half-wavelength resolution of 0.5 km is often considered to be the ultimate goal for resource exploration.

6.3 The low-resolution SINS errors

The presence of low-frequency errors owing to unmodeled accelerometer biases have traditionally been viewed as a major roadblock to successfully estimating the relative gravity field for geoid determination from data collected using a SINS airborne gravity system. Recall that these were lumped together as the variable **b** in equation 2.10. Although an attempt is made to estimate them in the Kalman filter during SINS/DGPS integration, they remain largely uncompensated and as shown in equation 2.5, go directly into the error budget of the estimated gravity disturbance. As a result of typical survey conditions, the errors in the upward pointing accelerometer make up the majority of this error for scalar gravimetry. The typical biases of the SINS system used at the University of Calgary are time dependent and vary by as much as 100 mGal over a flight lasting six hours. Given that the desired performance of an airborne gravity system is at the level of 1 mGal, this is obviously a major problem. It would be ideal if they could be modeled

appropriately and removed, permitting the estimation of the gravity field in all bandwidths.

Glennie (1999) should be consulted for typical time domain plots of these errors and a description of a number of attempts to remove them. Efforts were made there to model the biases as polynomials and as Gauss-Markov processes in the Kalman filter, unfortunately without success. Their characteristics were shown to be very similar whether the SINS was static in the lab or being used in flight, leading to the conclusion that they are most likely due to an unmodeled dependence on the temperature of the system. This conclusion is supported by several facts. The first is that the errors always reach a near steady-state value after a warm-up period of several hours (an unacceptable time delay in practice). The second is that there are several analyses that show negligible biases for data collected using similar accelerometers under more strict temperature control. See Ferguson and Hammada (2000) and Glennie (1999) for examples.

6.3.1 Estimating the relative gravity field in a more direct manner

The possibility of estimating the relative gravity field from data collected with the SINS used by the University of Calgary is discussed in this section. Conceptually, the method proposed herein uses high-pass filters to derive the relative information from airborne data without estimating or compensating for the accelerometer biases. In other words, instead of requiring that the measured gravity disturbance, δg, be corrected for the estimated biases before being used in equation 6.3, the method directly estimates the relative gravity disturbance, δg_r. Although the experiments described in the last section demonstrate that it is difficult to describe the biases as a function of time (and therefore space), their behaviour is relatively easy to characterize in the frequency domain. Consider Figure 6.6 that shows an estimate of the frequency spectrum of the accelerometer errors that was obtained empirically using data collected by the upward pointing accelerometer during a static session (in the aircraft, with the engines on). Note that the x-axis of the figure is presented using a logarithmic scale. The spectra derived from other data sets are very similar. The majority of the error is clearly very low-frequency in nature, implying that it can be removed by high-pass filtering the estimated

gravity disturbance. Although doing this will obviously make direct estimation of the absolute gravity field impossible, it permits the estimation of relative gravity, free from the effects of the biases. As shown, the low-frequency errors mainly influence the signal in the bandwidth below 0.00025 Hz. For a flying speed of 100 m/s, this implies that the system is capable of accurately measuring the relative gravity disturbance, δg_r, for half-wavelengths shorter than 200 km. Notice that the remaining accelerometer errors exhibit white noise behaviour except in the band between 0.00025 and 0.0005 Hz, implying that it should perform best for half-wavelengths shorter than 100 km.

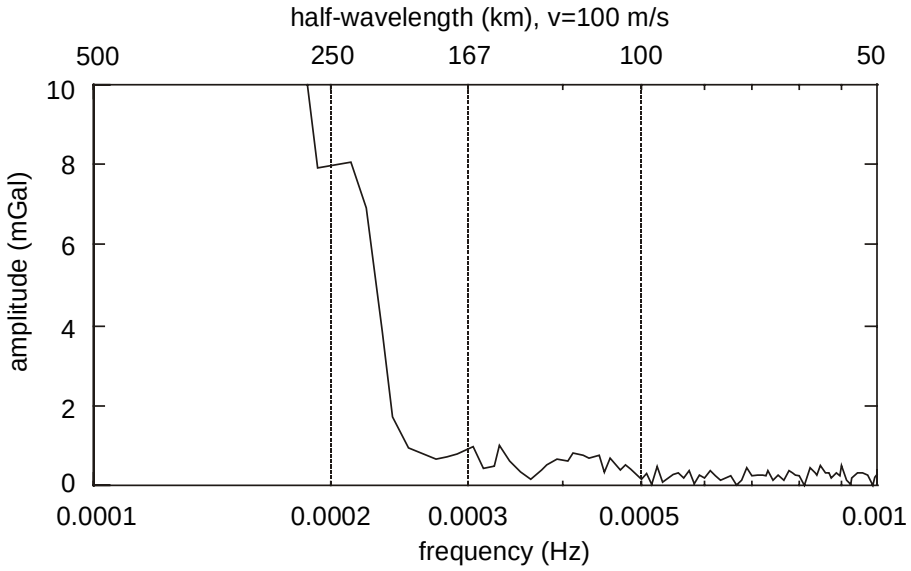

Figure 6.6 Frequency spectrum of the accelerometer bias errors

Consider Figure 6.7 that demonstrates this using real data. Figure 6.7a shows the true and estimated gravity disturbances for the first four flight lines surveyed on September 10 of the Kananaskis campaign. Figure 6.7b shows the results of band-limiting each so that they no longer contain information with half-wavelengths longer than 100 km. The level of agreement is significantly better in this bandwidth. In order to quantify the low-frequency errors affecting the performance of the SINS system for relative gravity determination, large portions of the data from the Kananaskis and Alexandria campaigns were processed in the same way for various bandwidths. This included the first 11 flight lines from each of September 9 and 10 of the Kananaskis campaign and the 10 flight lines

flown on April 19 and May 4 of the Alexandria campaign, making a total of 32 flight lines and nearly 13 hours of on-line data. Each data set was band-limited using 6 different high-pass filters designed to pass only frequencies corresponding to half-wavelengths shorter than 50, 100, 150, 200, 250 and 300 km when the flying speed was 100 m/s. Because the speed in the Alexandria campaign was close to 50 m/s, these filters correspond roughly to half-wavelengths of 25, 50, 75, 100, 125 and 150 km when applied to that data. The results are summarized in Figure 6.8, Table 6.1 and Table 6.2. Figure 6.8 shows the average offset and slope that exist between the estimates and the reference for each of the bandwidths mentioned above. The average offset values are calculated for each campaign by computing the mean difference between the estimate and the true solution for each of the flight lines flown in that campaign, within the bandwidths of interest. The values of average slope are calculated in the same way, based on the absolute value of the slope of the line that best fits the error in the estimate. The spatial resolution for the Alexandria campaign has been normalized to 100 m/s in Figure 6.8 so that the data from both campaigns can be directly compared. It is clear from this that the behaviour of the time dependent accelerometer biases is very similar for both campaigns. It is also clear that they are much larger for half-wavelengths longer than 200 km, confirming that the behavior seen in Figure 6.6 is also typical in the airborne environment. The performance of the system is obviously significantly better for half-wavelengths of 200 km and shorter. For example, Figure 6.8a demonstrates that the contribution of the accelerometer biases to the total error budget is less than 1 mGal for frequencies greater than 0.00025 Hz (which corresponds to half-wavelengths less than 100 km and 200 km for the Alexandria and Kananaskis data sets, respectively). Table 6.1 and Table 6.2 demonstrate the performance of the system for selected bandwidths within this range. The values given in these tables are the standard deviations of the agreement of the estimates with the reference within selected bandwidths. Agreement is clearly better for higher resolutions, i.e. when more of the biases have been removed. Note that the numbers in Table 6.2 can be roughly compared to other results published by the University of Calgary (where a best-fit line is typically removed from the data), by comparing the column corresponding to a half-wavelength equal to the length of the

flight line. For example, the numbers in Table 6.2 that represent the gravity field for half-wavelengths of 100 km can be roughly compared to the results presented in Glennie and Schwarz (1999) where a line of best fit was removed from the data. The agreement of those presented herein is about 0.3 mGal worse than those mentioned above for the same flight lines, perhaps implying that the high-pass filters being applied herein to the reference and the estimates have some imperfections. As shown later (in Figure 6.9), this difference might also be due to residual biases that would be distributed more evenly over a flight line if a line of best fit were removed. If the latter is true, then removing a line of best fit might be a slightly optimistic means of evaluating the relative performance of the gravimeter.

Note that the low-pass filter used to generate the data in Table 6.1 and Table 6.2 has a cut-off frequency of 0.011 Hz, meaning that the shortest half-wavelengths represented are approximately 2.3 and 4.5 km, respectively. Note also that the better performance seen for the Alexandria campaign is due to a combination of factors including slightly better GPS conditions, much more stable aircraft dynamics and the fact that the gravity field is significantly smoother in the Alexandria test range than over the mountains in the Kananaskis. The former two points are discussed in detail in Chapters 5 and 7, respectively.

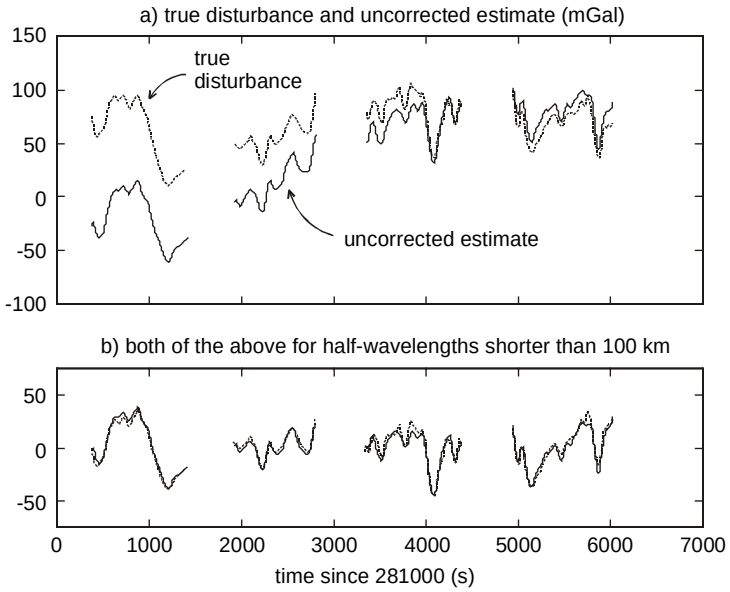

Figure 6.7 The quality of the solution for half-wavelengths less than 100 km

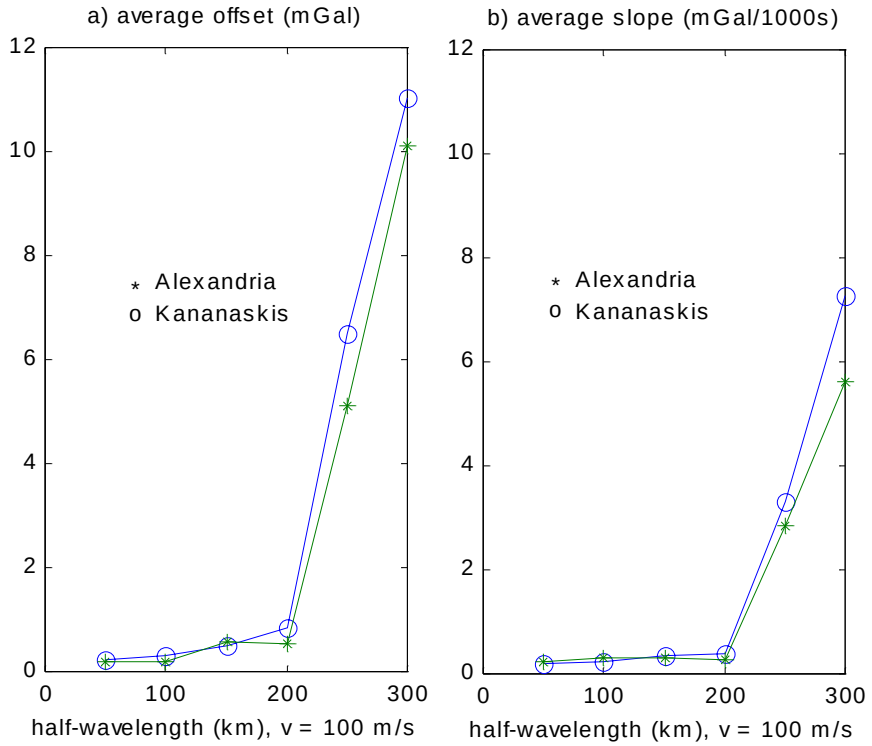

Figure 6.8 The effect of the low-frequency errors by bandwidth (normalized to a flying speed of 100 m/s)

Table 6.1 Performance for selected bandwidths for the Alexandria campaign

data set	line	minimum frequency (and corresponding half-wavelength)		
		0.00025 Hz (100 km	0.0005 Hz 50 km	0.001 Hz 25 km)
Alexandria April 19	T1001	3.7	3.4	3.3
	T1002	2.3	2.1	2.1
	T1003	2.7	2.6	2.5
	T1004	1.1	1.1	1.0
	T1005	2.1	2.0	2.0
	T1006	1.6	1.4	1.3
	average	**2.3**	**2.1**	**2.0**
Alexandria May 4	T1001	1.5	1.4	1.3
	T1002	1.2	0.9	0.9
	T1003	1.9	1.5	1.5
	T1004	1.1	1.1	1.0
	average	**1.4**	**1.2**	**1.2**

Table 6.2 Performance for selected bandwidths for the Kananaskis campaign

data set	line	minimum frequency (and corresponding half-wavelength)		
		0.00025 Hz (200 km	0.0005 Hz 100 km	0.001 Hz 50 km)
Kananaskis Sept 9	1	3.8	2.6	2.5
	2	3.0	2.7	2.4
	3	6.2	6.2	4.2
	4	3.2	2.9	2.7
	5	3.1	3.2	2.7
	6	2.9	2.5	2.5
	7	4.7	3.9	3.0
	8	4.2	3.7	3.5
	9	4.6	3.8	2.8
	10	4.1	3.4	2.9
	11	3.9	4.0	3.6
	average	**4.0**	**3.5**	**3.0**
Kananaskis Sept 10	1	4.6	4.1	3.5
	2	3.2	2.8	2.1
	3	4.7	4.9	3.9
	4	5.9	3.5	3.2
	5	4.9	3.8	2.4
	6	3.6	3.3	3.4
	7	3.2	2.9	2.4
	8	3.2	3.1	2.4
	9	2.5	2.3	2.3
	10	4.3	4.1	3.3
	11	4.1	3.7	3.7
	average	**4.0**	**3.5**	**3.0**

6.4 Relating estimates of the relative gravity field to the true gravity field

Generally speaking, relating estimates of the relative gravity field to an absolute scale requires some knowledge of the gravity field for resolutions below those represented by the relative information. On one extreme, if the sensors exhibit no bias-like behaviour, this relationship can be established using the true value of gravity for any point along the survey. In practice, this might be done using an absolute gravity control point at the start or end of a survey (on the tarmac, for example). If the sensors exhibit a linear drift (i.e. one that can be represented by a best-fit line), then absolute information can be obtained

if the true field is known for at least two points along the trajectory. In practice, the start and end points are typically used for this. The accuracy of the estimated absolute gravity is obviously directly dependent on the accuracy of the control point values used and the nature of residual biases. These cases are typical of surveys flown using 2-axis platform systems and the inertial platform system developed at Sander Geophysics Ltd., where the drift of the sensors is kept at or below the level of 1 mGal for an entire flight.

In the case of a SINS-based gravimeter, values of the gravity field are unfortunately required at a spatial resolution corresponding to the longest half-wavelength represented by the relative information. For example, if the data is band-limited such that it only contains half-wavelengths shorter than 100 km, then low-frequency information about the gravity field is required with half-wavelengths longer than 100 km. This can be likened to needing a sample of the absolute gravity field every 100 km. Figure 6.9 demonstrates this for two flight lines from the Kananaskis campaign. It shows values of the relative gravity disturbance (for half-wavelengths shorter than 100 km) determined using the system, and the quality with which the absolute field can be estimated when accurate samples are available at the endpoints (shown by the circles). Note that even when interpolating over only 100 km, low-frequency errors within this bandwidth can cause the solution to deviate from the truth (e.g. toward the middle of the flight line in Figure 6.9b). Although having samples of the true gravity field every 100 km can be used in this way, it is entirely impractical; one will not have such information in areas without very good ground gravity coverage. The following sections discuss two alternative means of solving this problem.

6.4.1 Using a geopotential model

If a global geopotential model is available that has sufficient accuracy in the required bandwidth, then it can be used to complement the relative information that is collected. Although this is not a straightforward concept if one wishes to consider the stochastic nature of the estimates provided by these two data sources, it can be accomplished from a deterministic point of view rather simply. This is especially true, considering the temporal nature of the data collection process in airborne gravity that lends itself well to one-dimensional signal processing.

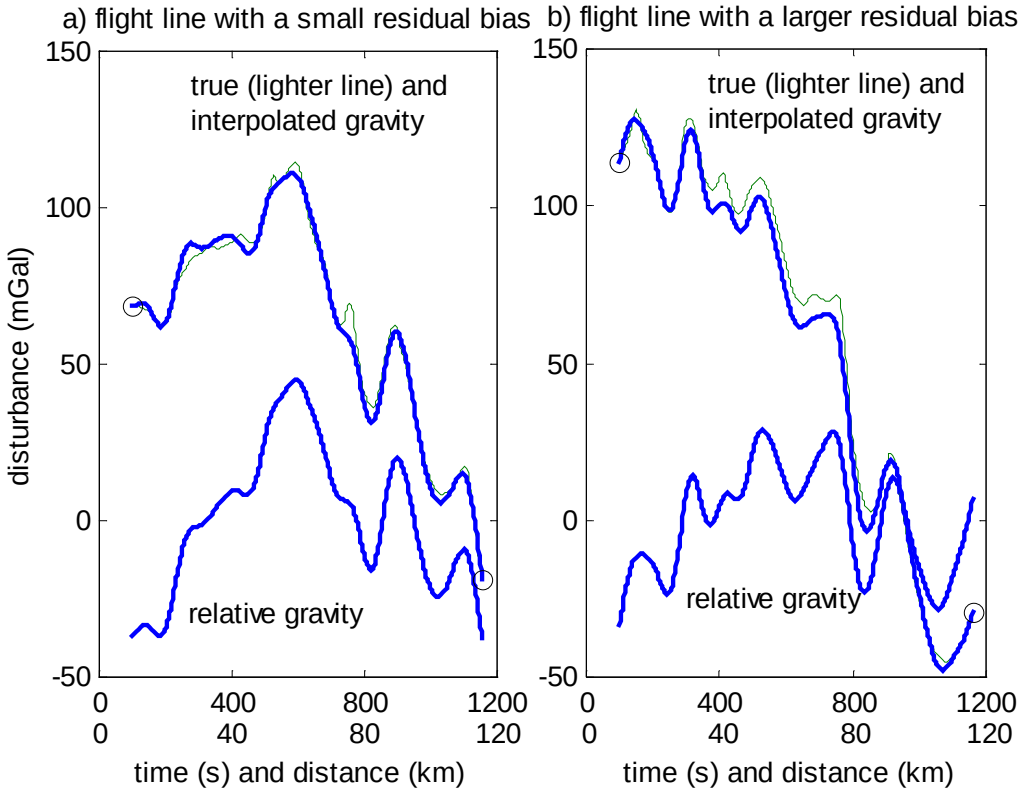

Figure 6.9 The interpolation of relative gravity using values at the endpoints

Consider Figure 6.10 that shows the magnitude of the frequency response of two complementary filters that can be used for this purpose. The first is a low-pass filter that is used to band-limit the information coming from the geopotential model to half-wavelengths longer than 100 km. The second is the high-pass filter that was used in Section 6.3.1 to extract relative information for half-wavelengths shorter than 100 km. These can be applied to the two time series to generate an estimate of the gravity field for all resolutions. This idea is attractive because it does not require a-priori knowledge of the gravity field, other than the global model. It also has several advantages over a crossover adjustment, including the facts that it does not require that tie lines be flown, it does not depend on any assumption about the nature of the accelerometer biases (e.g. linearity), and it does not limit the longest half-wavelength of the system to the length of the flight lines. (The crossover adjustment is discussed briefly in the next section.)

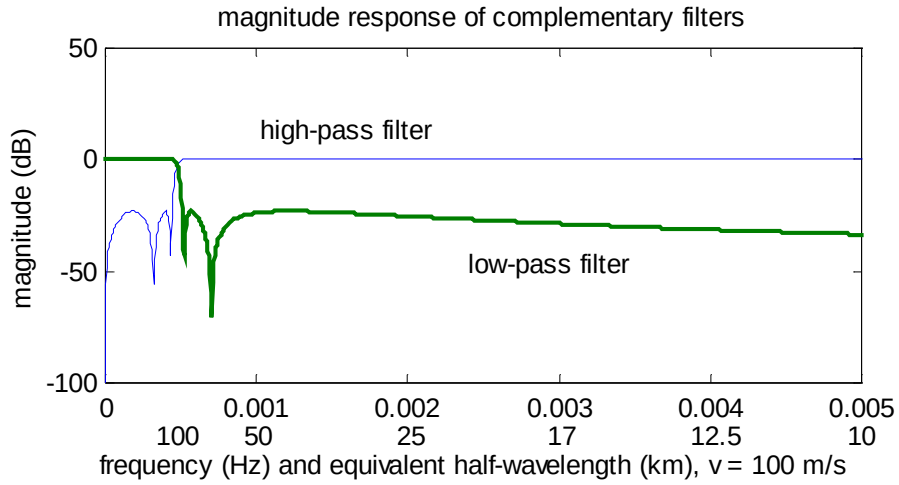

Figure 6.10 Complementary filters to combine high and low-resolution gravity information

For this method to be successful, the errors in the geopotential model and airborne data need to be small enough within the bandwidths in which they are each being used. Unfortunately, based on the results presented in Section 6.3, it is unlikely that the SINS system being used by the University of Calgary will be able to deliver accurate measurements within the range of frequencies required for combination with current global models (recall the plots of the accuracy of the latter that were presented in Section 6.1). To test this, selected data from the Kananaskis and Alexandria campaigns were processed in this way using the EGM96. As a worst-case example, consider the results obtained for the first four flight lines from the Kananaskis campaign that are shown in Figure 6.11 (this is the same data that was shown in Figure 6.7). It is clear that the EGM96 is not always accurate enough for this purpose. Biases as large as 20 mGal exist between the resulting solutions and the true gravity field in the Kananaskis region. Although the situation is somewhat better in areas with smoother topography, the errors can still be significant. For example, the average offset and slope between the true field and the global model for the data collected on May 4 of the Alexandria campaign are 5.8 mGal and 0.081 mGal/km, respectively, which may be significant depending on the application. It is interesting that all of the flight lines tested in the Alexandria area have a very similar bias, e.g. the standard deviation of the biases is only 0.9 mGal, implying that the biases might be dealt with by the removal of a single constant bias from the estimates for the whole area.

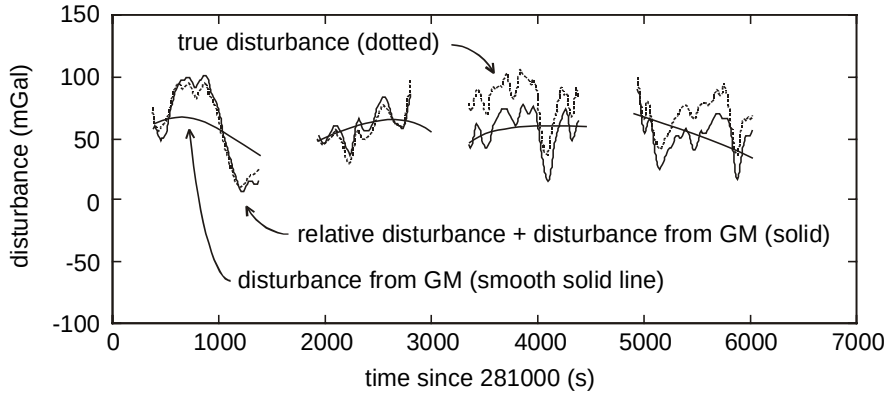

Figure 6.11 Complementary filtering in a mountainous area using the EGM96

These observations echo the call for better accelerometers that was made in Glennie (1999). On the other hand, although it might not be applicable at present in all situations, it is reasonable to expect that this method should constitute a very feasible approach once the data from the dedicated satellite missions becomes available. Recall from Figure 6.1 that models based on data from the GOCE mission are expected to provide an average accuracy of 1 mGal for half-wavelengths longer than 100 km. If this is achieved, it should be very useful in this complementary filter approach, except perhaps over very rugged terrain. This will have to be tested once that data is made available.

Finally, note that the author does not expect that a stochastic approach to combining current geopotential models with the relative airborne information from the LRF-III SINS would yield significantly better results than the deterministic approach presented here. This is because the fundamental problem is a lack of useful information in one part of the spectrum, which is independent of the understanding one has of the error characteristics of each source of information.

6.4.2 Using a crossover adjustment

In the absence of accelerometers with good temperature control and a global model with sufficient useful resolution, a well-known method called the crossover adjustment can be used to make the relative information consistent within a given area. Glennie and Schwarz (1997) demonstrated the use of the crossover adjustment as a method of dealing with the low-frequency errors present in the SINS gravimeter. For relative gravimetry, it

plays the role of making the data consistent, effectively estimating a consistent set of endpoint values for each flight line. It is an adjustment that takes advantage of the fact that the gravity disturbance is the same for a given point in space, independent of the time at which it is sampled. By flying a number of flight lines in the direction perpendicular to the main lines, a sufficient number of such points can be obtained to allow for the estimation of a residual offset and slope for each flight line. If the behaviour of the biases over the period of a flight line can be characterized by a line of best-fit, this should result in good levels of agreement between the solutions at the points where the flight lines intersect. In turn, this will mean that the relative information is fit to an arbitrary plane in space that can either be related to the absolute gravity field using known values for at least three locations in the test area, or presented as relative information by removing a plane of best fit. This is demonstrated in Figure 6.12 for the data collected during the Kananaskis campaign. The program used for this adjustment was written by C. Glennie. The agreement of these adjusted solutions with each other at the crossover points has already been demonstrated in Glennie (1999) to be at the level of 2.3 mGal (implying a standard error of 1.6 mGal coming from the data collected on each day). The figure is only included here for completeness and so that the reader can compare the concept of the resulting relative field to the relative field that was presented in Figure 6.2c.

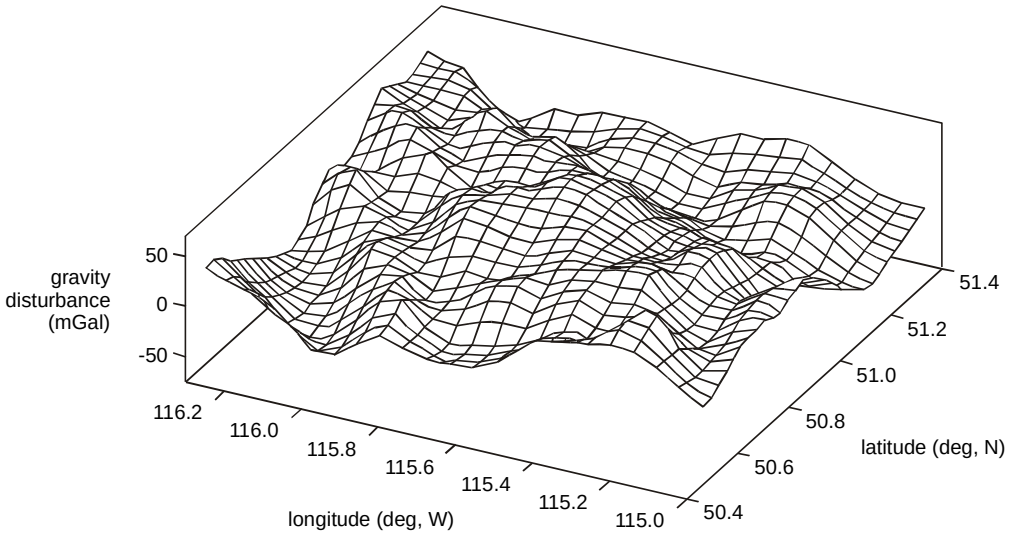

Figure 6.12 The residual field estimated by the SINS based gravimeter

In the case that the biases cannot be modeled by a line of best fit over the length of the flight lines being used, incorporation of a higher order polynomial fit into the crossover adjustment might be appropriate. As shown earlier, this can be the case for flight lines between 100 and 200 km in length. Discussion of this approach is left to Kennedy (1999), who has already implemented it with success.

6.5 Summary

This chapter has demonstrated the role that airborne gravimetry plays in the collection of medium and high-resolution information about the gravity field. It was compared to the method of data collection by land-based surveying and demonstrated to have several significant advantages. These include the speed, uniform accuracy and homogeneity of the data collection process. They also include the fact that if the flight pattern coincides roughly with the computation grid, then errors due to interpolation are negligible.

The accuracy with which the relative gravity field can be determined using the SINS system currently being used at the University of Calgary has been quantified for a variety of bandwidths. The system is currently only applicable for half-wavelengths up to 200 km, and performs best for half-wavelengths shorter than 100 km. Practically, these correspond to flight lines that are 200 and 100 km in length, respectively.

Methods of relating the measurements of the relative field to the absolute field have been discussed, including the proposal of a simple method for combining the medium and high-resolution information coming from airborne surveying with the low-resolution information coming from global models of the Earth’s gravity field. Although this method is only currently being applied for SINS systems with better accelerometers, it will be more generally applicable once global models from future dedicated satellite missions become available.

7 Natural Motion and Other High-Resolution SINS Errors

It is recalled from Chapter 2 that because of the SINS errors, the ideal flying conditions for airborne gravimetry occur when the survey aircraft travels with a constant velocity and attitude. This is an idealization that is only approximated in practice, however, due to effects such as vibration, turbulence, wind and changes in speed. Although it has long been recognized that the response of the aircraft to these perturbations results in errors in the estimates of the gravity field, their causes and characteristics are not well understood by airborne gravity researchers and there has never been a reliable means of suppressing their effects. Given this, the specific objectives of the current chapter are to:

a) present a theoretical study of the natural response of an aircraft to such disturbances,
b) demonstrate the actual motion experienced by typical survey aircraft during recent flights and estimate the effect it has on airborne gravimetry, and
c) discuss methods of removing the effect of this motion, including the proposal of a method of removing a major part of it.

As a starting point, consider Figure 7.1 that is included to demonstrate the performance of the system used by the University of Calgary, both under static conditions and during an airborne survey. The solid line is an error spectrum that was derived from SINS data collected over a period of several hours, while the SINS was static. The dotted line is a spectrum of the SINS errors, computed for the dynamics experienced during the Kananaskis campaign. (The derivation of the dynamic error spectrum will be discussed later in Section 7.3) Note that because the low-frequency errors were treated in detail in the last chapter, they will not be discussed herein (e.g. the spectrum in Figure 7.1 was not estimated for frequencies below 0.00025 Hz). Clearly, the estimate made by the system is significantly worse when it is in motion, demonstrating that errors induced by the dynamics play a significant role. Because of this and because the relationship between the SINS errors and the dynamics is not generally well understood, the band above 0.016 Hz

is not currently considered operational. Concentration in this chapter is therefore on defining this relationship and discussing methods of reducing the effects of the errors.

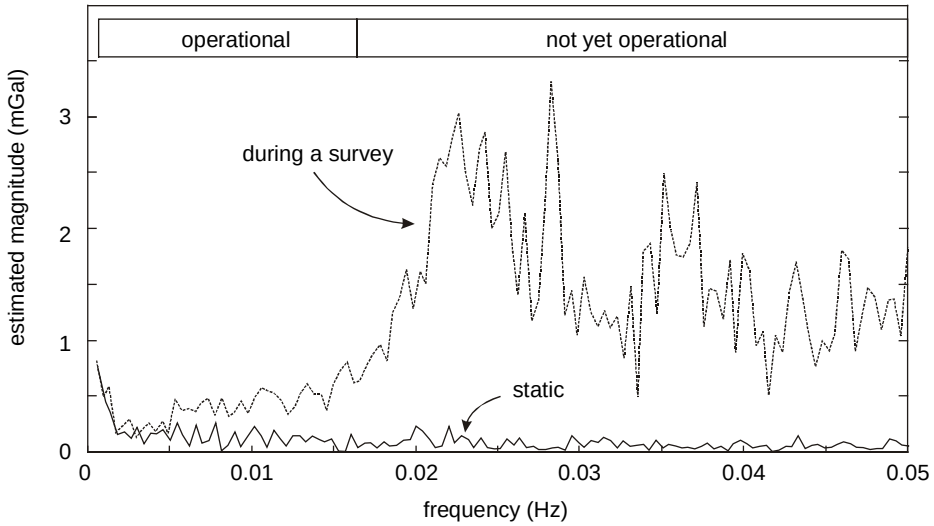

Figure 7.1 The SINS error spectrum while in motion and when static

Concentration is on the bandwidth between 0.01 and 0.05 Hz, which corresponds to spatial scales with half-wavelengths shorter than 5 km and longer than 0.5 km for typical flying speeds. As such, it serves the broader objective of increasing the performance of current airborne gravity systems within the bandwidth required for resource exploration.

The chapter is divided into four sections. The theoretical background required to understand the natural response of an aircraft to the perturbations that occur in flight is presented in Section 7.1. The major objective is to demonstrate which components of this characteristic motion can be expected to cause errors in the portion of the spectrum of interest in airborne gravimetry. In principle, this means identifying components of the motion that have frequency content between 0.01 and 0.05 Hz. In meeting this objective, it is explained why the well-known *Phugoid motion* is a limit to high-resolution gravity field estimation. Available SINS/DGPS navigation data is then used in Section 7.2 to demonstrate and analyze the actual motion that was observed for survey aircraft during two recent flights. Emphasis is on demonstrating the amount by which the natural motion is reduced by the interaction of a pilot and autopilot. The SINS errors that are induced by these residual dynamics are then characterized in Section 7.3. This is done in the

frequency domain using larger quantities of real data. Finally, practical ways of reducing these errors are discussed in Section 7.4, including the proposal and implementation of a method for modeling the errors that show up as a result of the misalignment of the SINS. This constitutes the first step in a more rigorous solution to the problem.

7.1 The natural motion of an aircraft

The motion of a body can be described by a dynamic model if the forces acting upon it are sufficiently well known. Although this is not generally possible for navigation because the external forces are far too complicated (see Schwarz et al. (1989)), such an approach does provide the theoretical basis for understanding the dynamics of flight. In other words, the dynamic model corresponding to a given aircraft can be used very successfully to characterize its natural motion. This is the topic of the current section. The discussion is based broadly on standard presentations of flight dynamics that can be found in textbooks such as Babister (1980), Etkin (1982) and Bryson (1994). However, the notation and reference frames used in Schwarz and Wei (1997) are adopted, to make the material more easily accessible to readers who are familiar with the fields of navigation and airborne gravimetry.

The motion being considered is that which occurs when the aircraft deviates from a state of equilibrium (steady motion with constant velocity and attitude) as a result of disturbances such as turbulence or changes in speed or wind conditions.

7.1.1 The equations of motion

It is standard to assume that the aircraft is a rigid body that is symmetric about the vertical plane passing through its fuselage when the wings are level. In such an approach, the motion of the aircraft is modeled as a translation of its centre of mass and a rotation of the body about the centre of mass. Using Newtonian mechanics, dynamic modeling of the centre of mass is based on the time derivative of the law of conservation of linear momentum as follows:

$$\frac{\mathrm{d}(m\mathbf{v})}{\mathrm{d}t} = \mathbf{F}\,, \tag{7.1}$$

where *m* is the mass of the aircraft, **v** is its velocity and **F** is the sum of the forces acting upon it. In other words, if the forces, **F**, acting on the aircraft are known, then given some initial conditions, the solution to this differential equation provides the position, velocity and acceleration of its centre of mass.

As described in Babister (1980), the field of flight dynamics deals with describing the motion of an aircraft under the influence of the gravitational attraction, aerodynamic forces (due to the linear and angular velocity of the aircraft and to the application of controls), and propulsive forces. Given this, and the fact that the mass of the aircraft can be assumed to be constant for short time periods (i.e. the amount of fuel used is negligible), the same equation can be expressed as follows:

$$\dot{\mathbf{v}} = \mathbf{f}_a + \mathbf{f}_p + \mathbf{g}, \tag{7.2}$$

where $\mathbf{f}_a$ are the aerodynamic forces and $\mathbf{f}_p$ are the propulsive forces, $\mathbf{g}$ is the gravitational attraction (all in units of acceleration) and the dot above the variable represents the derivative with respect to time.

At this point, a reference frame (r-frame), is adopted for convenience that has its origin at the centre of mass of the aircraft and that has constant orientation corresponding to its ideal path, e.g. if the intended direction of the aircraft from its starting point is north, then the reference axis can be likened to the l-frame. The r-frame and the b-frame are both depicted in Figure 7.2a. Equation 7.2 can then be expressed in the b-frame with respect to this r-frame as follows (Bryson (1994)):

$$\dot{\mathbf{v}}^b + \boldsymbol{\omega}_{rb}^b \times \mathbf{v}^b = \mathbf{f}_a^b + \mathbf{f}_p^b + \mathbf{g}^b, \tag{7.3}$$

where $\boldsymbol{\omega}_{rb}^b$ are the b-frame components of the angular velocity of the aircraft with respect to the r- frame. The second term on the left-hand side takes into account the fact that the b-frame is rotating with respect to the r-frame. Rotations of the r-frame with respect to the i-frame have been neglected.

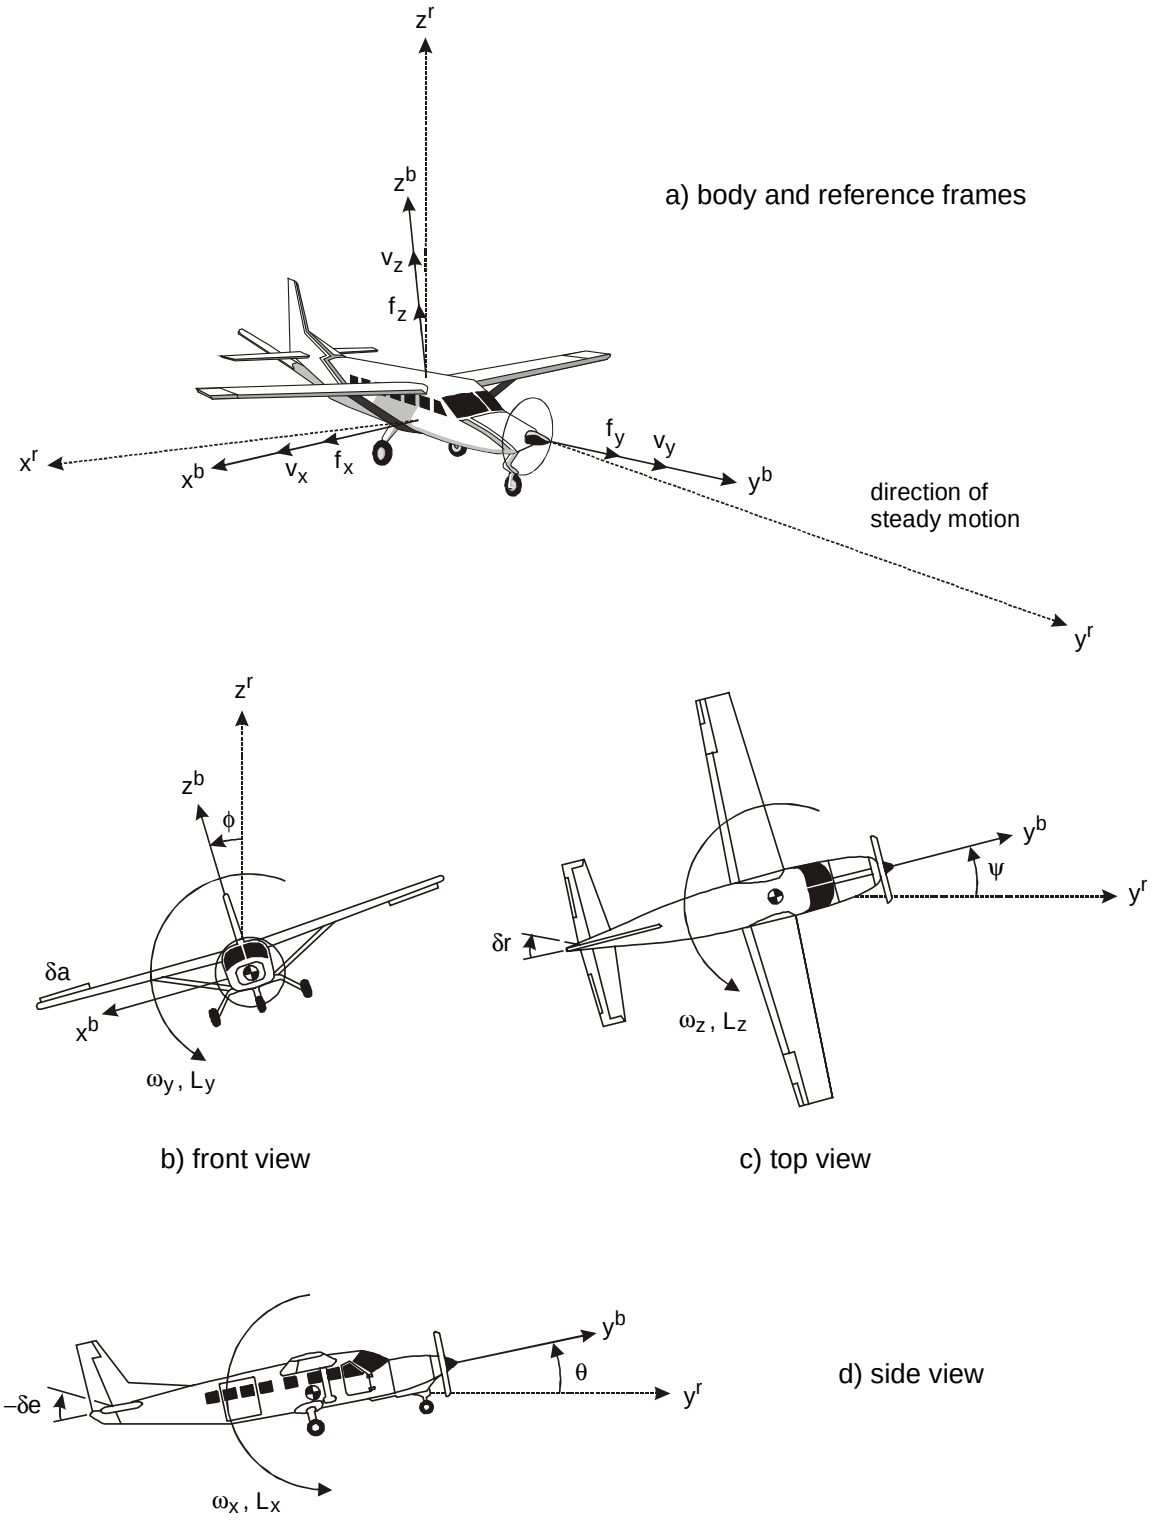

Figure 7.2 Variable and frame definitions for natural motion

As mentioned above, the rotation of the aircraft about the centre of mass must also be defined in order to completely represent its motion. This definition is based on the principle of conservation of angular momentum that relates the angular momentum, **H**, of the aircraft to the external torques that are applied to it. Assuming that the only torques present are due to aerodynamics, $\mathbf{L}_{\mathrm{a}}$, and propulsion, $\mathbf{L}_{\mathrm{p}}$, this takes the following form:

$$\frac{\mathrm{d}\mathbf{H}}{\mathrm{d}t} = \mathbf{L}_{\mathrm{a}} + \mathbf{L}_{\mathrm{p}}. \tag{7.4}$$

As was done above for equation 7.3, equation 7.4 can be expressed in the b-frame, with respect to the reference frame, to yield the following equation (Bryson (1994)):

$$\dot{\mathbf{H}}^{\mathrm{b}} + \boldsymbol{\omega}_{\mathrm{rb}}^{\mathrm{b}} \times \mathbf{H}^{\mathrm{b}} = \mathbf{L}_{\mathrm{a}}^{\mathrm{b}} + \mathbf{L}_{\mathrm{p}}^{\mathrm{b}}. \tag{7.5}$$

Note that this also directly relates the external torques to the angular rate of the body, $\boldsymbol{\omega}_{\mathrm{rb}}^{\mathrm{b}}$, because:

$$\mathbf{H} = \mathbf{I}\,\boldsymbol{\omega}_{\mathrm{rb}}^{\mathrm{b}}, \tag{7.6}$$

where **I** is the instantaneous inertia tensor given by:

$$\mathbf{I} = \begin{pmatrix} I_{xx} & I_{xy} & I_{xz} \\ I_{yx} & I_{yy} & I_{yz} \\ I_{zx} & I_{zy} & I_{zz} \end{pmatrix}, \tag{7.7}$$

that has diagonal elements known as the moments of inertia and off-diagonal elements known as the products of inertia. In other words, if the external torques, the inertia tensor of the aircraft and some initial conditions are known, then the solution to the differential equation given in equation 7.5 provides the angular rate and attitude of the aircraft with respect to the r-frame.

7.1.2 The practical form of the equations of motion

Although equations 7.3 and 7.5 provide the basis for describing the natural motion of an aircraft, they are usually applied in a simplified form. Specifically, each is linearized with respect to steady flight that is formally defined as having constant velocity and attitude and some constant forward speed, s (i.e. $\dot{\mathbf{v}}^{\mathrm{b}} = 0$, $\boldsymbol{\omega}_{\mathrm{rb}}^{\mathrm{b}} = 0$ and $\left|\mathbf{v}^{\mathrm{b}}\right| = s$). In steady motion, the x-axis of the b-frame will coincide with the x-axis of the r-frame and the y

and z-axes may differ only as a result of a possible constant rotation about the x-axis (i.e. pitch). Further, because of the symmetry of the aircraft, the equations of motion can be split into two uncoupled sets; one for longitudinal and one for lateral motion. These facts yield linear equations of the following form (Bryson (1994)):

$$\begin{pmatrix} \delta\dot{v}_y \\ \delta\dot{v}_z \\ \delta\dot{\omega}_x \\ \delta\dot{\theta} \end{pmatrix} = \mathbf{F}_{lon} \begin{pmatrix} \delta v_y \\ \delta v_z \\ \delta\omega_x \\ \delta\theta \end{pmatrix} + \mathbf{G}_{lon} \begin{pmatrix} \delta e \\ \delta t \end{pmatrix}, \quad \text{and} \tag{7.8}$$

$$\begin{pmatrix} \delta\dot{v}_x \\ \delta\dot{\omega}_z \\ \delta\dot{\omega}_y \\ \delta\dot{\varphi} \end{pmatrix} = \mathbf{F}_{lat} \begin{pmatrix} \delta v_x \\ \delta\omega_z \\ \delta\omega_y \\ \delta\varphi \end{pmatrix} + \mathbf{G}_{lat} \begin{pmatrix} \delta a \\ \delta r \end{pmatrix}, \tag{7.9}$$

which define the motion of the aircraft in the longitudinal (vertical) and lateral (horizontal) planes, respectively, with respect to the conditions of steady motion. The terms δe, δt, δa and δr represent the deflections caused by the elevator, throttle, aileron and rudder, respectively. In other words, they represent the interaction of the pilot (or auto-pilot) with the system. The variables $\delta\theta$ and $\delta\varphi$ represent deviations in the pitch and roll of the aircraft and the deviations in the yaw are given by $\delta\dot{\psi} \cong \delta\omega_z$.

Equation 7.8 defines the motion of the aircraft in the plane defined by the y and z-axes and equation 7.9 defines its motion in the plane defined by the y and x-axes. The state variables are given in the r-frame and the superscript, r, has been dropped for convenience. Figure 7.2 shows the elements of the two state vectors. The dynamic matrices, $\mathbf{F}_{lon}$ and $\mathbf{F}_{lat}$ are functions of a) the approximate magnitude of gravity, b) the moments and products of inertia, and c) the derivatives of the external forces and torques with respect to each of the state elements (that arise as a result of the linearization that is done). Parameter sets b) and c) are specific to a given aircraft and their derivation is a complex process based on a combination of observations made when the aircraft is stationary, in a wind tunnel and during well-observed flights. For example, consider Heffley and Jewell (1972) and Teper (1969) for examples of aircraft stability and control

data for several (mainly military) aircraft. The control matrices $\mathbf{G}_{\text{lon}}$ and $\mathbf{G}_{\text{lat}}$ are also aircraft and autopilot dependent. Finally, remembering the small angles that have been assumed, the velocities in the reference frame can be approximated by:

$$\delta\mathbf{v}^{\text{r}} \cong \begin{pmatrix} \delta v_x + s\delta\psi \\ \delta v_y \\ \delta v_z + s\delta\theta \end{pmatrix}. \tag{7.10}$$

The standard way of solving equations 7.8 and 7.9 is discussed in the next section and more details about their derivation and complete form can be found in Babister (1980) and Bryson (1994), for example.

7.1.3 Solving the equations of motion for natural motion

The natural motion of the aircraft is its response to perturbations when there is no interaction by means of the controls. This means that if equations 7.8 and 7.9 are represented in general form by $\dot{\mathbf{x}} = \mathbf{F}\mathbf{x} + \mathbf{G}\mathbf{u}$, then the natural motion is given by the solution to the homogeneous form of the same equation, $\dot{\mathbf{x}} = \mathbf{F}\mathbf{x}$. As discussed in Guterman and Nitecki (1991), if the coefficients of **F** are constant in such an equation, it is an ordinary linear differential equation that is well known to have solutions of the following form:

$$\mathbf{x} = e^{\lambda t}\mathbf{v}, \tag{7.11}$$

where λ is an eigenvalue of **F** and **v** is an eigenvector of **F** corresponding to λ. The complete natural motion will then be the sum of all such solutions. It is easy to see from equation 7.11 that the solutions corresponding to real eigenvalues will be exponentially increasing or converging functions of one of the forms shown by the solid lines in Figure 7.3. Whether they increase or converge will depend on the sign of the eigenvalue. These will be called Type A modes. It is also important to realize that although the longitudinal and lateral equations discussed in the last section will have complex eigenvalues, the corresponding solutions are always real. As discussed in Etkin (1982), complex eigenvalues will always appear in conjugate pairs that will result in a real-valued, exponentially weighted sinusoid of the form shown by the dotted lines in Figure 7.3

(depending on the sign of the real part of the eigenvalue). See Appendix D.3 for more about this. These will be called Type B modes.

Babister (1980) shows that the longitudinal motion of an aircraft will generally be composed of two Type B modes (i.e. it has two pairs of complex conjugate eigenvalues). The first is called the *Short-Period Oscillation* and the second is called the *Phugoid Mode* (or Long-Period Oscillation). For typical aircraft, the Short-Period Oscillation has a 4-10 s period and is very quickly damped (it will usually reach half-power within 1 s). The period and time to half-power are inversely proportional to the forward speed. The Phugoid mode, on the other hand, is a slow sinusoidal motion that is very slowly damped. It has a typical period of between 25 and 150 s (that increases with the flight speed) and can take thousands of seconds to reach half-power.

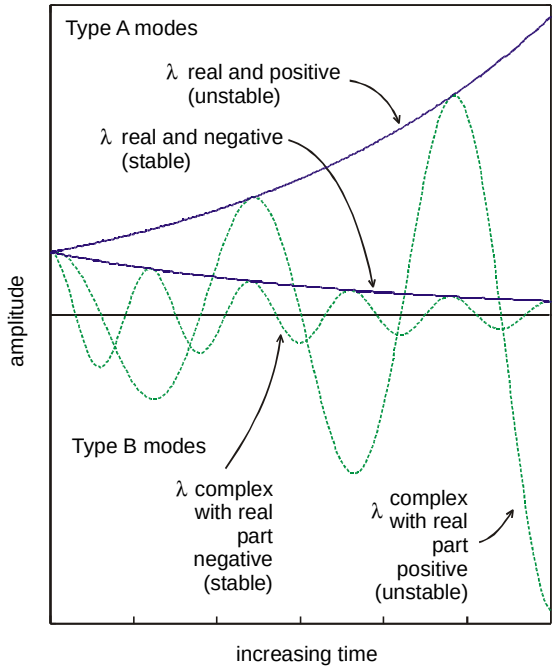

Figure 7.3 The forms that the natural motion can take

The lateral motion of an aircraft is made up of three basic modes. The first is known as the *Dutch Roll Mode* (or Yawing Oscillation). It is a Type B mode that has a period between 3 and 15 seconds and a half-time as short as 3 s. Generally speaking, the half-time and period of the Dutch Roll are inversely proportional to the forward speed and the latter is also proportional to the wing span, Babister (1980). The second mode, known as

the *Pure Rolling Motion* is a highly damped Type A mode with a typical half-time of less than 2.0 s. Finally, the last mode is known as the *Slow Spiral Motion* that is caused by a slowly damped Type A mode that can have a half-time of up to a few hundred seconds.

Recall from the last sub-section that the elements in the state vector **x**, correspond to the deviations of the aircraft motion from steady motion. The total natural motion will be made up of the sum of each solution of the form shown in equation 7.11. There will be one such solution for each eigenvalue of the system, **F**. For example, consider the first element of the longitudinal case (equation 7.8). Its solution implies that the deviation of the forward component of the velocity from *s* is given by the sum of two modes: the Short-Period and Phugoid modes. This point is further explored in Section 7.1.4.

7.1.4 What does natural motion typically look like?

This section briefly draws upon examples to provide the reader with a better understanding of the natural motion that was described above and to identify the modes with frequency content that might affect an airborne gravity system.

Because the dynamic stability derivatives and moments of inertia (i.e. matrices $\mathbf{F}_{\text{lon}}$ and $\mathbf{F}_{\text{lat}}$) for aircraft commonly used in airborne gravimetry were not available to the author at the time of printing, an example from Bryson (1994) is borrowed and built upon to demonstrate the characteristic motion of a small general aviation aircraft called the Navion. Consider Figure 7.4 that shows the longitudinal modes for this aircraft when it is traveling at a speed of about 50 m/s and note that different time scales have been used. Notice first, that the Short Period mode is damped very quickly, having virtually no influence after only 1.5 s. The Phugoid mode has a larger effect, however, with a period of approximately 25 s and taking some 75 s to reach half-amplitude. Consider Figure 7.5 that shows the lateral modes for the same aircraft. Although the Pure Roll and Dutch Roll modes are quickly damped, the Slow Spiral mode is not. Note that the period of the Dutch Roll mode is only about 3 s. These observations point to the conclusion that only the Phugoid and Slow Spiral modes have frequency content that overlaps with the gravity signal (i.e. their duration and period are long enough to influence the determination of gravity within the bandwidth of interest).

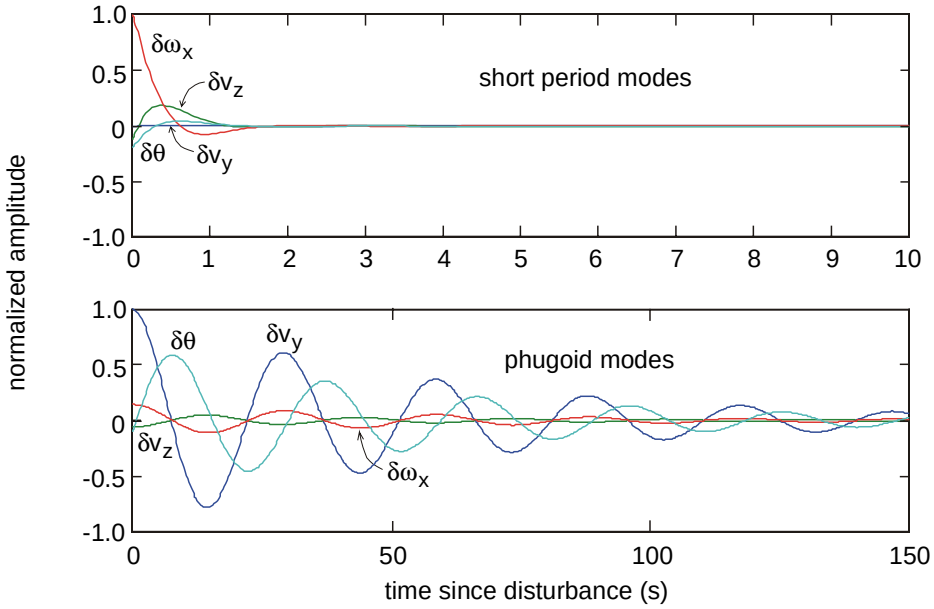

Figure 7.4 Longitudinal modes for the Navion

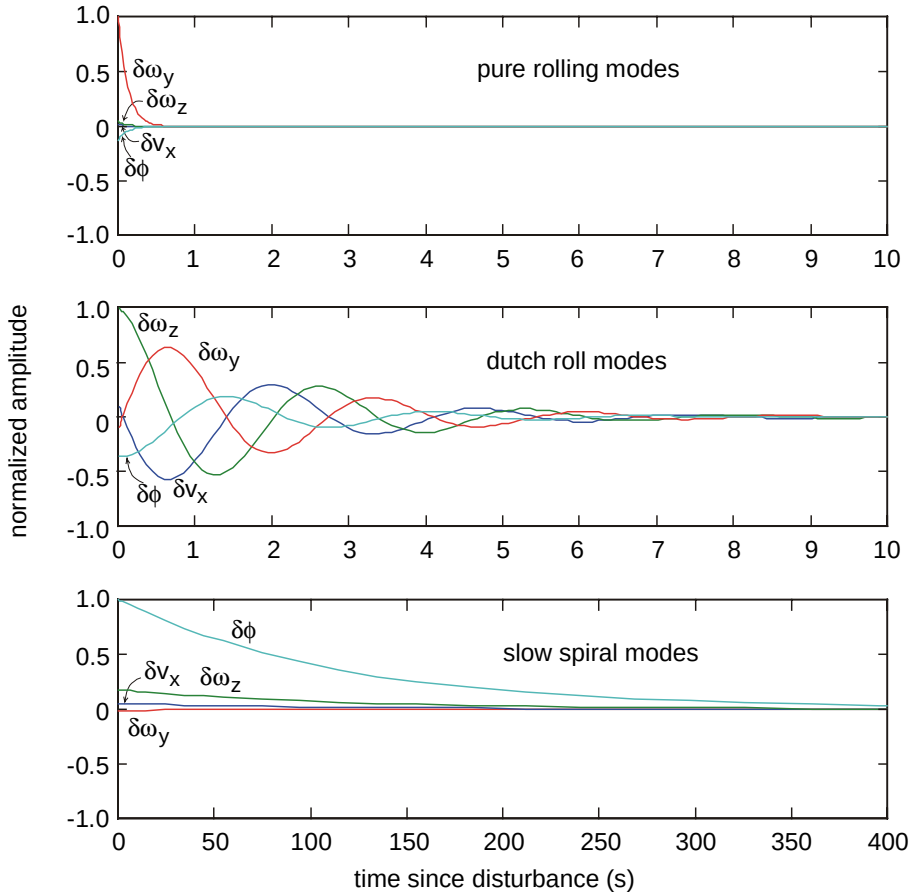

Figure 7.5 Lateral modes for the Navion

Because the airplanes used in airborne gravimetry are typically larger than the Navion, it is important to consider the modes of a larger aircraft. For this reason, the above example is extended to include the dynamic model for a Boeing 747. Although such an aircraft is undoubtedly much larger (and faster) than any practical surveying aircraft, it is included to represent an extreme case. For now, it will be assumed that the responses of typical survey aircraft will lie somewhere between those of the Navion and the 747. With this in mind, consider the longitudinal and lateral modes for the Boeing 747 flying at a speed of about 235 m/s that are presented in Figure 7.6 and Figure 7.7, respectively.

Although the Short Period and Pure Roll modes have longer damping times than for the Navion, their effects still last much less than 10 s, meaning that will average out and not influence the performance of an airborne gravity system. The Dutch Roll mode has a period of only about 7 seconds, implying that it will also average out and have no influence on an airborne gravity system. (Notice that it is divergent, however, meaning that some sort of feedback is needed to keep the aircraft close to steady flight.) Important again are the Slow Spiral and Phugoid modes. The former is very similar in nature to that observed for the Navion, but the latter is markedly different. In this case, the Phugoid motion has a period of about 100 s and displays almost no dampening.

Having identified the Phugoid and Slow Spiral modes as those most likely to affect an airborne gravity system, it is useful to consider the deviations in trajectory that they cause. This is done in Figure 7.8 where the scale of both plots has been exaggerated for clarity. As shown, the Slow Spiral mode causes the aircraft to deviate in horizontal position from its ideal path as it slowly approaches the intended heading. This can take hundreds of seconds, and as demonstrated for a large aircraft in Etkin (1982), the aircraft can deviate from the intended path by as much as 600 m. The magnitude of this deviation depends on the size and speed of the aircraft. The Slow Spiral mode also causes a smaller, but similar deviation in roll. The Phugoid mode results in a sinusoidal deviation of the height of the aircraft that is accompanied by similar deviations in pitch and (mainly forward) speed. Etkin (1982) shows that the natural deviations in height can be as large as 700 m for large aircraft.

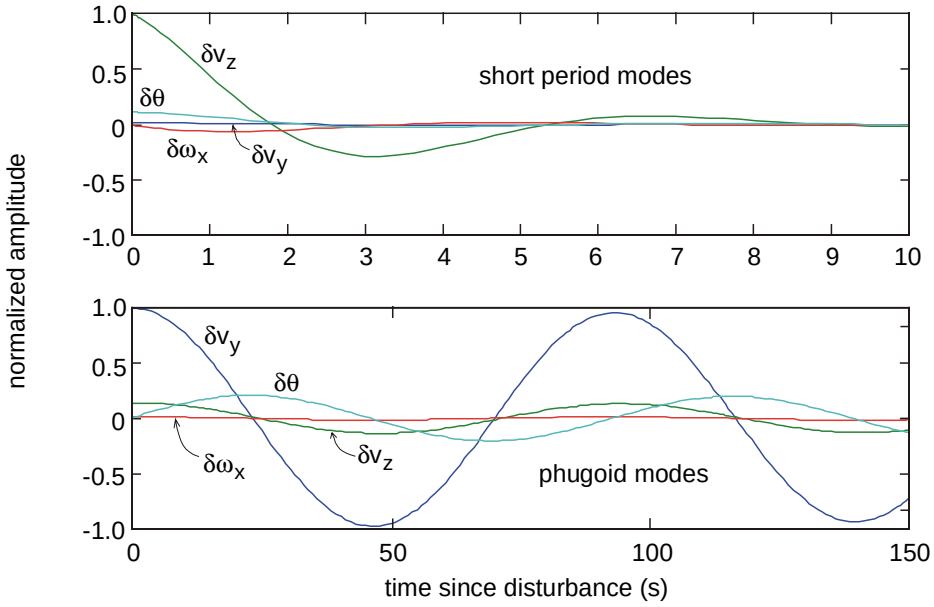

Figure 7.6 Longitudinal modes for a Boeing 747

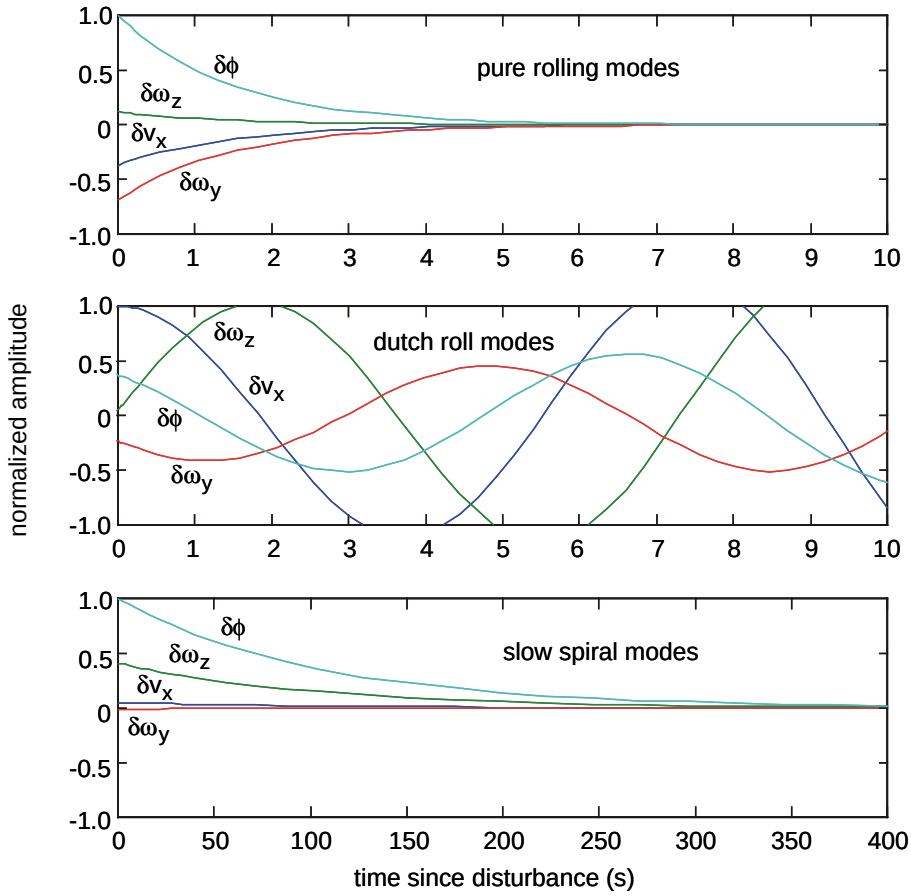

Figure 7.7 Lateral modes for a Boeing 747

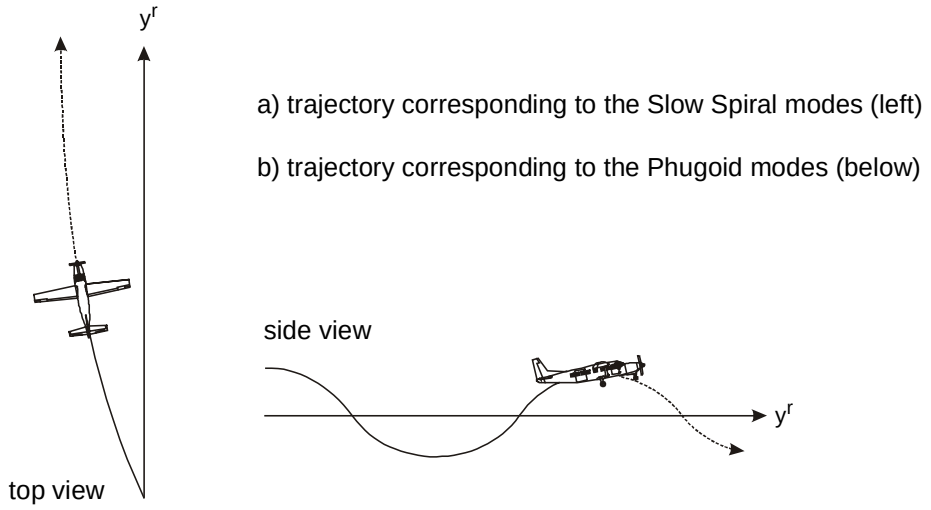

Figure 7.8 Trajectories corresponding to the Phugoid and Slow Spiral modes

The period of the Phugoid motion is approximated in Babister (1980) by:

$$T = \pi\sqrt{2}\,\frac{s}{g}, \tag{7.12}$$

where s is the forward speed of the aircraft and g is the magnitude of gravity. In other words, it is mainly dependent on aircraft speed and not on the aircraft itself. The amount of natural damping of the Phugoid modes is aircraft dependent, on the other hand, and depends specifically on the drag. It is typically small, however.

Details regarding the computations that were made to generate the data presented in Figure 7.4 through Figure 7.7 are given in Appendix D.2. Note that as Bryson (1994) suggests, the amplitude of the deviations has been normalized so that their relative sizes can be compared. Because the goal of this section is to identify problematic modes by frequency content, this is not considered to take away from the analyses. Further, it is the magnitude of the residual dynamics (i.e. the dynamics actually present after pilot or autopilot interaction) that are of importance in defining their effect on an airborne gravity system. Discussion of this is left until Section 7.2, where this is demonstrated using real data.

7.1.5 How might natural motion affect airborne gravimetry in practice?

In the last section, the Phugoid and Slow Spiral modes were identified as those most likely to affect airborne gravimetry. The identification of the Phugoid mode as a source

of error should have come as no surprise to readers who are familiar with the field of airborne gravimetry. In fact, errors associated with aircraft dynamics are often lumped together and simply referred to by airborne gravity researchers as those caused by Phugoid motion. At the same time, however, the same group of errors is also often referred to as being due to the autopilot used by the aircraft. These facts are somewhat contradictory and imply the need for clarifying how the natural motion causes errors in the estimated gravity field. This will be done below and in the next section.

In doing so, it is useful to recall the error equation for airborne gravimetry that was given in equation 2.6. For convenience, it is repeated below, leaving out the errors due to DGPS (that were treated in Part 2) and time synchronization (that will be addressed later, in Section 7.3):

$$d\delta\mathbf{g}^{l} = \mathbf{F}^{l}\boldsymbol{\varepsilon}^{l} - \mathbf{R}_{b}^{l}\, d\mathbf{f}^{b}. \tag{7.13}$$

The first term of this equation says that the misalignment errors, $\boldsymbol{\varepsilon}^{l}$, will present themselves as functions of the specific force, $\mathbf{F}^{l}$, resulting from the dynamics and the second term says that any dynamics induced errors in the measurements made by the accelerometers (e.g. scale factor errors) will present themselves directly in the solution after rotation from the b-frame to the l-frame. Both of these clearly depend on the specific force experienced by the aircraft. It is useful, therefore, to consider the characteristics of the specific force caused by each of the modes. This can be done by differentiation with respect to time of the deviations in velocity that were observed in Figures 7-4 to 7-7. The result is practically no signal for the Slow Spiral mode and another exponentially damped sinusoid for the Phugoid mode. As shown in Figure 7.4 and Figure 7.6, although the Phugoid mode affects both the upward (z) and forward (y) components, it will be much larger in the forward direction.

In the presence of accelerometer scale factors, this type of motion will result in errors in the gravity disturbance, as given by the second term in equation 7.13. Given that the variability of the accelerations is significantly larger in the forward direction, they are likely to be problematic for vector gravimetry because they will have their largest effect on the horizontal components of the gravity vector.

Recall from the full form the first term in equation 7.13 (see equation 2.11), that the forward specific force will act to amplify the effect of the misalignment error in all three components of the gravity disturbance estimate.

Although there will also be some errors due to the gyroscope scale factors, it is easy to show that they will be negligible in magnitude for the variations in angular rate caused by either the Phugoid or Slow Spiral modes. As a result, they will simply be lumped in as a part of the misalignment term, $\boldsymbol{\varepsilon}^l$, for the remainder of this chapter.

Although the Slow Spiral modes will not cause any dynamics related errors in the estimated gravity disturbance, the deviations from the intended trajectory may mean that the measurements of the gravity field are being taken at significantly different locations than intended (e.g. the anomalous gravity field can vary fairly significantly over a horizontal distance of 600 m). This problem will increase in severity for rough gravity fields and as the distance over which interpolation needs to be done increases. It would also play a greater role in the small area surveys typical of resource exploration. This issue was discussed in Section 6.1.

Again, because of the normalization that was done earlier, only the general form of the accelerations caused by the Phugoid motion has been discussed in this section. The magnitude of their effects will be quantified using real data in Section 7.2 that follows the brief discussion of motion control in the next section.

7.1.6 On the use of controls to reduce the natural motion

Anyone who has been in an airplane knows that its height does not typically deviate from level by anything like 600 m, as implied in Section 7.1.4. In fact, although they may not crash the aircraft, several of the deviations from steady flight caused by the natural modes would make a passenger very uncomfortable if not controlled in some reasonable way. The job of a pilot or autopilot is to ensure that none of these deviations present themselves in ways that cause discomfort to passengers (or gravimeters, for that matter).

The aircraft controls are used to counter the natural motion according to equations 7.8 and 7.9. In the case of a real pilot, he or she manipulates the elevator, throttle, aileron and rudder in order to accomplish this. In the case of an autopilot, measures of the observable

modes are used in a control system to automatically manipulate the same controls and achieve the desired effect. As described in Bryson (1994), autopilots can take many forms and the observables can include climb rate, altitude, yaw, roll and pitch that can be measured by sensors such as barometers, gyroscopes and full inertial systems. There, the following two types of autopilots are identified: a stability augmentation system (SAS) and a command (or hold) autopilot. The SAS only stabilizes one or more unstable modes via the feedback of the necessary parameters. The command autopilot is more sophisticated, involves many channels of feedback and can even fly the plane from shortly after takeoff to just before approach for landing. Further details about the process of feedback for automatic aircraft control are external to the scope of this research and the interested reader is referred to Bryson (1994) for details.

In airborne gravimetry, the goal of the pilot or autopilot interaction is to control the components of the motion that cause significant errors in the system. Ideally, such interaction would remove all accelerations along a flight line. Because this is not possible in practice, some acceleration must be tolerated. For a given application, the appropriate terms in equation 7.13 can be used to roughly estimate the magnitude of the maximum allowable accelerations in each component. In order to keep the effects of each term below 1 mGal for the scalar gravimetry case, for example, an accelerometer scale factor error of 25 ppm implies that the deviations in upward acceleration must remain below 0.40 m/s^2 (in the relevant bandwidth) and a horizontal misalignment of 30 arc seconds implies that the net horizontal acceleration must be less than 0.07 m/s^2 (also in the relevant bandwidth).

Using data from the onboard DGPS and very accurate SINS, the next section will evaluate how successfully the deviations from ideal motion were minimized in two recent airborne campaigns.

7.2 Examples of the aircraft motion observed during recent campaigns

This section uses the navigation data from two recent airborne gravity campaigns to demonstrate the characteristics of the dynamics that the aircraft underwent. Different approaches to approximating the ideal motion were employed in the two tests. The first

used a basic autopilot to control various aspects of the motion and the second was controlled entirely by the pilot. Each is discussed in turn below.

7.2.1 The Kananaskis campaign

Details about the survey specifications for this campaign can be found in Appendix A.1. Recall from there that the aircraft used for the survey was a Cessna Conquest turbo prop, that the average flying speed was 100 m/s and that the goal of the survey was to collect data for geoid determination. An autopilot was used during the campaign. While little is known about its design, it is known that it was a basic SAS autopilot that controlled the heading and altitude. Although it is not certain what parameters were used in the feedback systems, it is likely that the heading and altitude were controlled by separate systems. For the first, the heading was probably maintained by adjusting the roll, based on feedback of the heading to the compass card heading channel. For the second, it is likely that the altitude would have been maintained using the pitch trim, based on feedback of the altitude. It is known that the pilot used the throttle to maintain the speed of the aircraft. This information is based on personal communication with K. Tennant at Intermap Technologies Corporation.

Consider Figure 7.9 that shows the trajectory of a sample flight line from that campaign. Figure 7.9a shows the deviation in east position from the intended north-south trajectory (note that the scale in the east-west direction has been exaggerated for clarity) and Figure 7.9b shows the deviations in height from the average ellipsoidal height of 4357 m. While the autopilot system being used limits the deviations in height to well within ± 10 m, there is a clear deviation in horizontal position that is made up by a low-frequency component that causes the trajectory to vary by over 800 m and by a higher frequency periodic component that has an amplitude of roughly 50 m.

The low-frequency component looks suspiciously like the Slow Spiral mode that was identified in Section 7.1.3. Given that the heading was controlled by the autopilot, however, this is an unreasonable explanation. The observed deviations could also be explained if the heading was controlled by a magnetic compass, an idea that is supported by Figure 7.10. A straight line was fit to the observed trajectory for each flight line that was flown on September 10 to generate this figure. Clearly, such a pattern would arise if

the aircraft was flying according to a compass and if the magnetic declination had only been input once for the whole survey (near the central longitude in this case). While such low-frequency deviations in position have no effect on acceleration, they may cause errors due to interpolation to the intended grid, especially in areas with rough gravity fields or when closely spaced profiles are to be flown.

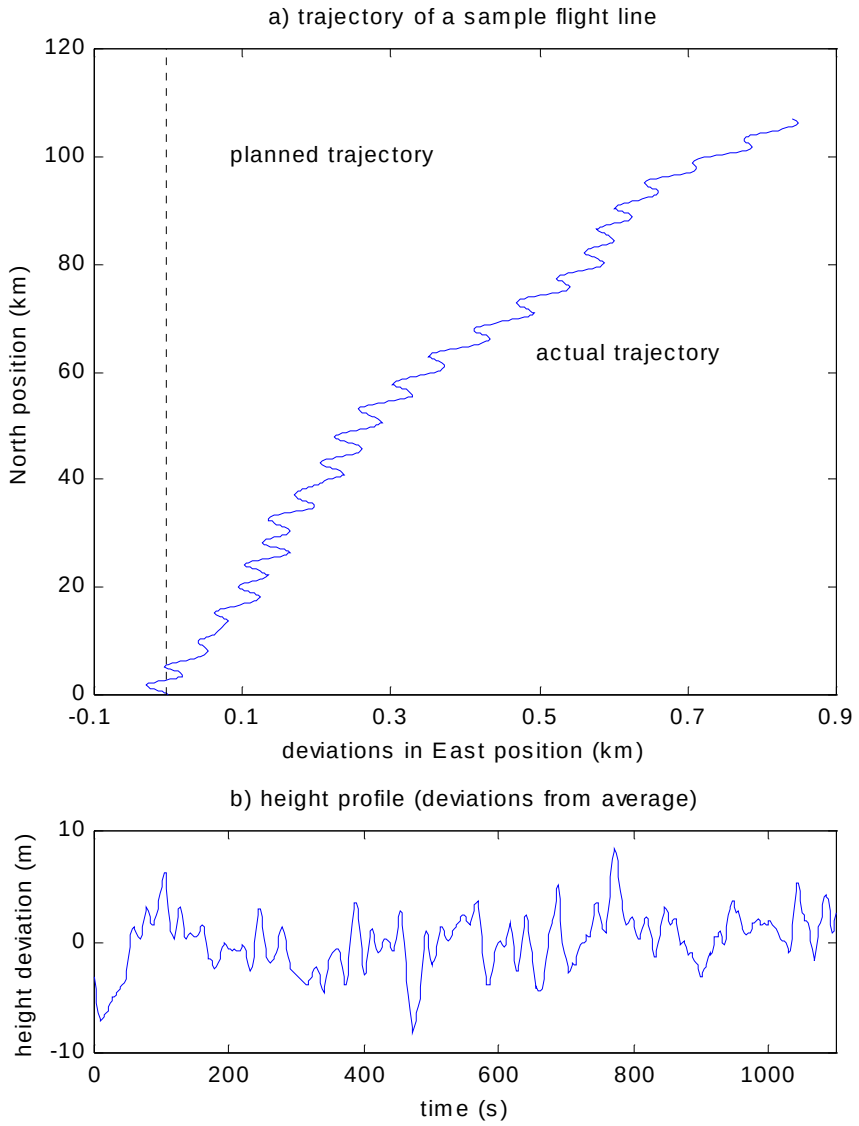

Figure 7.9 Trajectory of a sample flight line of the Kananaskis campaign

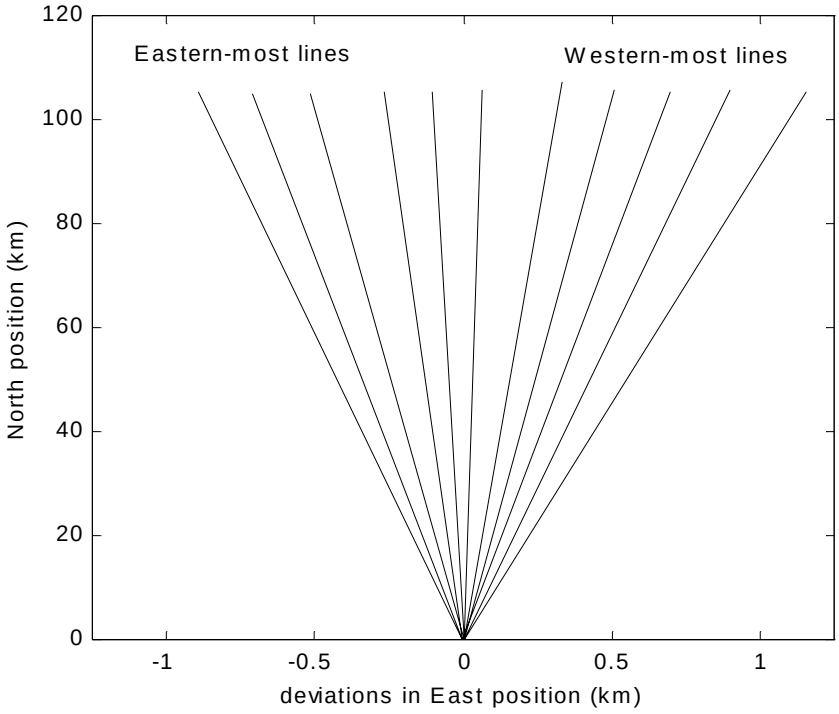

Figure 7.10 Observed horizontal deviation for all lines on September 10

More troublesome for airborne gravimetry are the sinusoidal deviations of the aircraft trajectory from these low-frequency curves. Consider Figure 7.11 that demonstrates the l-frame components of the specific force experienced by the aircraft for the flight line that was displayed in Figure 7.9. These were derived from the available DGPS position data by differentiating it and accounting for Coriolis acceleration and normal gravity. In other words, at the scale used here, the values shown differ negligibly from the true specific force experienced by the aircraft. The data in Figure 7.11 has been bandlimited to 0.05 Hz to isolate it to the bandwidth that is relevant for airborne gravimetry. For completeness, the total (i.e. unfiltered) specific force is shown in Appendix D.3. The deviations seen in Figure 7.9a clearly result mainly in sinusoidal deviations in specific force in the east component that have a maximum amplitude of 0.51 m/s^2. Observation of Figure 7.11 shows that the period of the deviations in the east direction varies between 40 and 50 s.

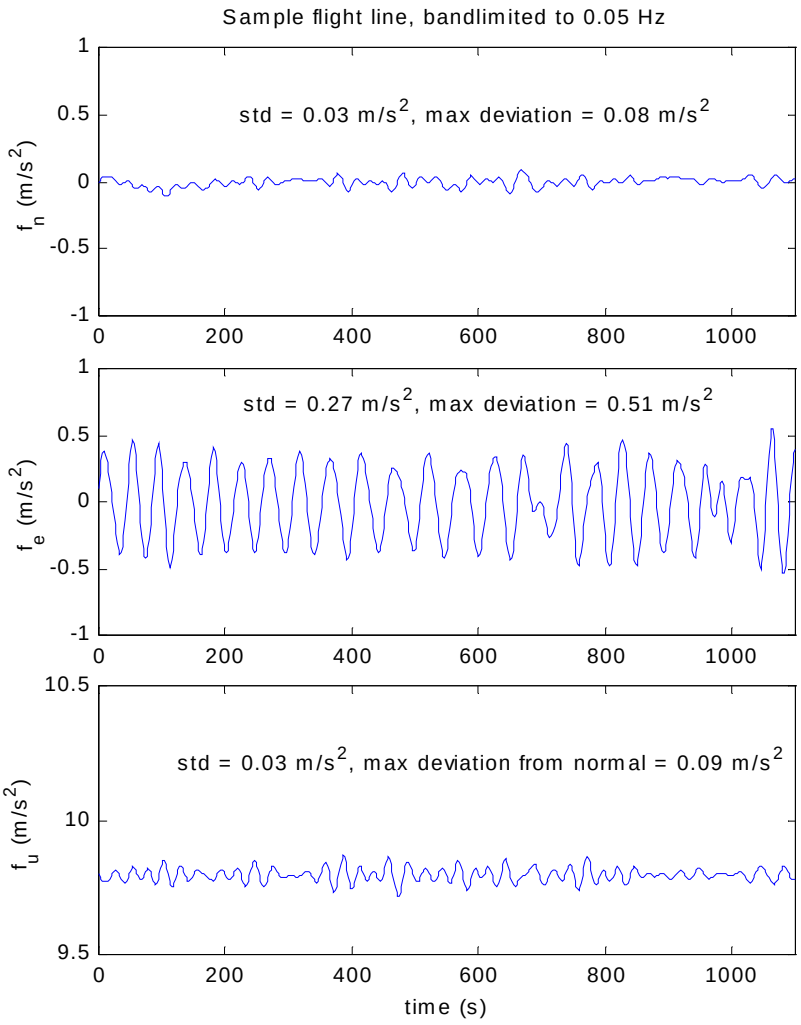

Figure 7.11 Acceleration experienced during the Kananaskis campaign (0.05 Hz)

There are at least two possible explanations for this lateral motion. The first and least likely is that the aircraft is undergoing a Dutch Roll motion that has an uncharacteristically long period (recall from Section 7.1.2, that the period of the Dutch Roll mode is typically no longer than 15 s). This seems especially unlikely in light of the fact that the period of the Dutch Roll for the much larger Boeing 747 was shown to be only 7 s in Section 7.1.3. A second possible explanation is that it comes indirectly from the Phugoid motion of the aircraft. This argument is supported by equation 7.12 that predicts the period of the motion to be 45 s for the speed at which the aircraft was flying; a value very close to the observed value (of between 40 and 50 s) in this case. It is

possible that the Phugoid mode (that was shown in Section 7.1 to be longitudinal) is causing the observed lateral motion because of some interaction with the autopilot. For example, in trying to counter the periodic change in altitude and pitch caused by the Phugoid mode, the autopilot may be causing a change in roll and therefore horizontal position and acceleration that all share the same period. For completeness, the speed and attitude of the aircraft are also shown in Appendix D.3 for the same flight line.

Whatever the explanation for this lateral motion, it can have severe effects on the estimation of gravity. Consider again the scalar gravimetry example. Recall the example from Section 7.1.6 that showed that the upward acceleration must be kept below 0.40 m/s^2 and the horizontal accelerations must be kept below 0.07 m/s^2 in order to keep their respective effects less than 1 mGal (given certain assumptions about the magnitude of the scale factor and misalignment errors). While this does not appear to be a problem for the upward acceleration, which only reaches 0.09 m/s^2, it is clearly a problem for the horizontal acceleration, which reaches 0.51 m/s^2.

The common way to deal with this problem is by low-pass filtering the gravity estimate. This works well because the magnitude of the specific force is significantly smaller as the bandwidth is reduced. Consider Figure 7.12, for example, that shows the specific force of the aircraft within the bandwidth corresponding to frequencies below 0.01 Hz. They are clearly much smaller in this bandwidth in all components, explaining why it is so much easier to estimate gravity for geodetic applications (frequencies lower than 0.01 Hz) than for geophysical applications (which try to include frequencies between 0.01 Hz and 0.05 Hz). Consider also Figure 7.13 that shows estimates of the gravity field after filtering to 90 and 30 s, and the obvious advantage of the heavier filtering.

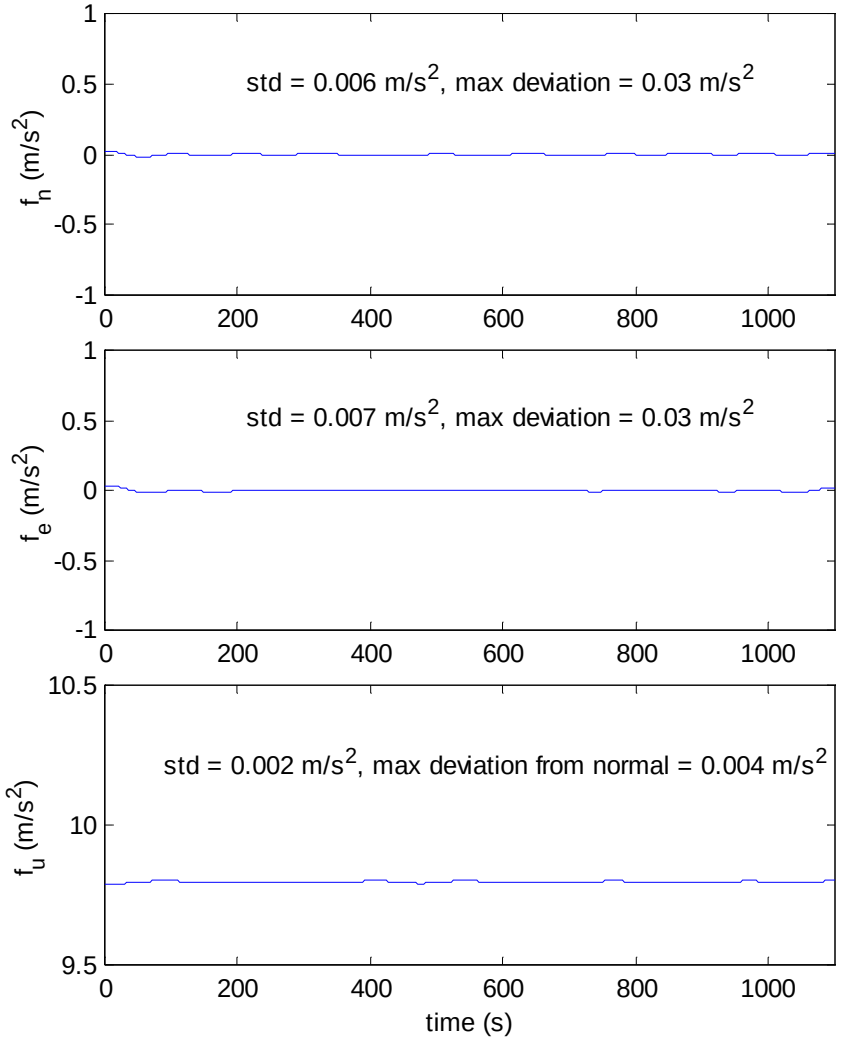

Figure 7.12 Acceleration experienced during the Kananaskis campaign (0.01 Hz)

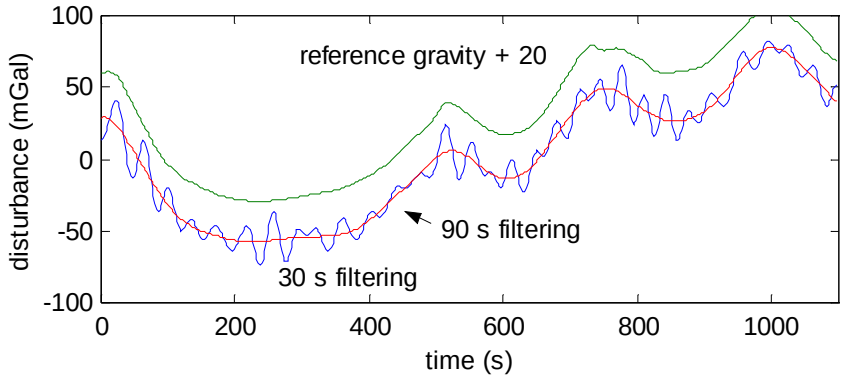

Figure 7.13 Estimates of the gravity field in different bandwidths

7.2.2 The Alexandria campaign

Details about this campaign can be found in Appendix A.3. It differed from the Kananaskis campaign in several important ways. The first is that the flying speed was half as much (45 m/s) because one of the goals was to test the performance of the system for higher resolution applications, i.e. for frequencies above 0.01 Hz. The second is that it was flown in a Cessna Caravan that was not using an autopilot. Instead, the pilot controlled the flight. The goal of this sub-section is to highlight the effects that those differences have on the dynamics of the aircraft.

Consider Figure 7.14 that shows the deviations in trajectory from the ideal straight-line case. By comparing this to Figure 7.9, one can see that the horizontal deviations from ideal are smaller for the Alexandria campaign than for the Kananaskis campaign. Figure 7.15 shows that the amplitude of the specific force occurring in the bandwidth around the Phugoid period is smaller than in the Kananaskis campaign. The standard deviation of the horizontal accelerations is only 0.06 m/s^2 in this case, compared to 0.27 m/s^2 in the Kananaskis case. Also, there are no low-frequency deviations, most likely due to the fact that the pilot was flying according to a heading derived from GPS.

For completeness, Appendix D.3 contains plots of the observed attitude, speed and specific force (without filtering). It can be noted from the appropriate figure in Appendix D.3 that although the deviations in height are similar in this campaign to what they were in the Kananaskis campaign, the unfiltered upward acceleration is considerably larger for the Alexandria campaign than for the Kananaskis campaign. Fortunately, as shown in Figure 7.15, it is considerably smaller in the relevant bandwidth.

The appropriate figure in Appendix D.3 shows that the specific force of the aircraft after low-pass filtering to 0.01 Hz is roughly the same as for the Kananaskis campaign.

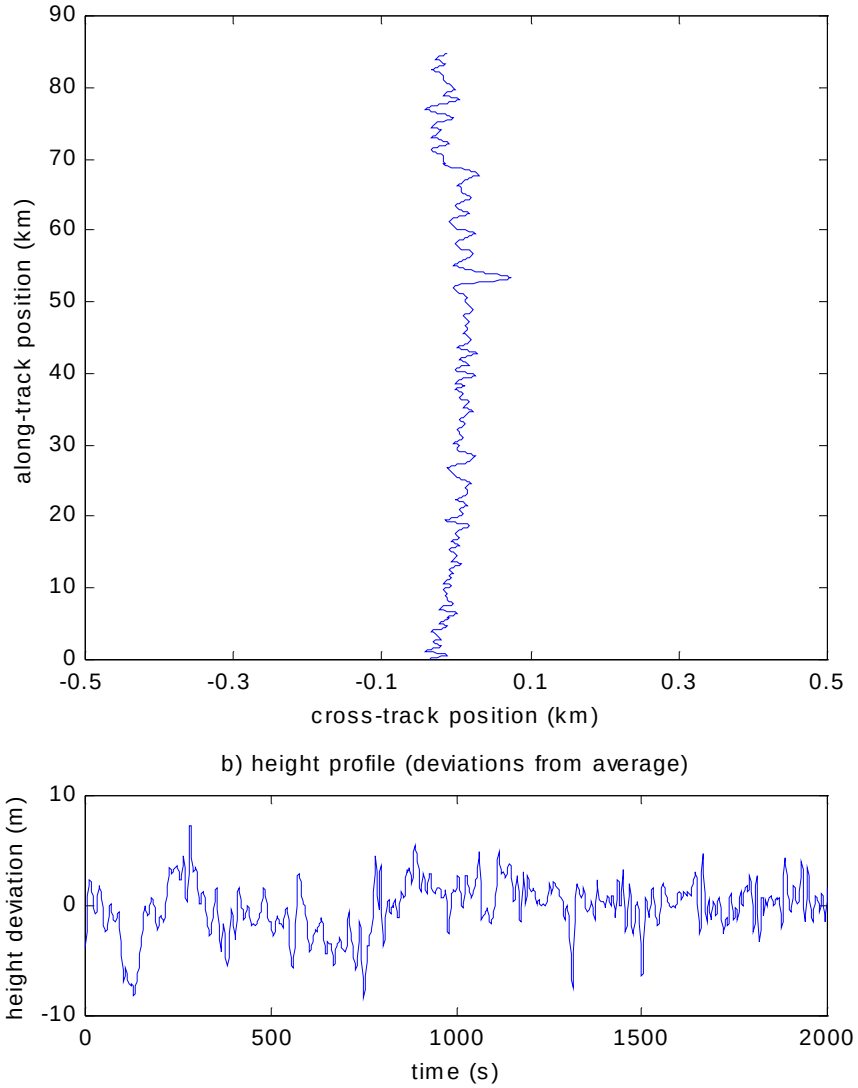

Figure 7.14 Trajectory of a sample flight line of the Alexandria campaign

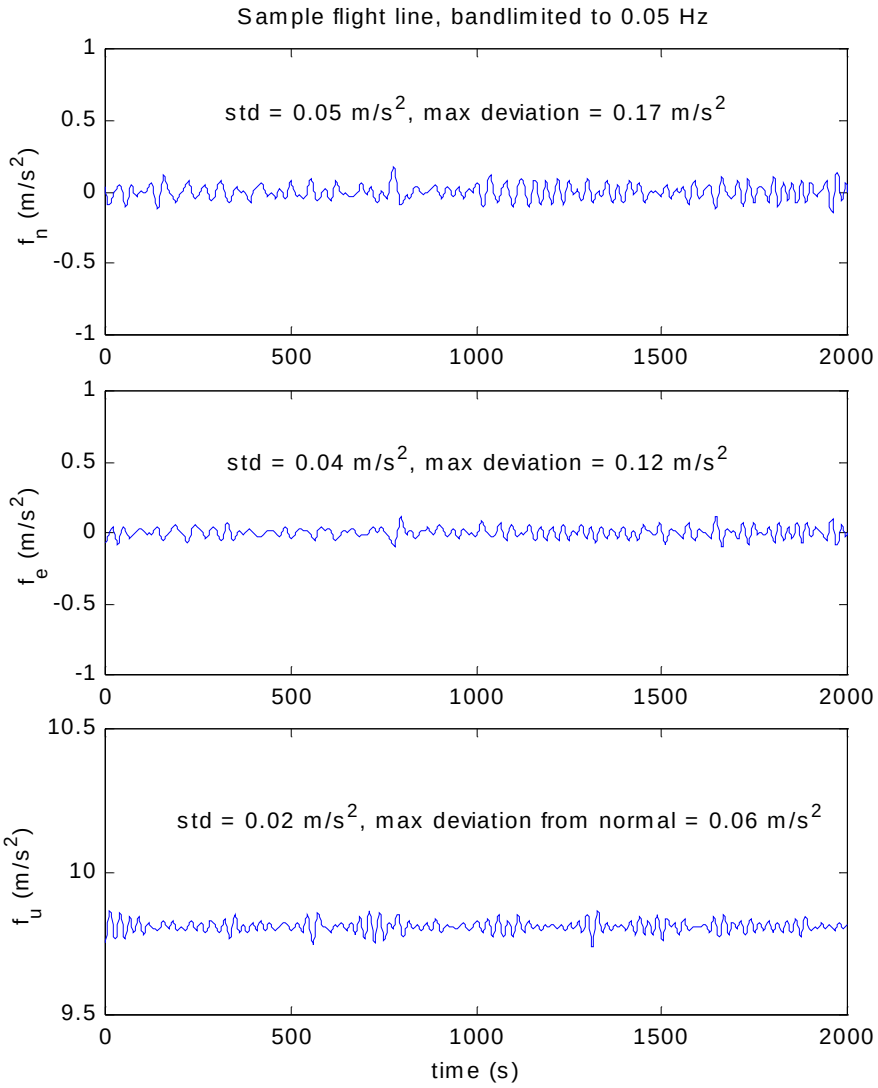

Figure 7.15 Acceleration experienced during the Alexandria campaign (0.05 Hz)

7.3 Estimating spectra of the SINS induced gravity errors

As touched upon in the last two sections, the dynamics experienced by the survey aircraft induce SINS sensor errors and highlight modeling errors that result in errors in the estimated gravity field. Although it is extremely difficult to define these as a function of time (if we could, we would simply subtract them from our estimates), it is possible to characterize them in the frequency domain if the dynamics are sufficiently well understood. This section does this by building on the knowledge of the aircraft dynamics

that was gained in the last section and using it to isolate and quantify the effect of each of the SINS related errors.

Recall from equation 2.7 that the errors in the upward component of the estimated gravity disturbance are given by the right hand side of the following:

$$d\delta g_u - d\dot{v}_u = -df_u + (f_e \varepsilon_n - f_n \varepsilon_e) + (\dot{\mathbf{A}}\mathbf{f}^b + \mathbf{A}\dot{\mathbf{f}}^b)\, dT, \qquad (7.14)$$

where the first term is due to SINS sensor errors, the second term is due to the horizontal misalignment errors and the last term is due to time synchronization errors. Figure 7.16 and Figure 7.17 show estimates of the individual error spectra for each of these for the Kananaskis and Alexandria campaigns respectively.

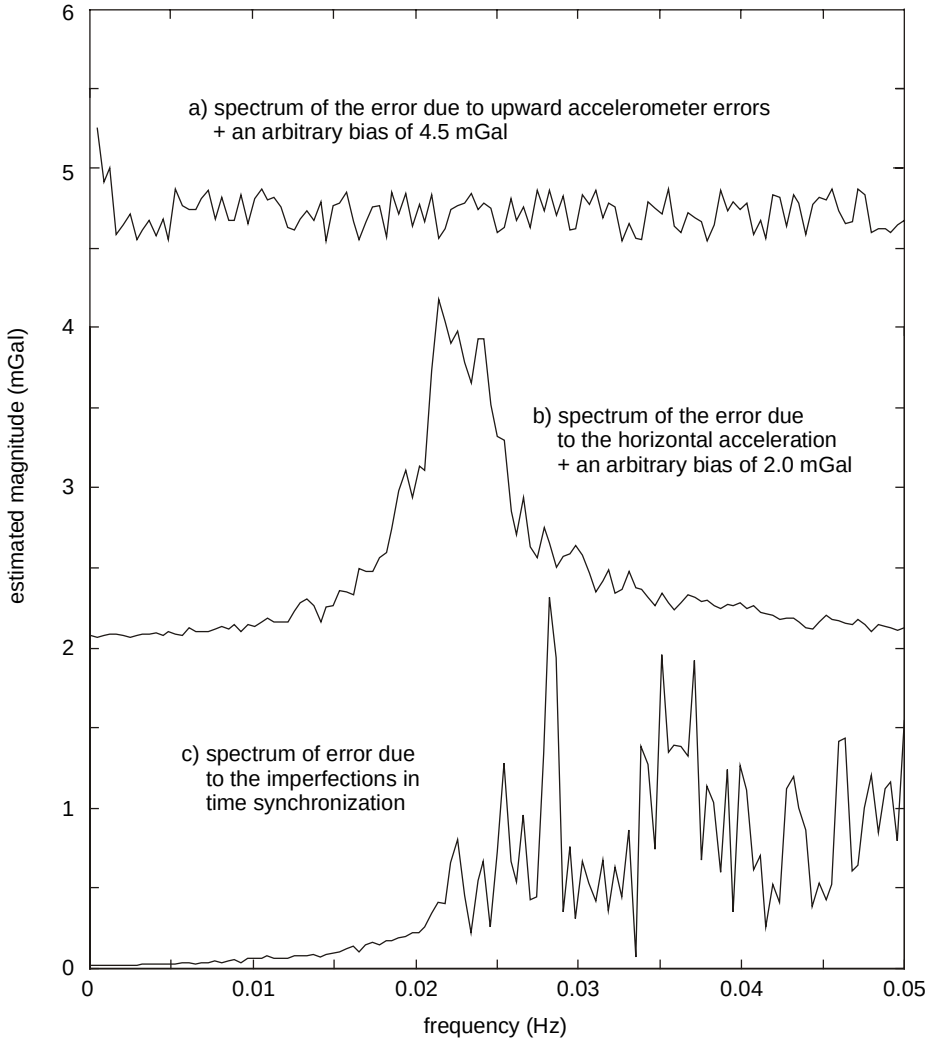

Figure 7.16 Spectra of the SINS induced gravity errors for the Kananaskis campaign

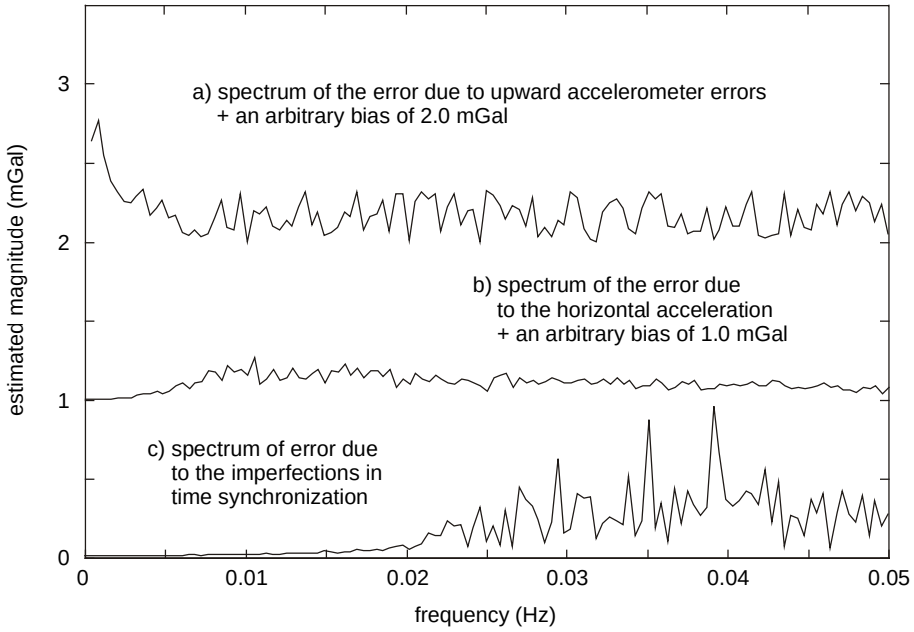

Figure 7.17 Spectra of the SINS induced gravity errors for the Alexandria campaign

The spectra of the SINS sensor errors, df_u, (i.e. Figure 7.16a and Figure 7.17a) were calculated from raw data collected using the upward pointing accelerometer during the static period at the outset of each of the surveys. The aircraft engines were on in both cases, meaning that the effect of the associated vibrations can be roughly evaluated by comparing these spectra to the solid line in Figure 7.1. For now, it will be assumed that these error spectra also represent the SINS errors while in flight. Strictly speaking, this is not true because scale factor errors may play a role (especially in the turns between flight lines).

The spectra of the errors owing to horizontal misalignments, $(f_e\varepsilon_n - f_n\varepsilon_e)$, were estimated by multiplying the average spectra of the horizontal specific force experienced during all of the flight lines by a constant approximate horizontal misalignment ε = 30 arc seconds. This was done for each campaign to generate the spectra shown in Figure 7.16b and Figure 7.17b. The choice of 30 arc seconds as an approximate value for the misalignment was made based on the experiences with a navigation grade SINS that are reported in Skaloud (1999). Despite this, the choice is somewhat arbitrary and will be evaluated in Section 7.4. The curve in Figure 7.16b is an effective demonstration of how

the interaction of misalignment errors with dynamics can cause a sharp rise in the error spectrum at a frequency corresponding to the period of the Phugoid motion of the aircraft (approximately 0.022 Hz = 1/45 s, in that case). Equation 7.12 is therefore a useful expression because it helps in identifying the bandwidth in which a gravity survey will be useful. For example, unless the Phugoid mode can be suppressed, the maximum useful resolution of the data coming from a given airborne gravity survey is given by this period.

Because the bracketed portion of the third term in equation 7.14 is a time dependent function of the vehicle dynamics, it can be computed for each epoch of the survey. This is done according to equations 2.8 and 2.9 using the values of specific force and attitude that were observed by the SINS/DGPS system (recall Section 7.2 and Appendix D.3 for examples). The vector of time synchronization errors, dT, is assumed to be due to registration errors that occur when the data acquisition and time tagging process is blocked by other processes with higher priority. As a result, they are assumed to vary randomly at the level of 0.5 ms and have a white-noise-like spectrum. Note that this is a rather pessimistic assumption about the nature of the synchronization errors and therefore represents a worst-case scenario. See Schwarz and Li (1996b) for more details about this type of error. Finally, the spectra in Figure 7.16c and Figure 7.17c are computed from the combination of the spectra of these two sequences.

The most noticeable differences between Figure 7.16 and Figure 7.17 are the facts that the errors due to the misalignment are not concentrated around some central frequency for the Alexandria campaign and that they are of a much smaller magnitude than for the Kananaskis campaign. The former observation supports the argument that it was the interaction of the autopilot that caused the Phugoid modes to show up as lateral motion in the Kananaskis test and the latter observation is probably due in part to the slower flying speed. Also evident is the fact that the synchronization errors are smaller for the lower dynamics of the Alexandria campaign. Overall, the result is that the combined SINS errors are significantly smaller for the Alexandria campaign, especially above about 0.02 Hz.

This demonstrates the importance of effectively controlling the motion of the aircraft. It appears that the pilot in the Alexandria campaign did a better job of controlling the natural motion of the aircraft than the autopilot in the Kananaskis campaign, especially in the horizontal. It is not possible to draw any broad conclusions regarding the appropriateness of automatic control at this point, however, because the autopilot used in the latter was not very sophisticated. For example, a different autopilot might do a better job of controlling the motion. This should be evaluated when more data becomes available.

Note that the individual components of Figure 7.16 were used in Chapter 2 to create Figure 2-2 and that the shaded error spectra that were shown in Figure 5-5 and Figure 5-10 are the sum of each of the spectra in parts a, b and c of Figure 7.16 and Figure 7.17, respectively. The former was also used in Figure 7.1.

7.4 Reducing the dynamics induced SINS errors

It was seen in Section 7.2 that despite efforts by pilots and autopilots to control the accelerations experienced by survey aircraft, they remain at potentially significant levels. Then, in Section 7.3, the effect of the remaining motion on the determination of gravity was characterized as a function of frequency. Ways of reducing the effects of these errors so that the bandwidth of airborne gravity systems can be widened are the subject of the current section.

There are three ways of reducing the dynamics induced SINS errors. The first is to reduce the level of motion occurring within the relevant bandwidth, essentially eliminating the cause of the errors. Although this is not straightforward, it might be of interest in future studies to consider the natural motion of various aircraft and to study the potential of more sophisticated types of autopilots. In theory, both of these can be done in pre-mission analyses (along the lines of the examples that were given in Section 7.1.4.), if sufficiently accurate models of the aircraft and the feedback systems are available. Currently, these have not been found for typical survey aircraft. Despite this, a thorough investigation may be appropriate and could result in a better understanding of the necessary qualities of both the aircraft and the autopilot. Extension of such analyses to include helicopters is a very realistic possibility, again assuming the required data can be found. Alternatively,

navigation data could be collected during airborne testing of various aircraft under various conditions and studied in much the same way as in the examples of Sections 7.2 and 7.3. In turn, either of these approaches may aid in the choice of an ideal airborne platform and pilot/autopilot configuration.

The second way to reduce the SINS errors within the bandwidth of interest is to improve the quality of the sensors and systems for those components that are sensitive to aircraft dynamics, essentially reducing the magnitude of the effect that the dynamics have. This is generally an expensive approach. For example, by using better quality gyroscopes and by upgrading the system responsible for the synchronization of the SINS and DGPS data streams, the corresponding misalignment and synchronization errors in Figure 7.16 and Figure 7.17 could be significantly reduced. Improvement of the quality of the accelerometers could also result in better performance in terms of biases, noise limits and scale factors.

If the dynamics cannot be reduced and in the absence of better sensor systems, the third way to widen the bandwidth of the system is by improving the modeling that is done. Although this is the least attractive alternative because it tries to treat the symptoms of the underlying problem rather than its causes, it offers considerable potential, especially for data that has already been collected. In concept, the idea is to make use of the relationship between the dynamics and the error characteristics of the system in such a way that the errors can be estimated as a function of time and used to improve its performance. Of the errors discussed in Section 7.3, only the misalignment and scale factors can be modeled; although the noise and synchronization errors can be roughly characterized in the frequency domain, they cannot be modeled as a function of time. A brief example of how the residual misalignment errors can be modeled is the subject of the next section.

7.4.1 An example of modeling: estimating the misalignment of the SINS

Recall from Chapter 2 that the role of the gyroscopes in a SINS-based gravimeter is to provide the orientation of the accelerometers in space, $\mathbf{R}_{b}^{l}$, so that the measured specific force can be rotated from the b-frame to the l-level frame. By now, the reader should be

familiar with the notion that the errors in this orientation matrix cause errors in the estimated gravity field that are a direct function of the dynamics (recall Figure 7.16b that shows this for the specific force experienced during the Kananaskis campaign, for example). Although this clearly defines the dependence of gravity on accurate estimates of the attitude of the aircraft, it is important to discuss the other side of the same coin; accurate estimates of attitude also depend on a good knowledge of the gravity field. Estimation of the attitude of the aircraft using a SINS is a complex process that is implemented in KINGSPAD and described in detail in Schwarz (1998) and Skaloud (1999). In short, it requires the solution of the following system of equations that relates both sets of measurements ($\mathbf{f}^b$ and $\boldsymbol{\omega}_{ib}^b$) to the variables of interest (the position and velocity of the aircraft, $\mathbf{r}^e$ and $\mathbf{v}^e$, and especially the attitude matrix, $\mathbf{R}_b^l$):

$$\begin{pmatrix} \dot{\mathbf{r}}^e \\ \dot{\mathbf{v}}^e \\ \dot{\mathbf{R}}_b^e \end{pmatrix} = \begin{pmatrix} \mathbf{v}^e \\ \mathbf{R}_b^e \mathbf{f}^b - 2\Omega_{ie}^e \mathbf{v}^e + \mathbf{g}^e \\ \mathbf{R}_b^e (\Omega_{ib}^b - \Omega_{ie}^b) \end{pmatrix}, \tag{7.15}$$

where the matrix $,\Omega_{ie}^b$, is the rotation rate of the Earth (expressed in the b-frame), the gravity vector, $\mathbf{g}^e$, is usually approximated by the normal gravity vector and as usual, the dot above the variables denotes their derivative with respect to time. $\mathbf{R}_b^e$ is the transformation matrix that relates the e-frame to the b-frame (the derivation of which is sufficient information to obtain $\mathbf{R}_b^l$ because the relationship between the e-frame and the l-frame is a simple function of the position that can be reliably estimated using DGPS). This system of equations is typically linearized, augmented by states that represent the biases of the SINS sensors and solved in a recursive manner using a Kalman filter that uses the position and velocity obtained from DGPS as regular updates.

The second set of the above equations says that the attitude that is derived depends on the gravity field itself. This implies a somewhat circular problem in any effort to solve simultaneously for attitude and gravity; a model of the gravity field is needed in order to solve for the attitude of the vehicle, but the attitude of the vehicle is needed in order to solve for the gravity field.

From the perspective of SINSs optimized for navigation and attitude determination, this problem has long been recognized. In this case, a common solution is to derive the gravity vector, $\mathbf{g}^e$, from a global model of the gravity field rather than using the normal gravity model. As discussed in detail in Chapters 1 and 6, the drawbacks of such an approach include the band-limited nature of current and future global models and the relatively large errors in the current models. This means that considerable portions of the gravity field are not well represented and can result in a poorer estimate of the attitude.

From the perspective of SINSs optimized for gravity determination, this problem was discussed in Schwarz and Li (1996b) where an iterative method for solving it was first proposed. In the first step of this method, an initial estimate of the gravity field is obtained using the process given by equation 7.15 and by low-pass filtering the difference between the SINS and DGPS streams (as it has been throughout this dissertation so far, according to the process outlined in Section 2.3.1). This estimate is then used in a second step to correct the specific force data before using it again in the same process. It was proposed that by repeating this procedure, the estimate of the gravity field would converge to a final and accurate solution.

As it is presented, however, this approach is difficult for several reasons. The first is that the corrected specific force data (that the second step yields) still contains the effects of the accelerometer biases; it was assumed in Schwarz and Li (1996b) that the Kalman filter would be able to estimate them. Detailed analyses presented in Glennie (1999) showed, however, that it cannot. The second reason is that even in the absence of biases and noise in the specific force data, the Kalman filter does not have a very accurate means of specifying the presence of different error sources by bandwidth. For example, as shown in Figure 7.16b, the majority of the misalignment error only influences a very small portion of the spectrum of the estimated gravity field (between 0.010 and 0.030 Hz). It is difficult to take advantage of this knowledge in a Kalman filter approach.

The remainder of this section provides an example of how modeling can be done to estimate and remove selected errors. It implements an alternative method of estimating the misalignment terms. It is similar in spirit to both approaches outlined above. Like the

first, it makes use of a relatively accurate representation of the gravity field in order to improve estimates of the attitude, and like the second, it uses a band-limited estimate of the gravity disturbance coming from the system as a starting point. It is best introduced by again using equation 2.6 and approximating it as follows for the bandwidth between 0.010 and 0.030 Hz:

$$\mathrm{d}\delta\mathbf{g}^{l}_{\text{bandlimited}} = \mathbf{F}^{l}\boldsymbol{\varepsilon}^{l} + \mathbf{e}\,, \tag{7.16}$$

where $\mathrm{d}\delta\,\mathbf{g}^{l}_{\text{bandlimited}}$ is the band-limited error in the gravity disturbance, and **e** is a white measurement noise sequence that represents the sum of the SINS errors, the DGPS errors and the synchronization errors. For the SINS errors, this approximation is justified rather heuristically for now by considering Figure 7.16 and keeping in mind that the synchronization errors in Figure 7.16c represent a worst-case scenario. For the DGPS errors, recall Figure 5-5 that implies that the DGPS errors can be very nearly characterized as white noise within this bandwidth if the acceleration is derived using a single frequency approach.

It is hypothesized at this point that if the true specific force, $\mathbf{F}^{l}$, and the errors in the estimated gravity disturbance within the same bandwidth, $\mathrm{d}\delta\,\mathbf{g}^{l}_{\text{bandlimited}}$, can be sufficiently well estimated, then equation 7.16 can be used to solve for the misalignment term, $\boldsymbol{\varepsilon}^{l}$.

This hypothesis is evaluated using a Kalman filter with the following state vector, **x**, design matrix, **H**, and measurement vector, **z**:

$$\mathbf{x}_k = \begin{pmatrix} \varepsilon_{\mathrm{e}} \\ \varepsilon_{\mathrm{n}} \\ \varepsilon_{\mathrm{u}} \end{pmatrix}, \tag{7.17}$$

$$\mathbf{H}_k = \begin{pmatrix} 0 & -\mathrm{f}_{\mathrm{u}} & \mathrm{f}_{\mathrm{n}} \\ \mathrm{f}_{\mathrm{u}} & 0 & -\mathrm{f}_{\mathrm{e}} \\ -\mathrm{f}_{\mathrm{n}} & \mathrm{f}_{\mathrm{e}} & 0 \end{pmatrix}, \text{ and} \tag{7.18}$$

$$\mathbf{z}_k = \delta\,\mathbf{g}^{l}_{\text{bandlimited}} = \delta\,\mathbf{g}^{l}\Big|_{30\mathrm{s}} - \delta\,\mathbf{g}^{l}\Big|_{90\mathrm{s}}, \tag{7.19}$$

at each time epoch, *k*. The values of specific force f_e, f_n, and f_u are derived from DGPS as they were in Section 7.2. As shown, the error in the estimate of the gravity disturbance is estimated from the difference between the solutions obtained by low-pass filtering to 30 and 90 s, respectively. As suggested in Schwarz and Wei (1997), the misalignment states are modeled as first-order Gauss Markov processes, meaning that their correlation decays exponentially. According to Gelb (1974), the appropriate transition matrix is given as follows:

$$\mathbf{\Phi}_k = \begin{pmatrix} e^{-\beta\Delta t} & 0 & 0 \\ 0 & e^{-\beta\Delta t} & 0 \\ 0 & 0 & e^{-\beta\Delta t} \end{pmatrix}, \tag{7.20}$$

where Δt is the sampling period, and the correlation period, *T*, of the process is given by $T = 1/\beta$. The corresponding process noise matrix is given by the following:

$$\mathbf{Q}_k = \begin{pmatrix} \sigma_x(1-e^{-2\beta\Delta t}) & 0 & 0 \\ 0 & \sigma_x(1-e^{-2\beta\Delta t}) & 0 \\ 0 & 0 & \sigma_x(1-e^{-2\beta\Delta t}) \end{pmatrix}, \tag{7.21}$$

where σ_x is the standard deviation of the misalignment state elements. Further, the initial state and associated covariance matrix are given by:

$$\mathbf{x}_0 = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \text{ and} \tag{7.22}$$

$$\mathbf{P}_0 = \begin{pmatrix} \sigma_x{}^2 & 0 & 0 \\ 0 & \sigma_x{}^2 & 0 \\ 0 & 0 & \sigma_x{}^2 \end{pmatrix}. \tag{7.23}$$

The covariance matrix of the measurements is given by:

$$\mathbf{R}_k = \begin{pmatrix} \sigma_h{}^2 & 0 & 0 \\ 0 & \sigma_h{}^2 & 0 \\ 0 & 0 & \sigma_v{}^2 \end{pmatrix}, \tag{7.24}$$

where σ_h and σ_v are the standard deviations of the horizontal and vertical components of the measurement vector, **z**.

The hypothesis is tested using the first four flight lines of data collected on September 10 of the Kananaskis campaign. The following values were used for the test:

T = 14400 s,

σ_x = 10 arc seconds,

σ_h = 25 mGal, and σ_v = 2 mGal.

The estimated misalignment state histories are shown in Figure 7.18 for the four flight lines. These agree well with the values that were presented in Skaloud (1999). Also, this confirms that the value of 30 arc seconds that was assumed for the horizontal misalignments in Section 7.3 was roughly correct. Following the modeling of these misalignment terms, the measured specific force of the system can be converted from the b-frame to the true l-frame using the following equation:

$$\mathbf{f}^{\mathrm{l}} = \mathbf{R}_{\mathrm{b}}^{\mathrm{l_{true}}} \mathbf{f}^{\mathrm{b}}, \tag{7.25}$$

where $\mathbf{R}_{\mathrm{b}}^{\mathrm{l_{true}}}$ is the rotation matrix between the b-frame and the true l-frame and is given by:

$$\mathbf{R}_{\mathrm{b}}^{\mathrm{l_{true}}} = \mathbf{R}_{\mathrm{l}}^{\mathrm{l_{true}}} \mathbf{R}_{\mathrm{b}}^{\mathrm{l}}, \tag{7.26}$$

where in turn, $\mathbf{R}_{\mathrm{l}}^{\mathrm{l_{true}}}$ is given by the misalignments as follows:

$$\mathbf{R}_{\mathrm{l}}^{\mathrm{l_{true}}} = \mathbf{R}_{\mathrm{z}}(-\varepsilon_{\mathrm{u}})\mathbf{R}_{\mathrm{y}}(-\varepsilon_{\mathrm{n}})\mathbf{R}_{\mathrm{x}}(-\varepsilon_{\mathrm{e}}). \tag{7.27}$$

The rotated specific force is then used as usual in equation 2.3 to yield an improved estimate of the gravity disturbance. The results of this are summarized in Table 7.1 that shows the agreement of the resulting solution with the upward continued reference both before and after correction for the misalignment terms. Also consider Figure 7.19 that demonstrates the improvement in the time domain.

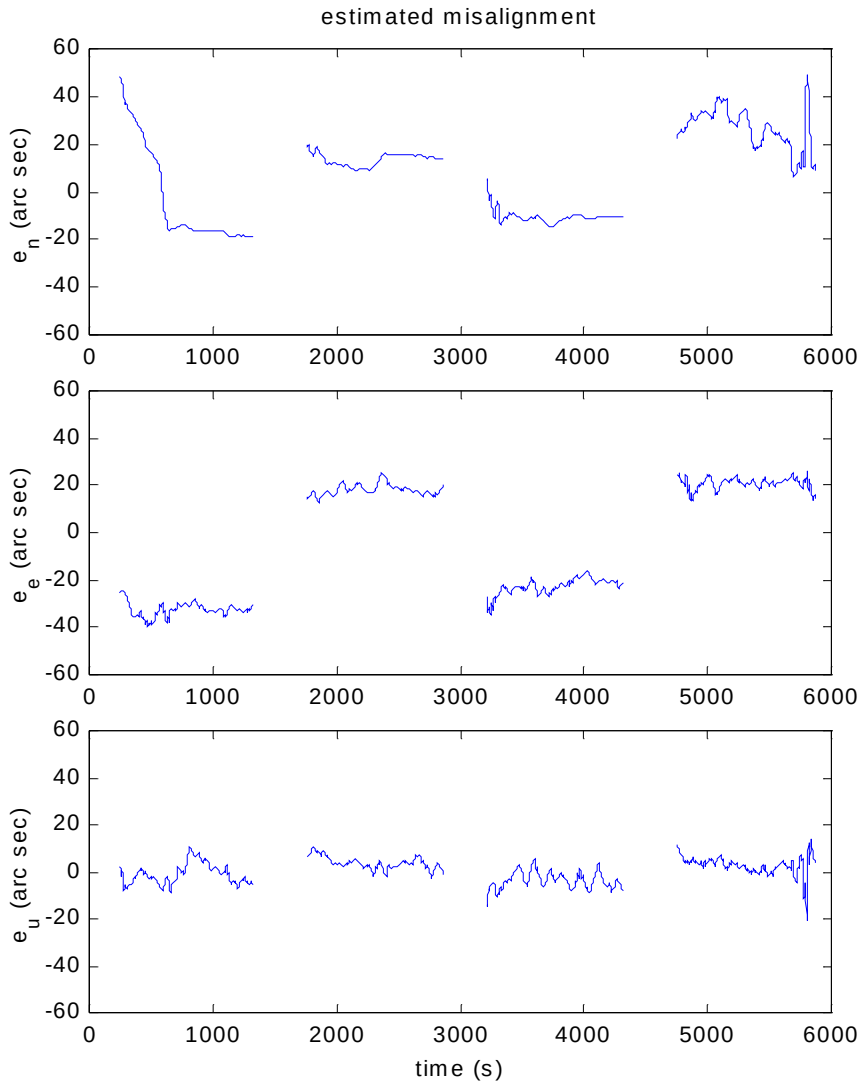

Figure 7.18 Estimated misalignments for the Kananaskis campaign

Table 7.1 Standard deviation of the agreement with reference after filtering to 30 s

Flight line	Residual (mGal)	
	original	corrected
1	11.5	5.5
2	9.0	3.6
3	11.9	5.8
4	13.9	7.8
average	11.6	5.7

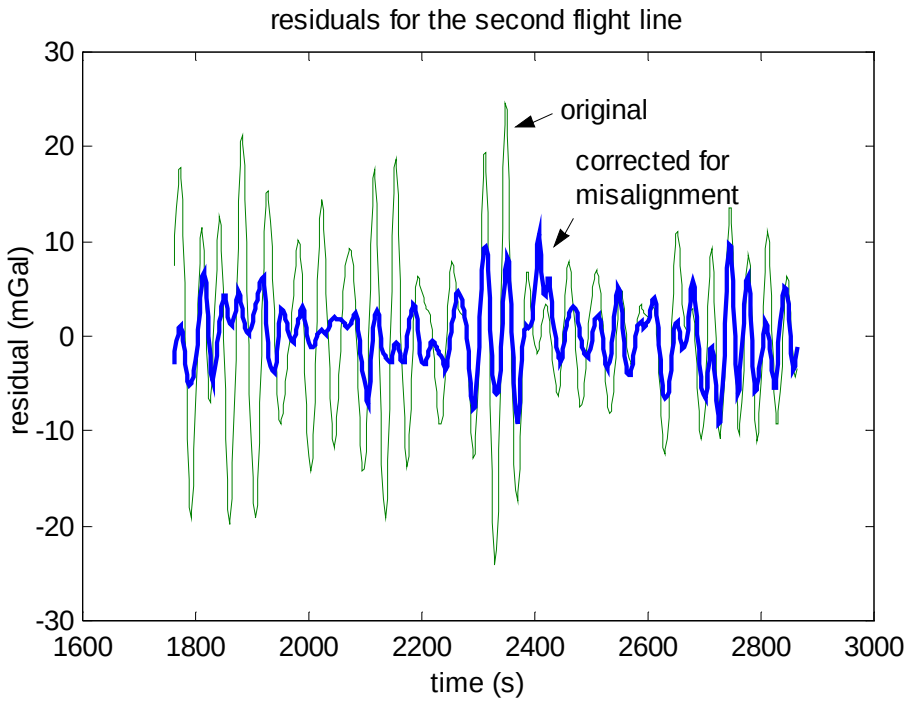

Figure 7.19 Samples of the residuals, before and after modeling of the misalignment

This preliminary result is very encouraging. It shows that it is indeed possible to model and estimate the error terms in an airborne gravity system, effectively removing the effect of the observed Phugoid motion. As a result, the standard deviation of the difference between the estimated gravity disturbance and the reference is approximately halved to just below 6 mGal. Because the aircraft was flying at a speed of 100 m/s, the corresponding resolution is a half-wavelength of 1.5 km. The remaining errors are due to errors in the upward continued reference, DGPS noise, synchronization errors and probably residual scale factor errors.

7.5 Summary

The errors affecting an airborne gravity system within the bandwidth corresponding to high-resolution gravity field estimation were dealt with in this chapter. The natural motion of the aircraft was studied in detail and the components of that motion that cause errors in the spectrum relevant to airborne gravimetry were identified as being due mainly to the Phugoid mode. Using real data, it was then seen that a pilot and autopilot can significantly reduce the magnitude of these dynamics. It was also seen that they still

exist at potentially dangerous levels despite this. In the studies presented here, the pilot did a better job of controlling the motion. It would be very useful to evaluate the use of a more sophisticated autopilot for controlling the motion. It is certainly not possible to say which aircraft is better in this case because it appears to depend heavily on the form of control used.

The relationships between the residual dynamics and the errors in gravity were then studied. This provides a good understanding of the behaviour of each error source in the frequency domain.

Methods of reducing the dynamics induced errors were discussed in the last section. This called for future studies into reducing the dynamics themselves and echoed calls made in other chapters for better sensors. Finally, the modeling of residual errors was discussed and a method of estimating the misalignment errors was proposed and implemented. It was shown to improve the accuracy of the SINS airborne gravity system by a factor of two at a very high resolution.

This suggests a possible compromise between new sensors and modeling. Because gyroscopes are generally the most expensive sensor set, it seems more convenient to treat the misalignment terms by modeling. Because accelerometers are cheaper and clearly the weakest component (especially considering the arguments about their low-resolution performance made in Chapter 6), their replacement is advised.

It is also concluded that the pursuit of methods for modeling SINS errors that are induced by dynamics should be a priority for future work. The model should incorporate scale factor errors that are likely to be problematic during vehicle turns.

Concluding Remarks

The primary objective of the research described in this dissertation is to simultaneously improve both the accuracy and the resolution of the SINS/DGPS gravimeter that is used by the University of Calgary. This goal has been met.

The major contributions of the research include a detailed analysis of the behaviour of the system and all of the components of the error budget under various dynamic conditions, an improvement of the performance of the system in the bandwidth corresponding to medium-resolution applications such as geoid determination, and a demonstration of the good performance of the system in the bandwidth corresponding to high-resolution applications such as resource exploration.

In the following, more detailed comments are made about the above contributions and improvements resulting from the research. Important results are highlighted, conclusions are drawn and recommendations are given. More detailed conclusions can also be found at the end of each chapter.

Specific contributions

The major contributions of the research include:

a) A frequency domain analysis of the errors affecting a SINS/DGPS airborne gravity system. Using a number of operational strategies, data processing strategies and estimation techniques, **error spectra have been estimated for each portion of the error budget of the estimated gravity disturbance**. This includes error spectra for each of the SINS errors in equation 2.7. In addition, a detailed understanding has been obtained about the effect that each of the DGPS errors in equation 2.14 has on the determination of acceleration. The spectra are derived for frequencies below 0.05 Hz, thereby offering a complete definition of the behaviour of the system for half-wavelengths as short as 450 m for a flying speed of 45 m/s.

b) The **accuracy of airborne positioning using carrier phase DGPS** has been assessed in Chapter 3. The assessment includes analyses that categorically isolate and quantify the effects of each portion of the error budget given by equation 2.14. Because the accuracy of this type of kinematic positioning is often quoted to be at or below the decimetre level, these analyses contribute to research in this field by:

- providing an evaluation of the observed positioning accuracy for a variety of realistic conditions during airborne gravity campaigns and showing that an accuracy of a decimetre cannot be assured, and
- confirming that DGPS is capable of comfortably providing the positioning accuracy of 50 cm that is required for airborne gravimetry, provided that certain conditions are not violated. See Section 3.4 for details.

c) **A fundamental consideration of the various processes of differentiation** has been given in Chapter 4 and particular differentiating filters have been proposed for the problems at hand. The performance and implementation of these filters has been investigated with real DGPS data that was collected under three realistic and different conditions: that is, for static receivers, for low dynamics and for dynamics with high frequency content. In each case, independent means were employed to assess the performance of the methods. The following conclusions have been reached:

- Appropriate differentiation methods must be employed, depending on a careful consideration of the vehicle dynamics and of the frequency domain characteristics of the differentiating filters.
- Simple methods of differentiation (such as those based on low-order Taylor series approximations and curve fitting) are only appropriate for use in static or very low-dynamic applications. Otherwise, their use is likely to lead to incorrect estimates of velocity and acceleration and therefore lead to major errors.
- Using a precise motion table, it has been verified that the velocity accuracy of 5 cm/s that is required for airborne gravimetry can be comfortably achieved.

d) The topic of **optimal estimation of acceleration for airborne gravimetry** has been treated in Chapter 5 where specific emphasis is placed on estimating acceleration for geoid determination and resource exploration. In addition to the error spectra that are mentioned in a) above, major conclusions of this research include a set of recommendations for estimating acceleration for each of these areas of applications. These recommendations are given in Section 5.4.

e) A detailed **analysis of the low-frequency biases that affect a SINS airborne gravity system** is given in Chapter 6. The limitations of the SINS in estimating relative gravimetry are quantified in terms of maximum spatial resolution. The SINS gravimeter has been shown to be useful for half-wavelengths shorter than 200 km and to perform best for half-wavelengths shorter than 100 km.

   A simple method has been proposed and implemented for combining the low-resolution information from a geopotential model with the medium and high-resolution information from the SINS gravimeter. Although the SINS that is used by the University of Calgary has been shown to be insufficiently accurate to be used in this way in all geographic areas, Intermap Technologies Corporation now uses it as part of the services they provide.

f) A thorough **investigation of the motion experienced by survey aircraft** has been given in Section 7.1 and Section 7.2. This study is in response to the need for a better understanding of the dynamics that the aircraft undergoes, including Phugoid motion. The observed dynamics of several survey aircraft have been analyzed in detail under conditions when the natural motion is either controlled by an autopilot or by a pilot.

g) A detailed **analysis of the high-frequency errors affecting a SINS airborne gravity system** is the subject of Sections 7.2 and 7.3. This analysis characterizes each error in the frequency domain, as discussed in a) above. The misalignment error has been confirmed to be a major barrier to high-resolution gravity field estimation when aircraft dynamics are high.

h) A method for **modeling and removing the misalignment** errors is given in Section 7.4. This increases the accuracy of the estimated gravity disturbance by a

factor of two for the Kananaskis data, to 6 mGal, for a spatial resolution of 1.5 km.

Demonstrated performance of the SINS/DGPS gravimeter

The performance of the SINS/DGPS gravimeter had been demonstrated in Wei and Schwarz (1998) to be 2-3 mGal for medium-resolution applications (i.e. having a half-wavelength of 5 km). The same level of performance was demonstrated in Glennie (1999). The research presented herein has demonstrated an improvement in the performance of the system for medium-resolution applications and has also demonstrated for the first time that the system performs very well for high-resolution applications. Important highlights of the new results include the following:

a) Using data from the Alexandria campaign, which was designed in part to test the performance of the system for high-resolution applications, it has been demonstrated that the performance of the SINS/DGPS gravimeter is at the level of 1.5 mGal for a resolution of 2.0 km and at the level of 2.5 mGal for a resolution of 1.4 km. See Section 5.5 for more details.
b) Using data from the Kananaskis campaign, it has been shown that the effect of Phugoid motion on a SINS gravimeter can be largely removed by employing the method that is proposed in Section 7.4 for modeling misalignment errors.
c) The performance of the system for the Kananaskis campaign has been improved by using an ionospheric-free approach for the DGPS processing. See Section 5.2.4 for details.
d) An improvement of about 40% has been demonstrated for the SINS data collected during the Greenland campaign, when compared to the shipborne reference data (see Appendix E). Agreement with that data is at the level of 1.3 and 2.0 mGal for the two flight lines where the reference was available. The performance of the SINS has been demonstrated to be slightly better than that of the LaCoste and Romberg for that campaign.

Recommendations

Finally, several interesting points have come to light as a result of the work presented herein, discussions with colleagues and discussions with the examining committee. These

form the basis of the following brief list of recommendations that may in turn form the basis of future work.

- In Chapter 3, the accuracy of current DGPS methods was evaluated for airborne positioning. It is recommended that for evaluation purposes, a future field campaign be planned to include an independent reference for the position of the aircraft. This might make use of photogrammetry or LIDAR, for example, and could allow for the further quantification of the influence of individual error sources, especially the differential troposphere.

  Also, it is recommended that the applicability of future Global Navigation Satellite Systems (GNSS) for positioning in airborne gravimetry is evaluated. This includes GPS modernization efforts (a new civil coarse- acquisition code on L2 and eventually a new civil frequency). This also includes the future Galileo system and combinations of such systems.
- In Chapter 4, the accuracy of current DGPS methods was evaluated for velocity and acceleration determination. This evaluation is one of few in the literature that uses real (not simulated) kinematic data and an independent truth for evaluation. Because real data often reflects realistic levels of performance, this approach is recommended for future work whenever possible. It is recommended that different receiver types be evaluated under such conditions in order to quantify the effects of the different data processing strategies they employ (to derive a Doppler and to reduce noise and multipath).
- In Chapter 5, the accuracy of current DGPS methods was evaluated for acceleration determination. This used examples from recent campaigns to quantify the components of the DGPS error budget as a function of frequency for half-wavelengths as short as 450 m. It was seen there, that major road blocks to widening the bandwidth of all airborne gravimetry systems include the effects of changing satellite geometry and receiver noise (especially if an ionospheric-free approach is required). It is recommended that, at least, the following research be conducted into reducing these error sources:

a) It was seen in Chapter 4 that accurate estimates of acceleration can be obtained using estimates of velocity that are based on a derived Doppler. Because this approach is simple (relative to the traditional approach of deriving position and differentiating it), it is recommended that future research be conducted into the optimal use of the derived Doppler for acceleration determination. Optimization should consider the requirements to minimize noise and the effects of changes in geometry, especially as the bandwidth is widened to include high-resolution applications.

b) Evaluation of the potential of GPS modernization efforts for improving the accuracy with which acceleration can be determined (e.g. the coarse-acquisition code on L2 might result in a better signal-to-noise ratio for the L2 carrier and in turn better ionospheric-free results).

- It remains a challenge to combine the relative information from airborne gravimetry with other sources of information about the gravity field. This was made clear in Chapter 6 and will become increasingly important as data from future satellite missions becomes available. Efforts in this direction are recommended.
- The work presented in this dissertation has resulted in an increased understanding of the frequency domain behaviour of the error sources influencing airborne gravimetry. This should form the basis of a set of standard error spectra for understanding and predicting the performance of airborne gravity systems under a variety of expected conditions and operational ranges. Also, it is felt that this detailed frequency domain analysis opens the door for appropriate statistical modeling of the remaining errors. A study of this approach to noise reduction is recommended as a complement to current band-limitation techniques.
- Future designs based on SINS technologies should include accelerometers of higher quality than those in the LRF-III system. Especially important is the compensation for changes in temperature.

- Further studies into the modeling of SINS/DGPS system errors should be conducted along the lines of the example that is presented in Section 7.4. This might include adding a model for scale factor errors.
- Finally, a general recommendation is reiterated here that has been made on several occasions by researchers at the University of Calgary over the last few years. The SINS and accelerometer triad approaches to airborne gravimetry have been proven to provide very accurate and high-resolution information about the relative gravity field. Because of this and because of the significant practical advantages of such systems (e.g. size, power consumption and especially cost), they should be the basis of future airborne gravimetry systems.

Overall, it is exciting and encouraging to see that high-resolution surveys of the gravity field can be carried out with good accuracy using the SINS/DGPS gravimeter. It is sincerely hoped that this research contributes to future work in this field.

Bibliography

Antoniou, A. (1993) Digital Filters: Analysis, Design and Applications (689 pp.) New York: McGraw-Hill Inc.

Argeseanu, V. (1995) *Upward Continuation of Surface Gravity Anomaly Data*, In Proceedings IAG Symposium G4 Airborne Gravity Field Determination, XXI General Assembly of the IUGG, Boulder Colorado, July 2-14, 1995, pp. 95-102 (Published by the University of Calgary).

Babister, A.W. (1980) *Aircraft Dynamic Stability and Response* (220 pp.) Permagon Press Ltd., Oxford, England.

Bastos, L. P. Tome, T. Cunha and S. Cunha (2000) *Gravity Anomalies from Airborne Measurements – Experiments Using a Low Cost IMU*, In Proceedings of IAG International Symposium on Gravity, Geoid and Geodynamics 2000, July 31-August 4 2000, Banff, Canada.

Bell, R.E., R. Anderson and L. Pratson (1997) *Gravity Gradiometry Resurfaces*, The Leading Edge, January, pp. 55-59.

Beutler, G. (1998) *GPS Satellite Orbits*. In GPS for Geodesy, Lecture Notes in Earth Sciences. In GPS for Geodesy, Lecture Notes in Earth Sciences, 2nd Edition, Teunissen and Kleusberg (Eds.), Springer Verlag, pp. 43-110.

Boedecker, G. (1998) *Configuration Optimization of Accelerometers for Strapdown Airborne Gravimeters*, Presented at International Workshop on Airborne Gravity and the Polar Gravity Field, Kangerlussuaq, Greenland, June 2-4, 1998.

Bona, P. and C. Tiberius (2000) *An Experimental Comparison of Noise Characteristics of Seven High-End Dual Frequency GPS Receiver-Sets*. Proceedings of PLANS 2000. San Diego, US. March 13-16, 2000: pp. 237-244.

Brown, J.M., E. Klingele, T.M. Niebauer (2000) *Towards an Airborne Absolute Gravity System*, In Proceedings of IAG International Symposium on Gravity, Geoid and Geodynamics 2000, July 31-August 4 2000, Banff, Canada.

Brozena, J.M. (1992) *The Greenland Aerogeophysics Project: Airborne Gravity, Topographic and Magnetic Mapping of an Entire Continent*. In Proc. IAG Symposium G3: Determination of the Gravity Field by Space and Airborne Methods, Vienna, Aug. 1991, Springer Verlag, 1992.

Brozena, J. and V.A. Childers (2000) *The NRL Airborne Geophysics Program*, In International Association of Geodesy Symposia, Vol. 121, Schwarz (ed.): Geodesy Beyond 2000 – The Challenges of the First Decade, Springer, Berlin, Heidelburg, pp. 126-130.

Brozena, J., V.A. Childers and J. Blaha (2000) *Applications of Airborne Gravity and Sea-Surface Topography to Coastal Oceanography*, In Proceedings of IAG International Symposium on Gravity, Geoid and Geodynamics 2000, July 31-August 4 2000, Banff, Canada.

Brozena, J.M., Mader G.L. & Peters M.F. (1989) *Interferometric Global Positioning System: Three-Dimensional Positioning Source for Airborne Gravimetry*, Journal of Geophysical Research, Vol. 94, No. B9, 12153-12162.

Brozena, J.M., M.F. Peters and R. Salman (1997) *Acrtic Airborne Gravity Measurement Program*, In Proc. of the Int. Symposium Gravity, Geoid and Marine Geodesy (GraGeoMar96), Tokyo, September 20 – October 5, pp. 131-138.

Brunner, F.K. (1997) Atmospheric Propagation Effects in Geomatics. Lecture Notes. TU Graz. Presented at the University of Calgary, June 9-19, 1997.

Bruton, A.M. (1997) *Reduction of GPS Receiver Noise Using Adaptive Filters*, In Proceedings of the 101th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GPS 97), pp. 645-653.

Bruton, A.M., C. Glennie and K.P. Schwarz (1999) *Differentiation for High-Precision GPS Velocity and Acceleration Determination*. GPS Solutions, Vol. 2, No. 4, 7-21.

Bruton, A.M., M. Kern, K.P. Schwarz, S. Ferguson, A. Simsky, K. Tennant, M. Wei, J. Halpenny, R. Langley, T. Beran, K. Keller, H. Martell, P. Mrstk, K. Kusevic, R. Faulkner (2000b) *On the Accuracy of Kinematic Carrier Phase DGPS for Airborne Mapping*, Accepted by Geomatica, September 2000.

Bruton, A. M. and K.P. Schwarz (1997) *Airborne Gravity Estimation using Adaptive Filters*, In Proceedings of International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97), Banff, Canada, June 3-6, pp. 605-612.

Bruton, A.M., K.P. Schwarz, S. Ferguson, M.Kern and M. Wei (2000a) *Deriving Acceleration from DGPS: Toward Higher Resolution Applications of Airborne Gravimetry*, Submitted to GPS Solutions, December, 2000.

Bryson, A.E. (1994) Control of spacecraft and aircraft (378 pp.) Princeton University Press, Princeton, New Jersey.

Cannon, M. E., G. Lachapelle, M.C. Szarmes, J.M. Hebert, J. Keith & S. Jokerst (1997) *DGPS Kinematic Carrier Phase Signal Simulation Analysis for Precise Velocity and Position Determination*, Journal of the Institute of Navigation, Vol 44, No. 2, 231-245.

Cannon, M.E., K.P. Schwarz and M. Wei (1992) *A Consistency Test of Airborne GPS Using Multiple Monitor Stations*, Bulletin Geodesique, Vol. 66, No. 1, pp.2-11.

Cheney, W., & D. Kincaid (1994) Numerical Methods and Computing, 3rd Edition (578 pp.). Pacific Grove, California: Brooks/Cole Publishing Company.

Childers, V.A., R.E. Bell and J.M. Brozena (1999) *Airborne Gravimetry: An Investigation of Filtering*, Geophysics, Vol. 64, No. 1, pp. 61-69.

Collins, J.P. and R. Langley (1999) *Nominal and Extreme Error Performance of the UNB3 Tropospheric Delay Model*. Final contract report for Nav Canada Satellite Navigation Program Office, Department of Geodesy and Geomatics Engineering, UNB, Technical Report No. 204, 173 pp.

Cunha, S.R. (1998) *On the INS and GPS Data Processing For Airborne Gravimetry Applications*, Presented at International Workshop on Airborne Gravity and the Polar Gravity Field, Kangerlussuaq, Greenland, June 2-4, 1998.

Czompo, J. (1991) *Acceleration Determination using Kinematic GPS Measurements*, In Proceedings of the 4th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GPS 91).

Czompo, J. (1994) *Airborne Scalar Gravimetry Systems in the Spectral Domain*, Ph.D. Thesis, UCGE Report No. 20067, Department of Geomatics Engineering at the University of Calgary, Calgary, Canada.

Czompo, J. and S. Ferguson (1995) *Design Considerations for a New Scalar Gravity Meter for Airborne Surveys*. In Proceedings IAG Symposium G4 Airborne Gravity Field Determination, XXI General Assembly of the IUGG, Boulder Colorado, July 2-14, 1995, pp. 3-11 (Published by the University of Calgary).

El-Sheimy, N. and K.P. Schwarz (2000) Kingpsad Users Manual, V.3.0, Department of Geomatics Engineering, University of Calgary, April, 2000.

ESA (1999) The Four Candidate Earth Explorer Core Missions – Gravity Field and Steady-State Ocean Circulation, Report ESA SP-1233 (1), ESA Publications Division, Noordwijk, The Netherlands.

Etkin, B. (1982) Dynamics of Flight – Stability and Control (370 pp.) John Wiley & Sons, Inc., Toronto.

Fenton, P. & B. Townsend (1994) *NovAtel Communications Ltd. – What's New?* In Proceedings of International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS94), Banff, Canada, pp. 25-29.

Ferguson, S. and Y. Hammada (2000) *Experiences with AIRGrav: Results from a New Airborne Gravimeter*, In Proceedings of IAG International Symposium on Gravity, Geoid and Geodynamics 2000, July 31-August 4 2000, Banff, Canada.

Fernandes, M.J., L. Bastos, R. Forsberg, A. Oleson, F. Leite (2000) *Geoid Modeling in Coastal Regions Using Airborne and Satellite Data: Case Study in the Azores*, In International Association of Geodesy Symposia, Vol. 121, Schwarz (ed.): Geodesy Beyond 2000 – The Challenges of the First Decade, Springer, Berlin, Heidelburg, pp. 112-117.

Forsberg, R. and S. Kenyon (1995) *Downward Continuation of Airborne Gravity Data*, In Proceedings IAG Symposium G4 Airborne Gravity Field Determination, XXI General Assembly of the IUGG, Boulder Colorado, July 2-14, 1995, pp. 73-80 (Published by the University of Calgary).

Forsberg, R., A.V. Oleson and K. Keller (1999) Airborne Gravity Survey of the North Greenland Shelf 1998, Technical Report 10, Danish National Survey and Cadastre, Denmark.

Froud, B. (1998) Good Faeries, Bad Faeries (192 pp.) Simon and Schuster, 1998.

Glennie, C. (1999) An Analysis of Airborne Gravity by Strapdown INS/DGPS, Ph.D. Thesis, Department of Geomatics Engineering, The University of Calgary, UCGE Report No. 20125.

Glennie, C. and K.P. Schwarz (1997) *Airborne Gravity by Strapdown INS/DGPS in a 100 km by 100 km Area of the Rocky Mountains*. In Proceedings of International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97), Banff, Canada, June 3-6, pp. 619-624.

Glennie, C. and K.P. Schwarz (1999) *A Comparison and Analysis of Airborne Gravimetry Results from Two Strapdown Inertial/DGPS Systems*, Journal of Geodesy, 73, 6, pp. 311-321.

Glennie, C., K.P. Schwarz, A.M. Bruton, R. Forsberg, A.V. Oleson, K. Keller (2000) *Comparison of Stable Platform and Strapdown Airborne Gravity*, Journal of Geodesy, 74, 5, pp. 383-389.

Gumert, W. R. (1998) *An Historical Review of Airborne Gravity*, The Leading Edge, January.

Guterman, M.M. and Z.H. Nitecki (1991) Differential Equations: A First Course, 3rd Edition, Saunders College Publishing, US.

Hammada, Y. (1996) *A Comparison of Filtering Techniques for Airborne Gravimetry*, M.Sc. Thesis, Department of Geomatics Engineering, The University of Calgary, UCGE Report No. 20089.

Hammada, Y. (1997) *Optimal Versus Non-Optimal Low-pass Filtering in Airborne Gravimetry*, In Proceedings of International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97), Banff, Canada, June 3-6, pp. 633-640.

Han, S.-C., C. Jekeli and J.H. Kwon (2000) *Precision Absolute GPS Positioning and its Use to Obtain Kinematic Acceleration*, In Proceedings of IAG International Symposium on Gravity, Geoid and Geodynamics 2000, July 31-August 4 2000, Banff, Canada.

Han, S., C. Rizos and R. Abbot (1998) *Flight Testing and Data Analysis of Airborne GPS LADS Survey*, In Proceedings of the 11th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GPS 98), pp. 1211-1221.

Harrison, J.C., J.D. MacQueen, A.C. Rauhut and J.Y. Cruz (1995) The LCT Airborne Gravity System, In Proceedings IAG Symposium G4 Airborne Gravity Field Determination, XXI General Assembly of the IUGG, Boulder Colorado, July 2-14, 1995, pp. 163-168 (Published by the University of Calgary).

Hartinger, H. and F. Brunner (1999) *Variances of GPS Phase Observations: The Sigma-ε Model*, GPS Solutions, Vol. 2, No. 4, 35-43.

Hebert, J. (1997) *High Accuracy GPS Velocity Using the Carrier Phase Observable*, In Proceedings of International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97), Banff, Canada, June 3-6, pp. 265-269.

Heffley, R.K. and W.F. Jewell (1972) Aircraft Handling Qualities, Systems Technology Inc., Technical Report 1004-1, Hawthorne, CA, US.

Hehl, K., G.W. Hein, H. Landau, M. Ertel, J. Fritsch and P. Kewitsch (1990) *An Integrated Precise Airborne Navigation and Gravity Recovery System – Verification of GPS-Determined Vertical Disturbing Acceleration*, In Proceedings of International Symposium in Geodesy, Surveying and Remote Sensing (KIS90), Banff, Canada, September 10-13, pp. 477-487.

Hein, G.W., K. Hehl, H. Landau and M. Ertel (1990) Experiments for an Integrated Precise Airborne Navigation and Gravity Recovery System, Proceedings of IEEE PLANS 1990. pp. 279-285.

Hein, G.W. (1995) *Progress in Airborne Gravimetry: Solved, Open and Critical Problems*. In Proceedings IAG Symposium G4 Airborne Gravity Field Determination, XXI General Assembly of the IUGG, Boulder Colorado, July 2-14, 1995, pp. 3-11 (Published by the University of Calgary).

Heiskanen, W.A. and H. Moritz (1969) Physical Geodesy, W.H. Freeman and Co., San Francisco, 1967.

Hoffmann-Wellenhof, B., H. Lichtenegger, J. Collins (1994) GPS Theory and Practice, 3rd Edition, Springer-Verlag Wien, New York.

Jekeli, C. and R. Garcia (1997) *GPS Phase Acceleration for Moving-Base Vector Gravimetry*, Journal of Geodesy, 71, pp 630-639.

Jekeli, C. and J.H. Kwon (1999) *Results of Airborne Vector (3-D) Gravimetry*, Geophysical Research Letters, Vol. 26, No. 23, pp 3533-3536, December 1.

Johnson, E.A.E. (1998) *Gravity and Magnetic Analyses Can Address Various Petroleum Issues*, The Leading Edge, January.

Jones, P. (1997) *Towards a New Free Air Anomaly Map of the Arctic Peninsula*, In Forsberg et al. (Eds.): Geodesy on the Move, IAG Symposia, Vol. 119, pp. 517-522, Springer Verlag.

Kennedy, S. (1999) Progress of Crossover Difference Post-Mission Processing, Internal Unpublished Report, Department of Geomatics Engineering, University of Calgary.

Kleusberg, A., D. Peyton and D. Wells (1990) *Airborne Gravity and the Global Positioning System*, in Proceedings of IEEE PLANS 90, pp. 394-401.

Klingele, E. E., M. Cocard, M.E. Halliday, H.-G. Kahle (1995) The Airborne Gravimetric Survey of Switzerland. Contribution to the Geology of Switzerland Series, Swiss Geophysical Commission, Report No. 31.

Klobuchar, J.A. 1996. *Ionospheric Effects in GPS*. In Global Positioning System: Theory and Applications. American Institute of Aeronautics and Astronautics, Inc., Vol. I: 485-515.

Lachapelle, G. (1998) GPS Theory and Applications, ENGO 625 Lecture Notes, Department of Geomatics Engineering, The University of Calgary.

Langley, R. (1998) *GPS Receivers and the Observables*, In GPS for Geodesy, Lecture Notes in Earth Sciences, 2nd Edition, Teunissen and Kleusberg (Eds.), Springer Verlag, pp. 151-185.

Lemoine, F.G., S.C. Kenyon, J.K. Factor, R.G. Trimmer, N.K. Pavlis, D.S. Chinn, C.M. Cox, S.M. Klosko, S.B. Luthcke, M.H. Torrence, Y.M. Wang, R.G. Williamson, E.C. Pavlis, R.H. Rapp, T.R. Olson (1998) The Development of the Joint NASA GSFC and the National Imagery and Mapping Agency (NIMA) Geopotential Model EGM96, NASA TP-1998-206861.

Li, Y. (2000) Airborne Gravimetry and Geoid Determination, Ph.D. Thesis, Department of Geomatics Engineering, The University of Calgary, UCGE Report No. 20141.

Liu, Z. (1992) Comparison of Statistical Methods for the Alignment of Strapdown Inertial Systems, Ph.D. Thesis, Department of Geomatics Engineering, The University of Calgary, UCGE Report No. 20047.

Mendes, V.B. (1998) Modeling the Neutral Atmosphere Propagation Delay in Radiometric Space Techniques, Ph.D. Dissertation, Department of Geodesy and Geomatics Engineering, University of New Brunswick, Fredericton, New Brunswick.

Mendes, V.B. and R. Langley (1998) *Tropospheric Zenith Delay Prediction Accuracy for Airborne GPS High-Precision Positioning*, In Proceedings of 54th Annual Meeting of the ION, pp. 337-347.

Mostafa, M. M. R. and K.P. Schwarz (1999) *An Autonomous Multi-Sensor System for Airborne Digital Image Capture and Georeferencing*, Proceedings of the ASPRS Annual Convention, Portland, Oregon, May 17-21, pp. 976-987.

Novak, P., M. Kern and K.P. Schwarz (2000a) Continuation of Airborne Gravity, Unpublished internal report. Department of Geomatics Engineering. University of Calgary, Calgary, Canada.

Novak, P., M. Kern and K.P. Schwarz (2000b) *On the Determination of the Relative Geoid from Airborne Gravimetry*, In Proceedings of IAG International Symposium on Gravity, Geoid and Geodynamics 2000, July 31-August 4 2000, Banff, Canada.

NRC (National Research Council) (1995) Airborne Geophysics and Precise Positioning: Scientific Issues and Future Direction. Proceedings of the Workshop on Airborne Geophysics, July 12-14, 1993, Washington, D.C., National Academy Press.

NRC (National Research Council) (1997) Satellite Gravity and the Geosphere: Contributions to the Study of the Solid Earth and Its Fluid Envelope. National Research Council, Washington, D.C., National Academy Press.

Oppenheim, A.V. and R.W. Schafer (1989) Discrete-Time Signal Processing, Prentice Hall Signal Processing Series, Prentice Hall, New Jersey, US.

Orfanidis, S.J. (1996) Introduction to Signal Processing, (798 pp.). Toronto: Prentice Hall Signal Processing Series.

Parkinson, B.W. and P.K. Enge. (1996) *Differential GPS*. In GPS: Theory and Applications, American Institute of Aeronautics and Astronautics, Inc., Vol. II: 3-50.

Parkinson, B.W. and J.J. Spilker (Eds.) (1996) GPS: Theory and Applications, American Institute of Aeronautics and Astronautics Inc., (2 Volumes).

Press, W.H., Teukolsky S.A, Vetterling W.T. & Flannery B.P. (1992) Numerical Recipes in C, Second Edition (994 pp.). New York: Cambridge University Press.

Rapp, R. H. and N.K. Pavlis (1990) *The Development and Analysis of Geopotential Coefficient Models to Spherical Harmonic Degree 360*, Journal of Geophysical Research, Vol. 95, No. B13, pp. 21885-21911, December 10.

Raquet, J. (1998) Development of a Method for Kinematic GPS Carrier-Phase Ambiguity Resolution Using Multiple Reference Receivers, Ph.D. Thesis, Department of Geomatics Engineering, University of Calgary, UCGE Report No. 20116.

Ray, J. K. 2000. Mitigation of GPS Code and Carrier Phase Multipath Effects Using a Multi-Antenna System. Department of Geomatics Engineering, Ph.D. Thesis, The University of Calgary. UCGE Report No. 20136.

Reynolds, J. M. (1997) An Introduction to Applied and Environmental Geophysics, John Wiley & Sons, Chichester, England.

Salychev, O.S., A.V. Bykovsky, A.V. Voronov, K.P. Schwarz, Z. Liu, M. Wei (1994) *Determination of Gravity and Deflections of the Vertical for Geophysical Applications Using the ITC-2 Platform*. In Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation. Banff, Canada, August 30 – September 2, pp. 521-529.

Salychev, O.S. and K.P. Schwarz (1995) *Airborne Gravimetric Results Obtained with the ITC-2 Inertial Platform System*. In Proceedings IAG Symposium G4 Airborne Gravity Field Determination, XXI General Assembly of the IUGG, Boulder Colorado, July 2-14, 1995, pp. 125-141 (Published by the University of Calgary).

Schwarz, K.P. (1984) *Data Types and Their Spectral Properties*, In Proceedings of Local Gravity Field Approximation, Beijing International Summer School, Beijing, China, Aug 21 – Sept 4, 1984, pp. 1-66.

Schwarz, K.P. (1998) Sensor Integration and Image Georeferencing. Lecture Notes for Duane C. Brown International Summer School in Geomatics, The Ohio State University, Columbus, Ohio, USA, July 9-11.

Schwarz, K.P., M.E. Cannon and R.V.C. Wong (1989) *A Comparison of GPS Kinematic Models for the Determination of Position and Velocity Along a Trajectory*, Manuscripta Geodetica, Vol. 14, pp. 345-353.

Schwarz, K.P., O. Colombo, G. Hein and E.T. Knickmeyer (1991) *Requirements for Airborne Vector Gravimetry*, In Proceedings of IAG Symposium No. 110, From Mars to Greenland: Charting with Space and Airborne Instruments, Springer Verlag, New York, pp. 273-283.

Schwarz, K.P. and Y.C. Li (1996a) *What Can Airborne Gravimetry Contribute to Geoid Determination?* Journal of Geophysical Research, Vol. 101, B8, pp. 17873-17881.

Schwarz, K.P., Y.C. Li and M. Wei (1994) *The Spectral Window for Airborne Gravity and Geoid Determination.* In Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation. Banff, Canada, August 30 – September 2, pp. 445-456.

Schwarz, K.P. and Z. Li (1996b) An Introduction to Airborne Gravimetry and its Boundary Value Problems. Lecture Notes, IAG International Summer School, Como Italy, May 26-June 7.

Schwarz, K.P. and M. Wei (1990) *A Framework for Modeling Kinematic Measurements in Gravity Field Applications*. Bulletin Geodesique, 64, 331-346.

Schwarz, K.P. and M. Wei (1994) *An Error Analysis of Airborne Vector Gravimetry*, In Proceedings of the International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation. Banff, Canada, August 30 – September 2, pp. 509-520.

Schwarz, K.P. and M. Wei (1997) Inertial Surveying and INS/GPS Integration, ENGO 623 Lecture Notes, Department of Geomatics Engineering, The University of Calgary.

Shi, J. (1994) High Accuracy Airborne Differential GPS Positioning Using a Multi-Receiver Configuration, M.Sc. Thesis, Department of Geomatics Engineering, The University of Calgary, UCGE Report No. 20061.

Shin, E. H. (2000) IMU Calibration Without Aligning to the Local-Level Frame, Internal Unpublished Report for ENGO 623, Department of Geomatics Engineering, University of Calgary.

Sideris, M.G. (1993) The Gravity Field in Surveying and Geodesy, ENGO 527 Lecture Notes, Department of Geomatics Engineering, The University of Calgary.

Sideris, M.G. (1997) Geoid Determination by FFT Techniques, Lecture Notes in the IAG International School for the Determination and Use of the Geoid, Rio de Janeiro, September 10-16, 1997.

Sideris, M.G., M. Wei, M.E. Cannon and K.P. Schwarz (1992) *Airborne Gravimetry and Gravity Gradiometry for Geophysical Prospecting*, In Proc. of the International Workshop on Global Positioning Systems in Geosciences, Chania, Greece, June 8-10, pp. 327-346.

Skaloud, J. (1999) Optimizing Georeferencing of Airborne Survey Systems by INS/DGPS, Ph.D. Thesis, Department of Geomatics Engineering, The University of Calgary, UCGE Report No. 20126.

Skaloud, J. and B. Merminod (2000) *DGPS-Calibrated Accelerometric System for Dynamic Sports Events*, In Proceedings of the 13th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GPS 2000).

Skaloud, J. and K.P. Schwarz (1998) *Accurate Orientation for Airborne Mapping Systems*, ISPRS Commission II, Symposium on Data Integration: Systems and Techniques, Cambridge, UK, July 13-17, pp. 283-290.

Teper, G.L. (1969) Aircraft Stability and Control Data, Systems Technology Inc. Technical Report 176-1, Hawthorne, CA, US.

Teunissen, P.J.G. (1998) *GPS Carrier Phase Ambiguity Fixing Concepts*, In GPS for Geodesy, Lecture Notes in Earth Sciences, 2nd Edition, Teunissen and Kleusberg (Eds.), Springer Verlag, pp. 319-388.

Teunissen, P.J.G., A. Kleusberg (Eds.) (1998a) GPS for Geodesy, Lecture Notes in Earth Sciences, 2nd Edition, Teunissen and Kleusberg (Eds.), Springer Verlag, pp. 187-229.

Teunissen, P.J.G., A. Kleusberg (1998b) *GPS Observation Equations and Positioning Concepts*, In GPS for Geodesy, Lecture Notes in Earth Sciences, 2nd Edition, Teunissen and Kleusberg (Eds.), Springer Verlag, pp. 187-229.

Tiberius, C.C.J.M. (1998) Recursive Data Processing for Kinematic GPS Surveying, NCG-Nederlandse Commissie voor Geodesie, Delft, Netherlands.

Tiberius, C., N. Jonkman and F. Kenselaar (1999) *The Stochastics of GPS Observables*. Innovation Column, GPS World, February, 1999.

Tiemeyer, B. M.E. Cannon, G. Lachapelle and G. Schanzer (1994) *High Precision Aircraft Navigation with Emphasis on Atmospheric Effects*, Proceedings of IEEE PLANS 94, New York, pp. 394-401.

Timmen, L., L. Bastos, R. Forsberg, A. Giskehaug, U. Meyer (2000) *Airborne Gravity Field Surveying for Oceanography, Geology and Geodesy – The Experiences from AGMASCO*, In IAG Symposia, Vol. 121, Schwarz (ed.): Geodesy Beyond 2000 – The Challenges of the First Decade, Springer, Berlin, Heidelburg, pp. 118-123.

Titterton, D.H. and L.L. Weston (1997) Strapdown Inertial Navigation Technology, Peter Peregrinus Ltd., United Kingdom.

Tscherning, C.C., F. Rubek and R. Forsberg (1997) *Combining Airborne and Ground Gravity Using Collocation*, In IAG Symposia Vol 119, Forsberg et al. (Eds.): Geodesy on the move, Springer Verlag, pp. 18-23.

Van Dierendonck, J.K., M.E. Cannon, M. Wei and K.P. Schwarz (1994) *Error Sources in GPS-Determined Acceleration for Airborne Gravimetry*, In Proceedings of the ION National Technical Meeting, San Diego, CA, January 24-26.

Verroneau, M., P. Vanicek, P. Novak, J. Huang, J. Janak, M.G. Sideris, O. Esan (2000) *Canadian Gravimetric Geoid Model 2000 (CGG2000): Preliminary Results*, In Proceedings of IAG International Symposium on Gravity, Geoid and Geodynamics 2000, July 31-August 4 2000, Banff, Canada.

Wang, J. 1999. Modelling and Quality Control for Precise GPS and GLONASS Satellite Positioning. Ph.D. Thesis, School of Spatial Sciences, Curtin University of Technology, Perth, Australia, 171 pp.

Wei, M. (1999) *From Airborne Gravimetry to Airborne Geoid Mapping*, Report of IAG SSG 3.164. In Report of IAG Section III: Determination of the Gravity Field, compiled and edited by R. Forsberg, published by Danish National Survey and Cadastre (KMS), pp.25-32.

Wei, M., S. Ferguson and K.P. Schwarz (1991) *Accuracy of GPS-Derived Acceleration from Moving Platform Tests*. In Proceedings of IAG Symposium No. 110, From Mars to Greenland: Charting with Space and Airborne Instruments, Springer Verlag, New York, pp. 235-249.

Wei, M. and K.P. Schwarz (1994) *An Error Analysis of Airborne Vector Gravimetry*, In Proceedings of International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation (KIS97), Banff, Canada, June 3-6, pp. 509-520.

Wei, M. and K.P. Schwarz (1997) *Comparison of Different Approaches to Airborne Gravimetry by INS/DGPS*, In Proc. of the Int. Symposium Gravity, Geoid and Marine Geodesy (GraGeoMar96), Tokyo, September 20 – October 5, pp. 155-162.

Wei, M. and K.P. Schwarz (1998) *Flight Test Results from a Strapdown Airborne Gravity System*, Journal of Geodesy, 72, pp. 323-332.

Wei, M. and K. Tennant (2000) *STAR-3i Airborne Gravity and Geoid Mapping System*, In Proceedings of IAG International Symposium on Gravity, Geoid and Geodynamics 2000, July 31-August 4 2000, Banff, Canada.

Wu, L. and M. Sideris (1996) *Accuracy Improvement of Airborne Vector Gravimetry With and Without Noise PSDs*, In Proc. of the Int. Symposium Gravity, Geoid and Marine Geodesy (GraGeoMar96), Tokyo, September 20 – October 5, pp. 147-154.

APPENDICES

There are five appendices that follow.

The goal of Appendix A is to describe the airborne field campaigns that are used for the research presented herein. They are described in the order in which they took place and specific reference is made to the author's contributions to their planning and execution. Reference is also made to other publications that have employed the data from each field campaign.

The technical specifications of the Honeywell Laseref III SINS are given in Appendix B.

In Appendix C, a number of scientific assumptions that are made at different places in the dissertation are described.

The goal of Appendix D is to supplement the analyses of aircraft dynamics that are presented in Chapter 7.

Finally, Appendix E updates previously published results of the SINS data processing for the Greenland campaign. The appendix describes work done by the author to improve the performance of the LRF-III for that campaign.

Appendix A: Description of the Airborne Field Campaigns

The objective of this appendix is to describe the airborne field campaigns that are used in the research presented herein. These SINS/DGPS data sets constitute a portion of the database that has been created at the University of Calgary.

The field campaigns that are described below were preceded by an airborne test that was carried out by the University of Calgary in June 1995. In that initial field test, a single east-west profile over the Canadian Rocky Mountains was flown four times with the goal of assessing the potential accuracy of a SINS/DGPS gravimetry system for geoid determination. The test is the first known use of a SINS for gravimetry. As shown in Wei and Schwarz (1998), it yielded encouraging results in the very rough gravity field of that mountainous area. As a result, it prompted the further tests that are described in the following. Other publications that make use of the data from the 1995 campaign are Bruton and Schwarz (1997) and Jekeli and Kwon (1999).

Figure A-1 shows the location of each of the campaigns that are described in the following sections; namely, the Kananaskis, Greenland and Alexandria campaigns.

Details about the Kananaskis campaign are given in Section A.1. It took place in September 1996 with the goals of assessing the long-term accuracy and repeatability of the SINS system and further evaluating its potential for geoid determination. The author was responsible for data collection at one of the DGPS master stations. The other individuals from the University of Calgary who participated in the data collection were Y. Li, J. Skaloud, G. Chevalier, C. Glennie, G. Kerschbaumer and A. Mohamed. Flight services were provided by Intermap Technologies Corporation. Detailed processing of the data collected during this campaign is described in Glennie and Schwarz (1997) and Glennie and Schwarz (1999). The latter two publications also make up parts of Glennie (1999). Another publication that makes use of the data from the 1996 campaign is Li (2000).

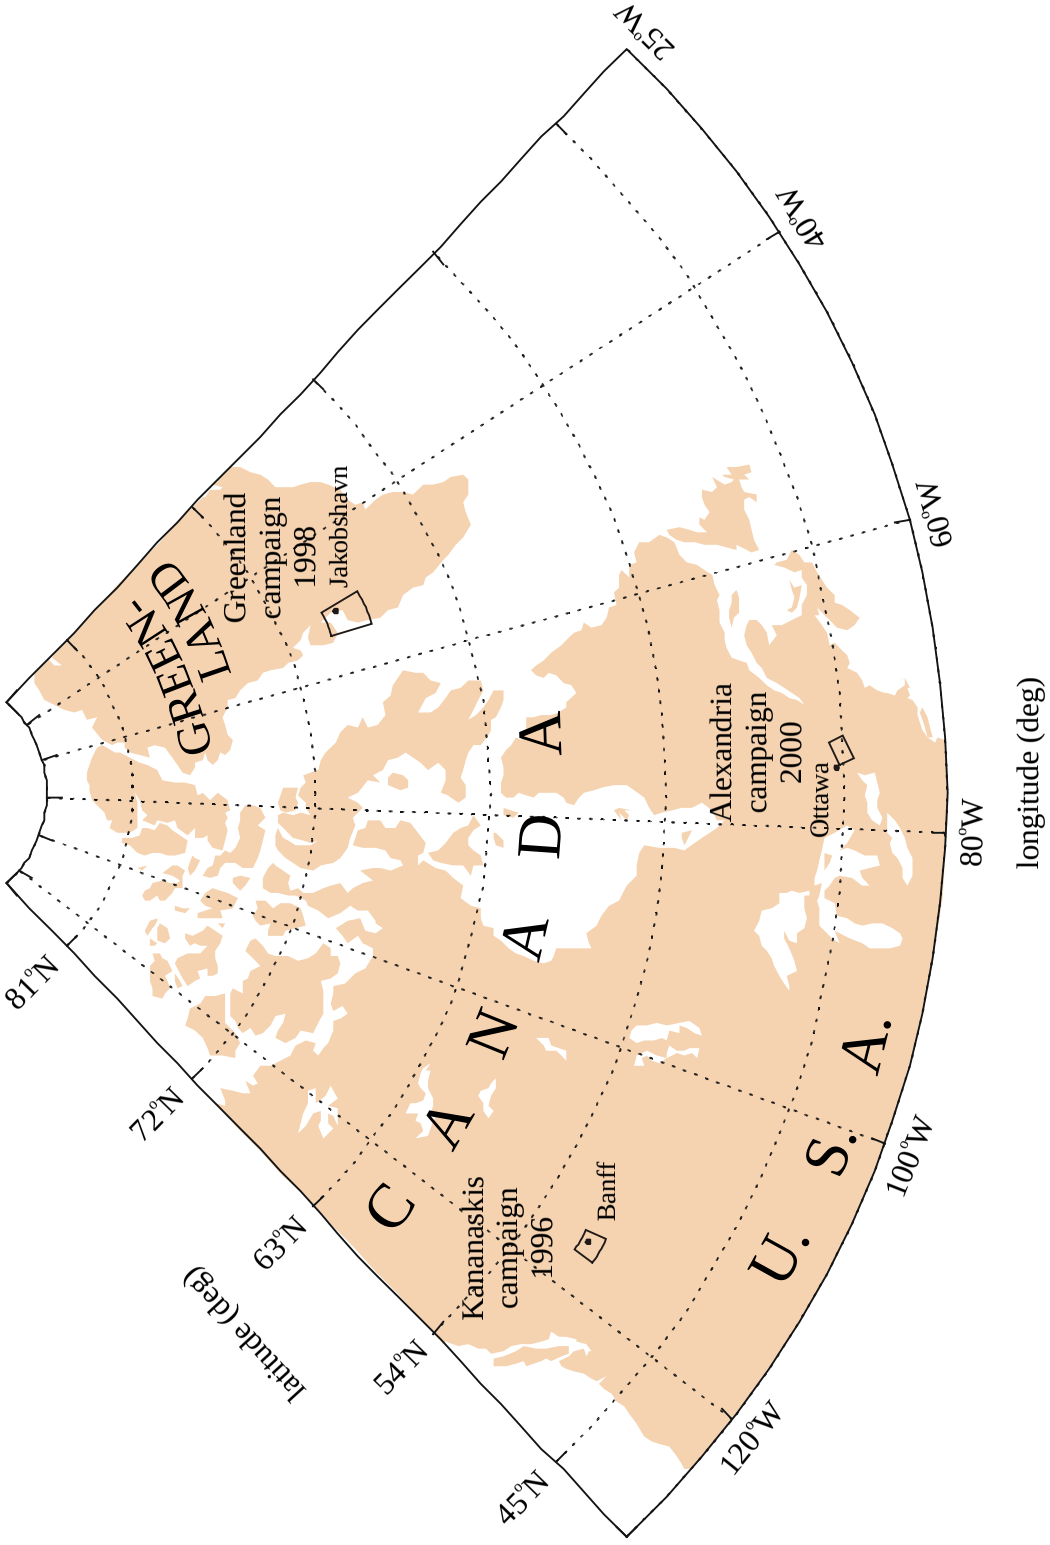

Figure A-1 Airborne gravity campaigns used in this research

In September 1999, the author coordinated an effort to increase the density of the surface gravity data available in the vicinity of the Kananaskis campaign. Using a helicopter, the densification resulted in the addition of over 70 high-elevation points to the database of surface measurements available in that area. The author was responsible for planning the test, hiring the helicopter company and for the positioning component of this densification campaign. The gravity measurements were made and processed by P. Salib from the Geodetic Survey Division of Geomatics Canada (GSD). Some processing was also done by P. Novak of the University of Calgary.

Details about the second airborne campaign are given in Section A.2. It took place after the *Airborne Gravity and the Polar Gravity Field Workshop* in Kangerlussuaq, Greenland between June 2 and 4, 1998. The University of Calgary was invited by R. Forsberg from Kort & Matrikelstyrelsen (KMS) to participate in the workshop and the joint flight that followed it. The goal was to compare the LRF-III SINS gravimeter to their stable platform LaCoste & Romberg gravimeter on the same aircraft. In this test, the author shared the responsibility with C. Glennie for planning, testing and execution with respect to the role of the University of Calgary. Flight planning and services were carried out by KMS. First results from this campaign can be found in Glennie et al. (1999) and an improvement of the LRF-III results can be found herein, in Appendix E.

The third field campaign is described in Section A.3. It was carried out near Ottawa, Canada in April and May 2000, as part of the AGEM project of the GEOIDE NCE. Flight services were carrier out by Sander Geophysics Ltd. The first goal of the campaign was to compare each of the three available airborne gravity system concepts on a single aircraft (see Section A.3 for a list of them). The second goal of the campaign was to assess the performance of the systems for geophysical applications. This was possible because the aircraft used is capable of flying as slow as 45 m/s. The author represented the University of Calgary in this campaign. He was responsible for testing, planning and execution with respect to the role of the University of Calgary. Since the campaign took place, he also has been responsible for coordinating a project that compares the results obtained with each gravity system. Although the comparison is still underway,

preliminary results from the SINS system can be found in Bruton et al. (2000a) and Chapter 5 of this dissertation.

A.1 The Kananaskis campaign

Following the promising results obtained from the initial flight test carried out in June 1995, this second more extensive test took place on September 9, 10 and 11, 1996. It was carried out over the Canadian Rocky Mountains, centred above Banff, Alberta. As in 1995, the test area was chosen for the very high variability of the gravity field and because of the dense surface gravity coverage available in the area. The height of the terrain varies between approximately 800 and 3600 m in the area and the gravity disturbances at flight height vary from about –70 to 100 mGal. For plots of the topography and the anomalous gravity field in the area, the reader is referred to Glennie and Schwarz (1999) and Li (2000).

On each day of testing, data was collected in a 100 by 100 km area with flight lines spaced 10 km apart. The average ellipsoidal flying height was 4357 m on September 9 and 10, and 7300 m on September 11. Data was collected at night (24:00 to 6:00 local time) in an effort to minimize the effects of the atmosphere on the GPS signals and to minimize turbulence. As shown in Figures A-2 and A-3, the bulk of the flight lines were flown in an east-west direction on September 9 and 11 and a north-south direction on September 10. The legend for the maps shown in this appendix is given in Table A-1.

A Cessna Conquest turbo prop airplane was used for the testing and the average flying speed on all three days was 360 km/h (100 m/s). An Ashtech Z-12 and a Trimble 4000 SSI receiver were both used on the aircraft. They shared the same antenna. GPS master stations were located at the Calgary airport (one NovAtel GPS Card), Banff (one Ashtech Z-12 and one Trimble SSI) and Invermere (one Trimble SSI). All GPS data was collected at a sampling frequency of 1 Hz.

Two SINS were onboard: LRF-III and a Litton-101 Flagship. The data from the LRF-III was acquired using a dedicated system that is owned and operated by Intermap. The dedicated system simultaneously obtains data from the GPS and INS using the pulse-per-second (PPS) timing trigger that is provided as output by the GPS. The data from the Litton-101 was acquired using a University of Calgary personal computer that emulates

the behaviour of the dedicated system owned by Intermap. Only data from the LRF-III and the Banff and Invermere master stations is used in this dissertation. Data from September 11 is not used herein.

A reference gravity field was computed by V. Argeseanu by upward continuing the available ground gravity data in the area. For details about this process, see Argeseanu (1995). The RMS accuracy of the upward continued reference after interpolation to the flight lines is about 1.5 mGal.

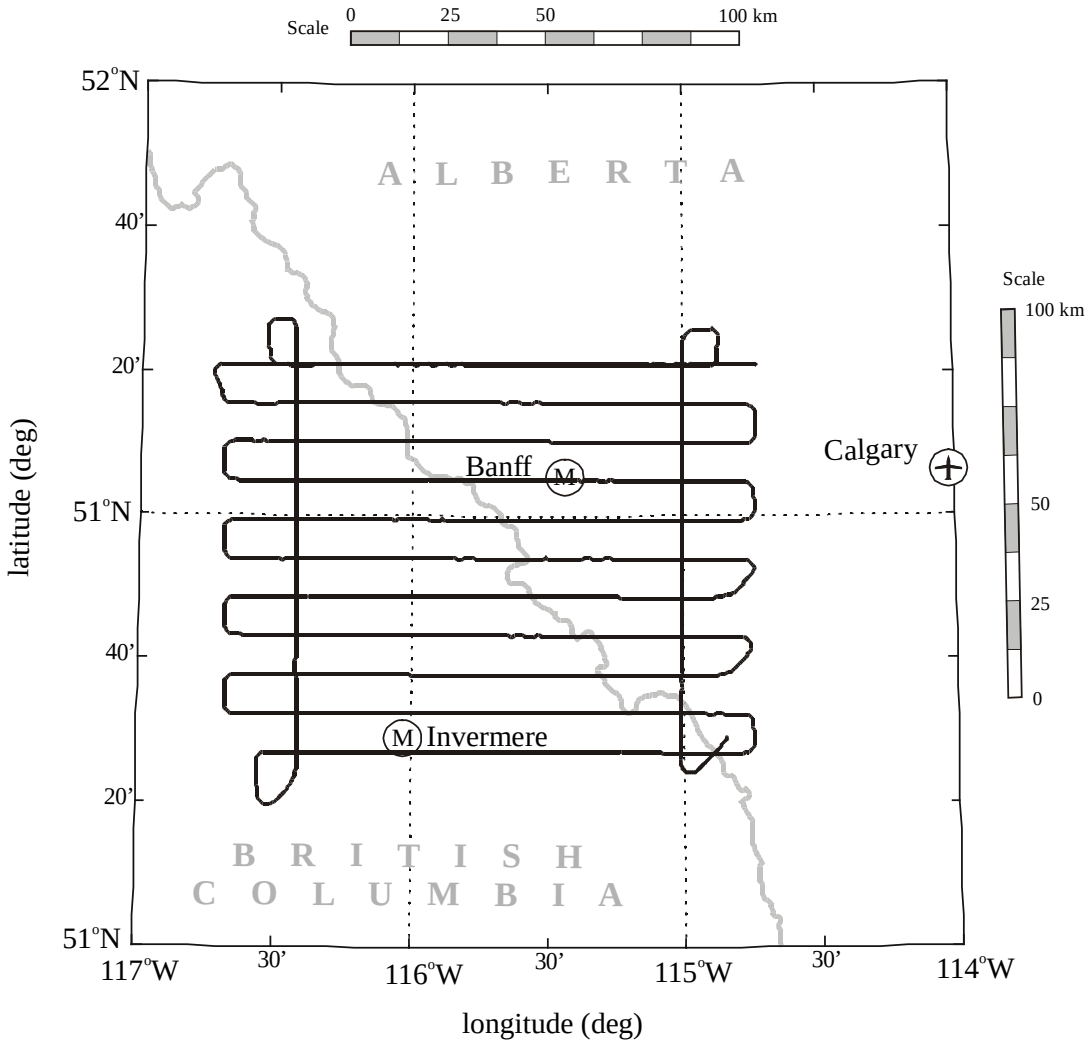

Figure A-2 Flight pattern for September 9 and 11 of the Kananaskis campaign

Table A-1 Legend used in Figure A-2 through Figure A-6

Projection: Lambert	flight lines
neighboring cities or towns	coastlines and rivers
GPS master station	provincial borders
airport and GPS master station	international borders

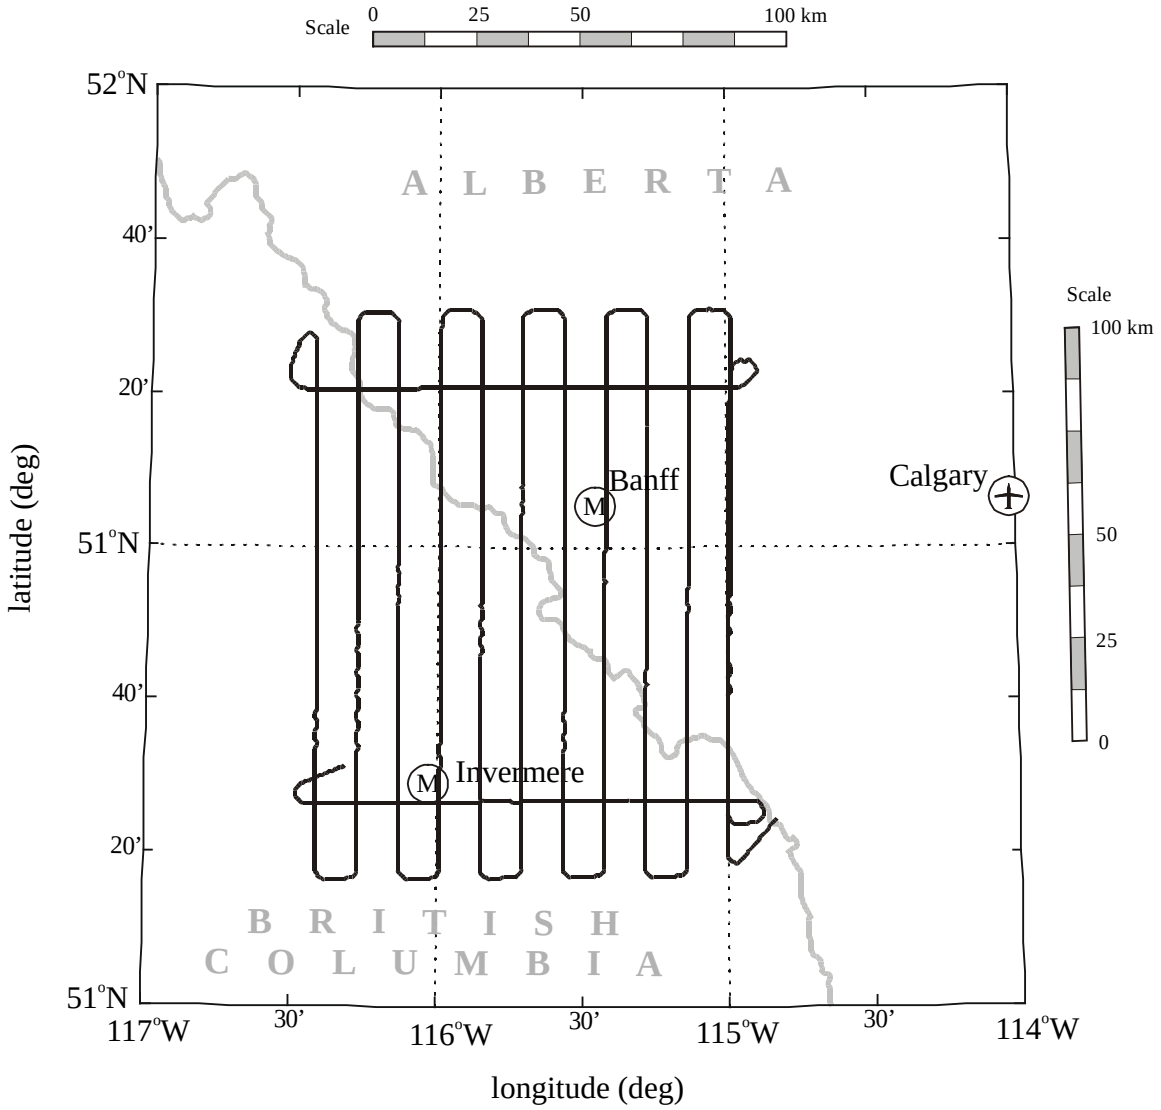

Figure A-3 Flight pattern for September 10 of the Kananaskis campaign

A.2 The Greenland campaign

The University of Calgary and KMS carried out this joint airborne gravity survey off the west coast of Greenland on June 6, 8 and 9, 1998. The LRF-III and the LaCoste and Romberg were mounted together in a Twin Otter airplane so that they could be directly compared. The LRF-III was on loan to the University of Calgary from Intermap. The LaCoste and Romberg meter is owned by the University of Bergen, Norway. Data from the LRF-III was acquired using the same University of Calgary personal computer that was mentioned in Section A-1. Also on board was an orthogonal triad of accelerometers that is owned and operated by the Bavarian Academy of Sciences and Humanities in Munich.

As shown in Figure A-4, the flights in this campaign took place mainly over the ocean. Flight lines A, B and C were flown on June 6. Flight lines F and G1 were flown on June 8. The lines flown on June 9 are not shown because the data acquisition system used for the LRF-III malfunctioned. The aircraft took off from Kangerlussuaq and landed at Jakobshavn on June 6 and the airstrip at Jakobshavn was used for all subsequent flights. Data was collected in the afternoon in all cases (between approximately 15:00 and 18:00 local time) at an average flying speed of 250 km/h (70 m/s) and an average ellipsoidal height of 300 m.

The GPS master stations at Kangerlussuaq and Aasiaat were each equipped with Trimble 4000 SSI receivers and the station at Jakobshavn was equipped with an Ashtech Z-Surveyor receiver. On June 6, GPS data was collected at Kangerlussuaq and Jakobshavn. On June 8 and 9, GPS data was collected at all three master stations. Two dual frequency GPS antennas were mounted on the fuselage of the aircraft, separated by over 3 m. The front antenna was attached to a Trimble 4000 SSI receiver and the signal from the rear antenna was split to a Trimble 4000 SSI and an Ashtech Z-12 receiver. All GPS data was collected at a sampling frequency of 1 Hz.

Flight lines A and G1 were in part flown over existing shipborne gravity profiles, in order to obtain an independent reference. As mentioned in Glennie (1999), the accuracy of the shipborne profiles is 1 mGal.

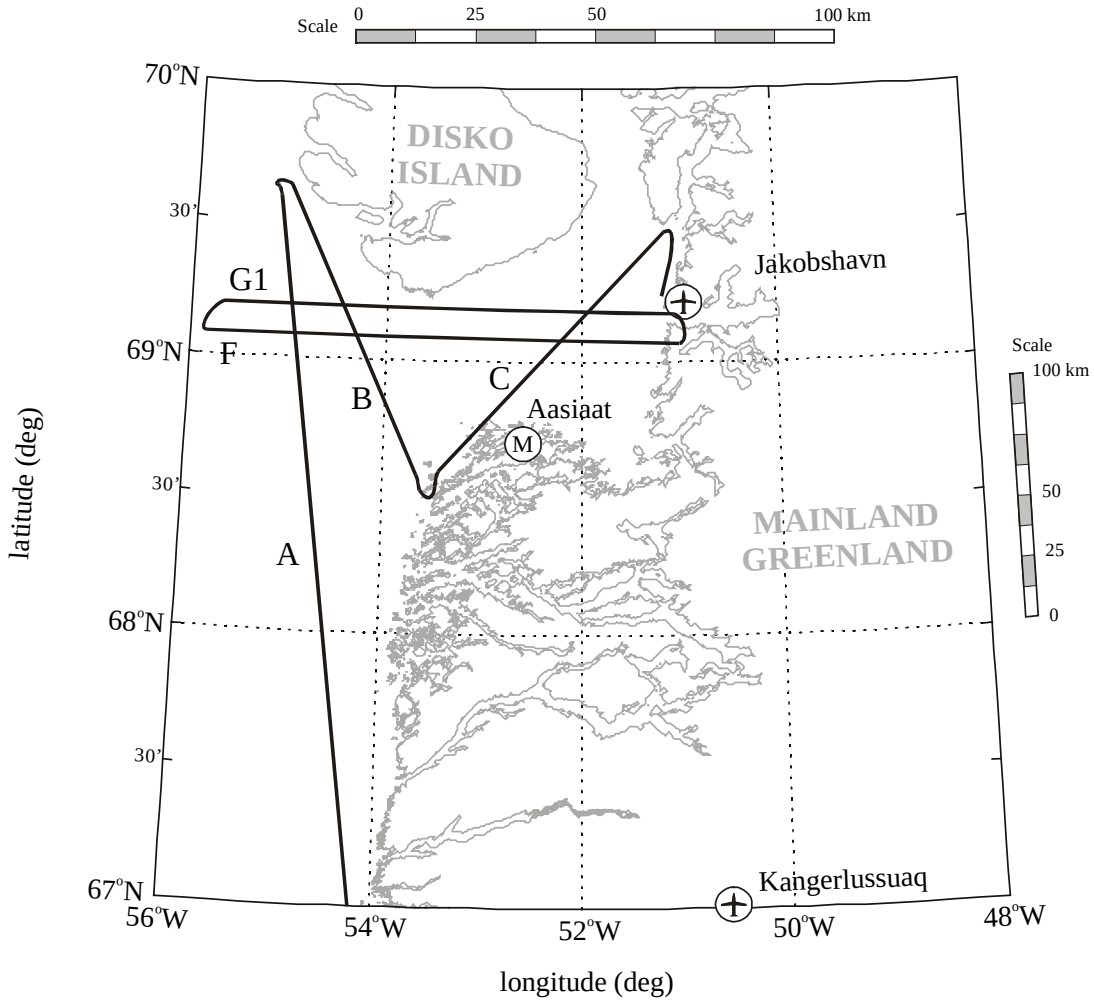

Figure A-4 Flight patterns for the Greenland campaign

A.3 The Alexandria campaign

The Alexandria campaign took place on six days between April 17 and May 4, 2000 over a 70 by 120 km test field conveniently located with respect to the main office of Sander Geophysics Ltd. in Ottawa, Canada. The test field was chosen because the variations in the gravity field are considerable, but the variations in the height of the terrain are small, meaning that the effect of the terrain on the gravity signal is small. Consider Figure A-5 that shows the gravity disturbance at an altitude of 600 m that varies between –41 and +19 mGal. This was calculated and supplied to the author by Sander Geophysics Ltd. using surface data that is available at a spacing of 1-2 km. The height of the terrain varies between approximately 27 and 253 m.

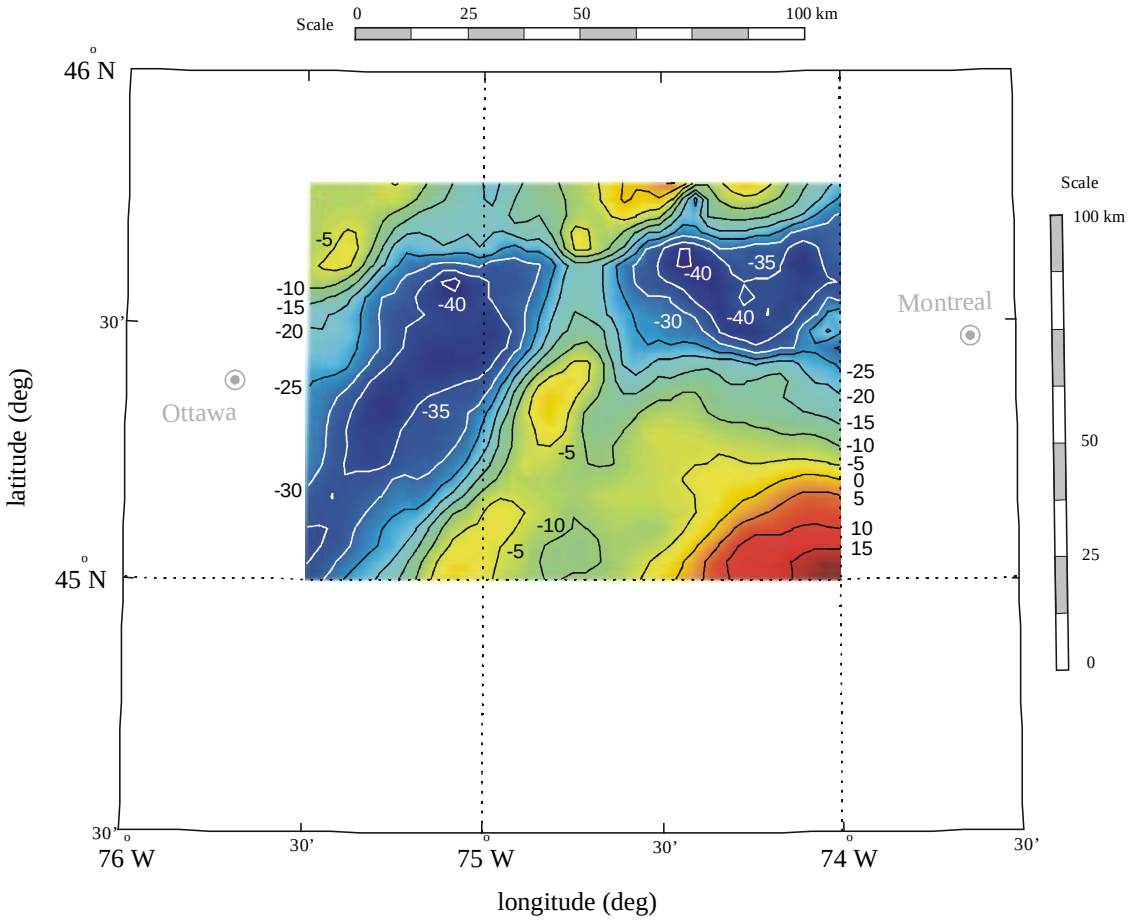

Figure A-5 Variations in the gravity field within the Alexandria range (contours in mGal)

A Cessna 208B Grand Caravan airplane was used to house the three different types of gravity systems. The participants of the test included Intermap Technologies Corporation (Intermap) and the University of Calgary (who jointly operated the LRF-III SINS), Sander Geophysics Ltd. (who operated AIRGRAV, the custom inertially stabilized three-axis platform gravity system that they have developed) and the Geodetic Survey Division of Geomatics Canada (who operated a Lacoste and Romberg gravimeter, model SL-1). Data from the LRF-III was acquired using the University of Calgary personal computer.

Two GPS antennas were used on the aircraft; one on the fuselage directly above the gravimeters and one on the tail. Each of these antennas was each attached to a NovAtel MiLLennium GPS receiver collecting data at a sampling frequency of 10 Hz. The signal from the antenna on the fuselage was also split to an Ashtech Z-12 receiver that was logging data at a sampling frequency of 1 Hz. On the ground, there were three GPS

master stations. Two of these master stations were located near the Ottawa airport as shown in Figures A-6 and A-7 and the third was located in the centre of the test area. The third master station is not shown in the figures because it is not used in this dissertation. Only two surveys that were carried out during the campaign are used herein. The first was flown on April 19, 2000 between 16:00 and 21:00 and the second on May 4, 2000 between 09:30 and 12:30. The May 4 flight was flown in the morning in an effort to minimize the effect of turbulence and the atmosphere. As shown in Figures A-6 and A-7, the flight lines were flown parallel to a line running from south-west to north-east. The spacing between lines was 10 km. The average flying height was 600 m and the average flying speed was just under 160 km/h (45 m/s).

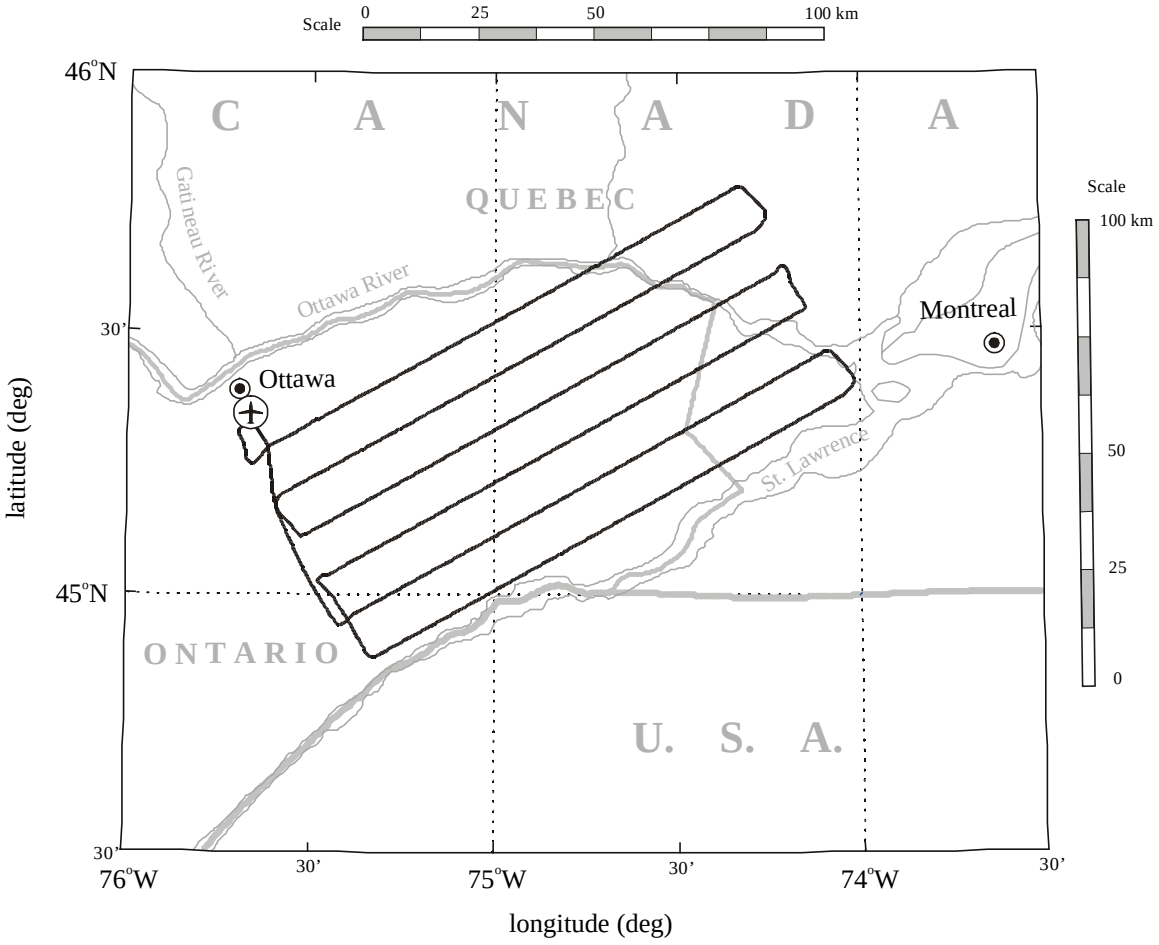

Figure A-6 Flight trajectory for April 19, 2000

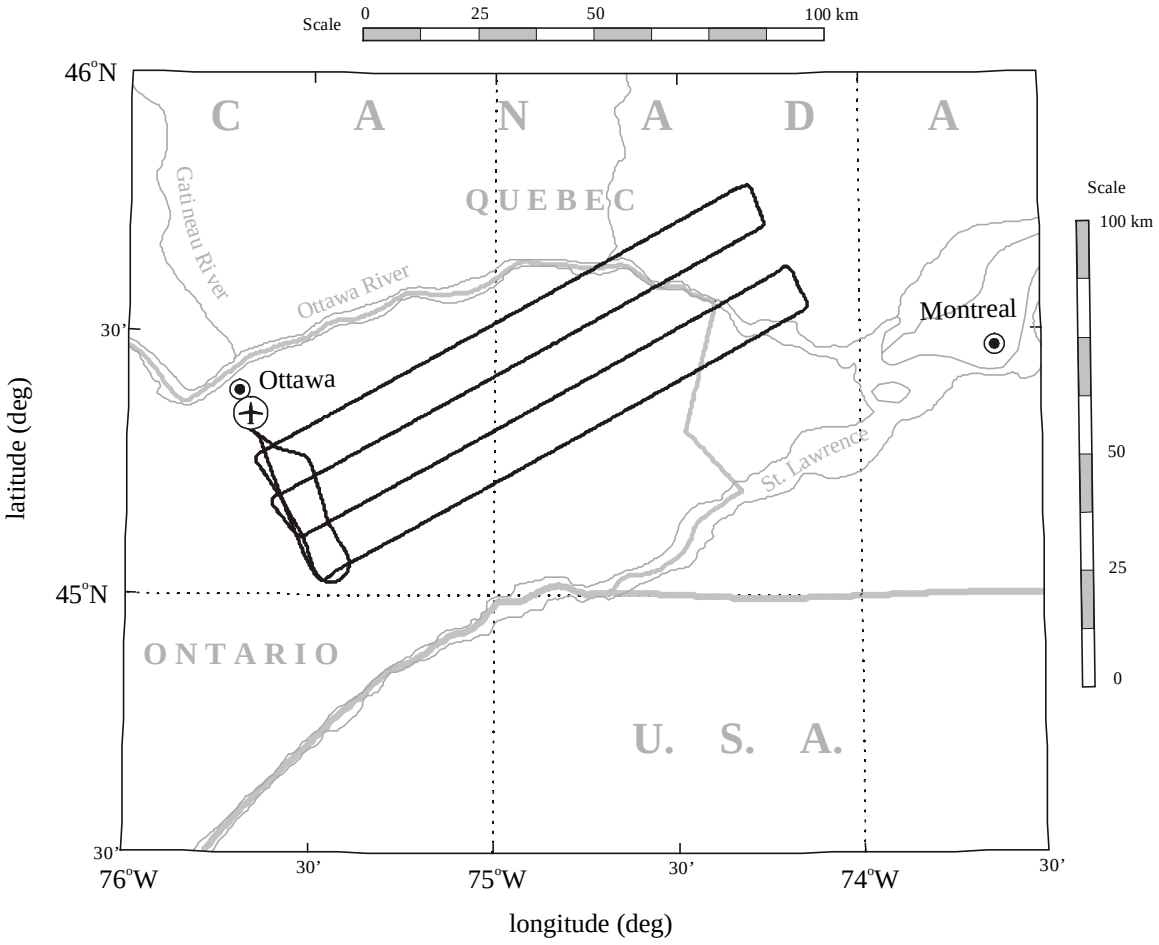

Figure A-7 Flight trajectory for May 4, 2000

Appendix B: Technical Specifications of the Honeywell Laseref III SINS

The Honeywell Laseref III is a navigation-grade strapdown inertial navigation system. It is owned by Intermap Technologies Corporation of Calgary, Alberta. It contains QA-2000 (10) accelerometers and dithered GG1342 ring laser gyroscopes, and has the following technical specifications:

General:	grade:	Navigation
	free-inertial spec.:	1 nm/h
	data rate:	50 Hz
Gyroscopes:	drift:	0.003 deg/h
	scale factor:	1 ppm
	misalignment:	2 arc sec
	random walk:	0.001 deg/$\sqrt{h}$ (1-σ)
Accelerometers:	bias:	10-25 mGal
	scale factor:	25-50 ppm
	misalignment:	5 arc sec
	random walk:	5 mGal (1-σ)

Appendix C: Scientific Background and Assumptions

In this appendix, a number of scientific assumptions that are made at different places in the dissertation are described. Background information is provided in each case.

C.1 On the resolution of the true and measured gravity fields

The gravity field of the Earth has spectral energy at all spatial frequencies. However, representation of the gravity field using discrete, spatial samples results in a band limited estimate of the continuous gravity field. Fortunately, band limited representations of the gravity field are sufficient for many applications because the spectral energy of the anomalous gravity field decreases rapidly with increasing frequency.

The accuracy and resolution requirements for different applications were introduced in Chapter 1 in terms of the high, medium and low-resolution components of the gravity field. These were defined such that they correspond to spatial resolutions below 5 km, between 5 and 150 km and above 150 km, respectively. For the purposes of this research, data represented by a given spatial resolution implies having an image of the field corresponding to a sample distance, x, given in km/sample. For example, a low-resolution image of the gravity field corresponds to sample distances greater than x = 150 km. According to Nyquist's sampling theorem, this means that such an image only represents gravity signals with wavelengths longer than 300 km. It should be pointed out that it is also common in airborne gravimetry to characterize spatial resolution in terms of either

a) the full-wavelength, which is 300 km in this example, or
b) the half-wavelength, which is 150 km in this example.

In this dissertation, the terms *spatial resolution* and *half-wavelength* are used synonymously.

In airborne applications, discrete samples of the gravity field are collected as a function of time. The spatial data is acquired at a specific uniform temporal sampling frequency, f_s, in samples per second (i.e. Hz). The corresponding spatial resolution, x, of the raw

sampled data is therefore approximated by $x = s / f_s$, where s is the nearly constant aircraft speed. Typically, the effective value of f_s for GPS receivers is 1 Hz and the flying speed varies between 180 km/h (50 m/s) and 720 km/h (200 m/s). This means that the raw spatial resolution varies between 0.05 and 0.20 km. If the temporal data were free of errors, the spatial resolution of the gravity data along the flight path would therefore be as high as 50 m. However, this is not the case because of the noise characteristics of the sensors and the inaccuracies of the kinematic models that are used, see Chapter 2. Because the typical way of handling these errors is to apply a low-pass filter to the estimates, the upper limit of their spatial resolution is given by $x_{\text{limit}} = s/(2 f_c)$, where f_c is the cut-off frequency of the low-pass filter. Typical low-pass filters used in current airborne gravity systems have cut-off frequencies between 0.0050 and 0.0166 Hz. They are often referred to by their corresponding periods, $T_c = 1/ f_c$, of between 200 and 60 s, respectively. For example, at a speed of 180 km/h, these correspond to effective maximum spatial resolutions of 5 and 1.5 km, respectively. At a speed of 720 km/h they correspond to 20 and 6 km, respectively.

Section 6.1 in Chapter 6 includes a detailed discussion of the sample spacing of surface and airborne data and a comparison of the spatial resolution that is implied in each case.

C.2 Relating spatial resolution to harmonic degree

Global models of the gravity potential are typically expressed as an expansion in spherical harmonic functions, see Heiskanen and Moritz (1969). The maximum degree, *n*, of that expansion implies a minimum spatial resolution $x_{\min} \approx (40000km)/(2n)$. For example, a spherical harmonic expansion to degree 360 implies a sample spacing of approximately 55 km. This is discussed in more detail in Sideris (1993) and NRC (1997), for example.

C.3 Attenuation of gravity with altitude

According to Newton's inverse-square law, the strength of the gravitational attraction decays as a function of distance from the attracting mass. This decay is also a function of resolution; for a given distance from the source, high-resolution signatures are attenuated

less than low-resolution signatures. This was demonstrated in Section 6.1.2 where a plot of the attenuation factor, *A*, was generated for several examples using the following equation:

$$A = \left(\frac{a}{a+h} \right)^{(n+2)}, \qquad \text{(C.1)}$$

where *h* is the height above the surface *a*. The spatial resolution is obtained from the degree of the expansion, *n*, according to the equation given in the last section. Equation C.1 is quoted from NRC (1997) where it is derived from the spherical harmonic expansion for the anomalous gravity field.

Appendix D: More About Aircraft Dynamics

D.1 Solutions with complex eigenvalues

It was stated in Chapter 7 that the solutions to the characteristic equations for natural motion are always real, even when the eigenvalues are complex. The following simple proof of this fact is based on a similar proof presented in Etkin (1982). Because the complex eigenvalues always appear in conjugate pairs, $\lambda = a \pm \omega j$, the solution (i.e. equation 7.11) will be given by:

$$\mathbf{x} = e^{(a+\omega j)t}\mathbf{v}_1 + e^{(a-\omega j)t}\mathbf{v}_2 . \tag{D1.1}$$

where $\mathbf{v}_1$ and $\mathbf{v}_2$ are the eigenvectors corresponding to the complex eigenvalues. Because the real and imaginary parts of a complex exponential sequence vary sinusoidally with *t*, this can also be written as follows:

$$\mathbf{x} = e^{at}(\mathbf{c}_1 \cos \omega t + \mathbf{c}_2 \sin \omega t), \tag{D1.2}$$

where the coefficients $\mathbf{c}_1 = \mathbf{v}_1 + \mathbf{v}_2$ and $\mathbf{c}_2 = j(\mathbf{v}_1 - \mathbf{v}_2)$ are both real. The latter is real because the elements of $\mathbf{v}_1$ and $\mathbf{v}_2$ are themselves complex conjugate numbers. In other words, the solution corresponding to a pair of complex eigenvalues is a real-valued, exponentially weighted sinusoid, as shown by the dotted lines in Figure 7-3.

D.2 Calculation of the natural modes for sample aircraft

Chapter 7 includes examples where the natural motion of the Navion and Boeing 747 aircraft are demonstrated. The characteristic equations for each of these are borrowed for from Bryson (1994) who calculates them from dynamic control data available for those aircraft. They are given as follows for the Navion (for the state vector that has units of ft/s, centirads/s and centirads):

$$\mathbf{F}_{\text{lon}} = \begin{pmatrix} -0.045 & 0.036 & 0.000 & -0.322 \\ -0.370 & -2.020 & 1.760 & 0.000 \\ 0.191 & -3.960 & -2.980 & 0.000 \\ 0.000 & 0.000 & 1.000 & 0.000 \end{pmatrix}, \tag{D2.1}$$

$$\mathbf{F}_{\text{lat}} = \begin{pmatrix} -0.254 & -1.760 & 0.000 & 0.322 \\ 2.550 & -0.760 & -0.350 & 0.000 \\ -9.080 & 2.190 & -8.400 & 0.000 \\ 0.000 & 0.000 & 1.000 & 0.000 \end{pmatrix}, \tag{D2.2}$$

and as follows for the Boeing 747:

$$\mathbf{F}_{\text{lon}} = \begin{pmatrix} -0.003 & 0.039 & 0.000 & -0.322 \\ -0.065 & -0.319 & 7.740 & 0.000 \\ 0.020 & -0.101 & -0.429 & 0.000 \\ 0.000 & 0.000 & 1.000 & 0.000 \end{pmatrix}, \tag{D2.3}$$

$$\mathbf{F}_{\text{lat}} = \begin{pmatrix} -0.056 & -0.774 & 0.000 & 0.322 \\ 0.733 & -0.115 & -0.318 & 0.000 \\ -0.394 & 0.039 & -0.465 & 0.000 \\ 0.000 & 0.000 & 0.000 & 0.000 \end{pmatrix}. \tag{D2.4}$$

The temporal histories of the state variables presented in Figures 7-4 through 7-7 were obtained by solving equations 7.8 and 7.9 as outlined in Section 7.1.3. The Matlab function ODE45 was used for this purpose. The eigenvalues and their corresponding eigenvectors were determined using the Matlab function EIG and the latter were normalized using the largest element of each as recommended in Bryson (1994). They are given in Tables D.1 to D.4. When the eigenvectors are complex, their real parts are used as initial conditions.

Table D.1 Eigensystem for the natural longitudinal modes of the Navion

	Eigenvalues	
	-2.5054 ± 2.5949 *i*	-0.0171 ± 0.2134 *i*
modes	Short period	Phugoid
normalized eigenvectors	-0.0029 ∓ 0.0195 *i*	1.0000 ± 0.0000 *i*
	-0.1200 ∓ 0.6562 *i*	-0.0592 ∓ 0.0013 *i*
	1.0000 ± 0.0000 *i*	0.1430 ∓ 0.0086 *i*
	-0.1926 ∓ 0.1994 *i*	-0.0934 ∓ 0.6627 *i*

Table D.2 Eigensystem for the natural lateral modes of the Navion

	Eigenvalues		
	-8.4327	-0.4862 ± -2.5949 *i*	-0.0088
modes	Pure roll	Dutch roll	Pure spiral
normalized eigenvectors	0.0135	-0.0947 ± 0.0195 *i*	0.0506
	0.0411	1.0000 ± 0.0000 *i*	0.1759
	1.0000	-0.0924 ∓ 0.8835 *i*	-0.0088
	-0.1186	-0.3550 ± 0.1136 *i*	1.0000

Table D.3 Eigensystem for the natural longitudinal modes of the 747

	Eigenvalues	
	-0.3750 ± 0.8818 *i*	-0.0005 ± 0.0674 *i*
modes	Short period	Phugoid
normalized eigenvectors	0.0111 ∓ 0.0079 *i*	1.0000 ± 0.0000 *i*
	1.0000 ± 0.0000 *i*	0.1382 ∓ 0.0124 *i*
	-0.0071 ± 0.1139 *i*	0.0142 ± 0.0007 *i*
	0.1123 ∓ 0.0396 *i*	0.0088 ∓ 0.2107 *i*

Table D.4 Eigensystem for the natural lateral modes of the 747

	Eigenvalues		
	-0.6837	-0.4862 ± -2.5949 *i*	-0.0098
modes	Pure roll	Dutch roll	Pure spiral
normalized eigenvectors	-0.3678	1.0000 ± 0.0000 *i*	0.0522
	0.1177	0.0449 ∓ 0.9845 *i*	0.4129
	-0.6837	-0.2328 ± 0.3248 *i*	-0.0098
	1.0000	0.3707 ± 0.2852 *i*	1.0000

D.3 Observed specific force, attitude and speed

For completeness, this section of Appendix D contains a number of plots that are supplemental to the information presented in Chapter 7. Figures D3-1 and D3-2 show the observed specific force for the sample flight lines of the Kananaskis and Alexandria campaigns, respectively, i.e. they have not been band-limited. Figure D3-3 shows the specific force during the Alexandria campaign, after filtering to 0.01 Hz. Figures D3-4 to D3-7 show the attitude and speed of the aircraft during the same flight lines.

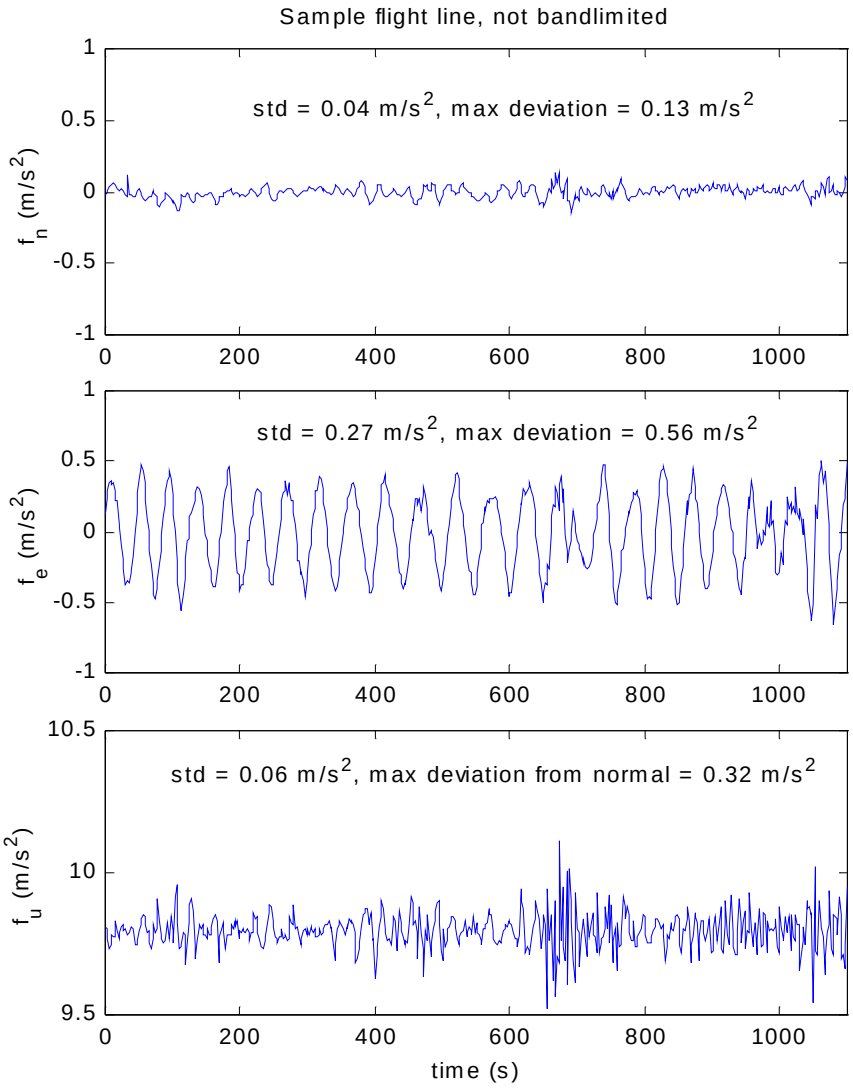

Figure D3-1 Unfiltered specific force for the sample flight line of the Kananaskis campaign

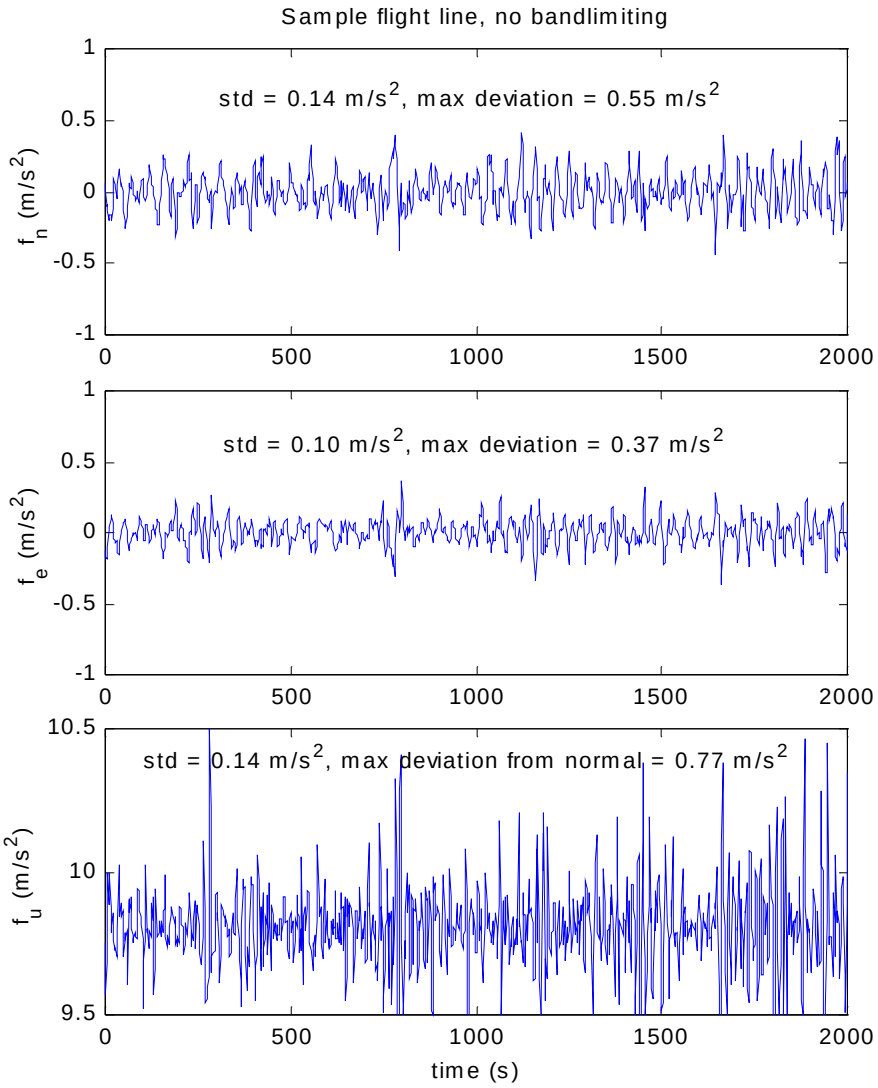

Figure D3-2 Unfiltered specific force for the sample flight line of the Alexandria campaign

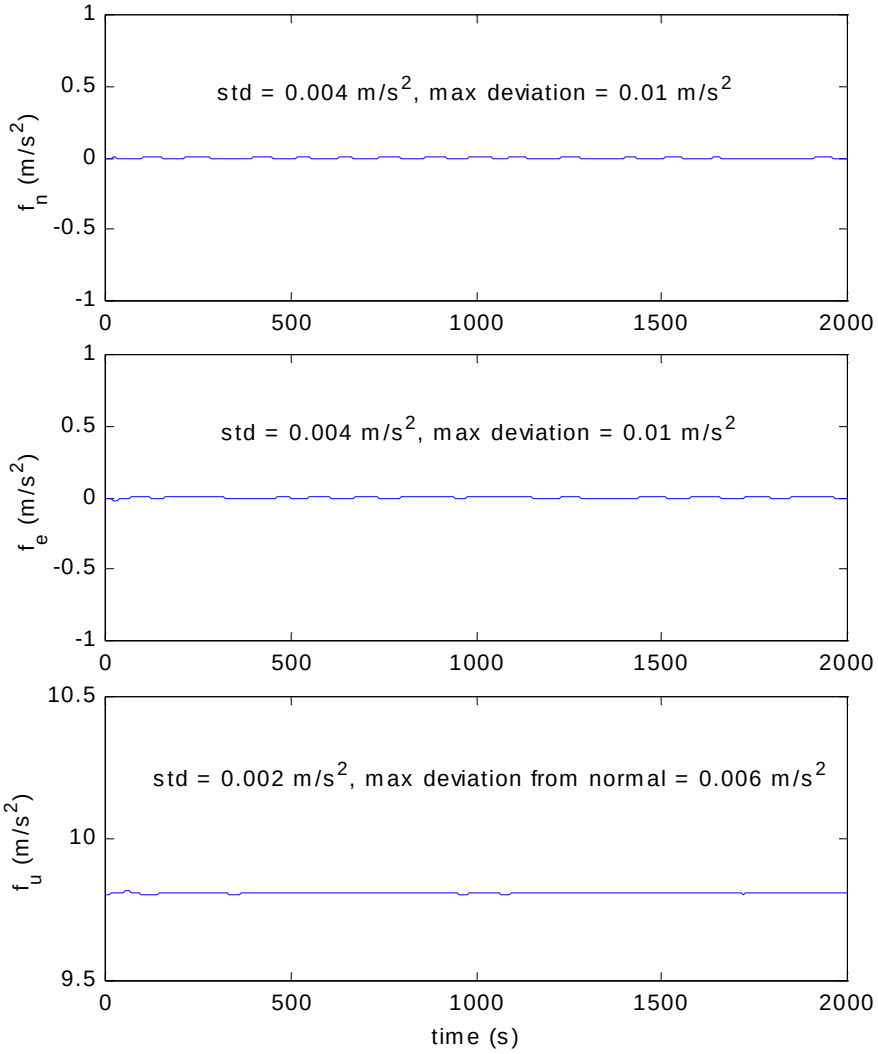

Figure D3-3 Filtered specific force for a sample flight line of the Alexandria campaign (0.01 Hz)

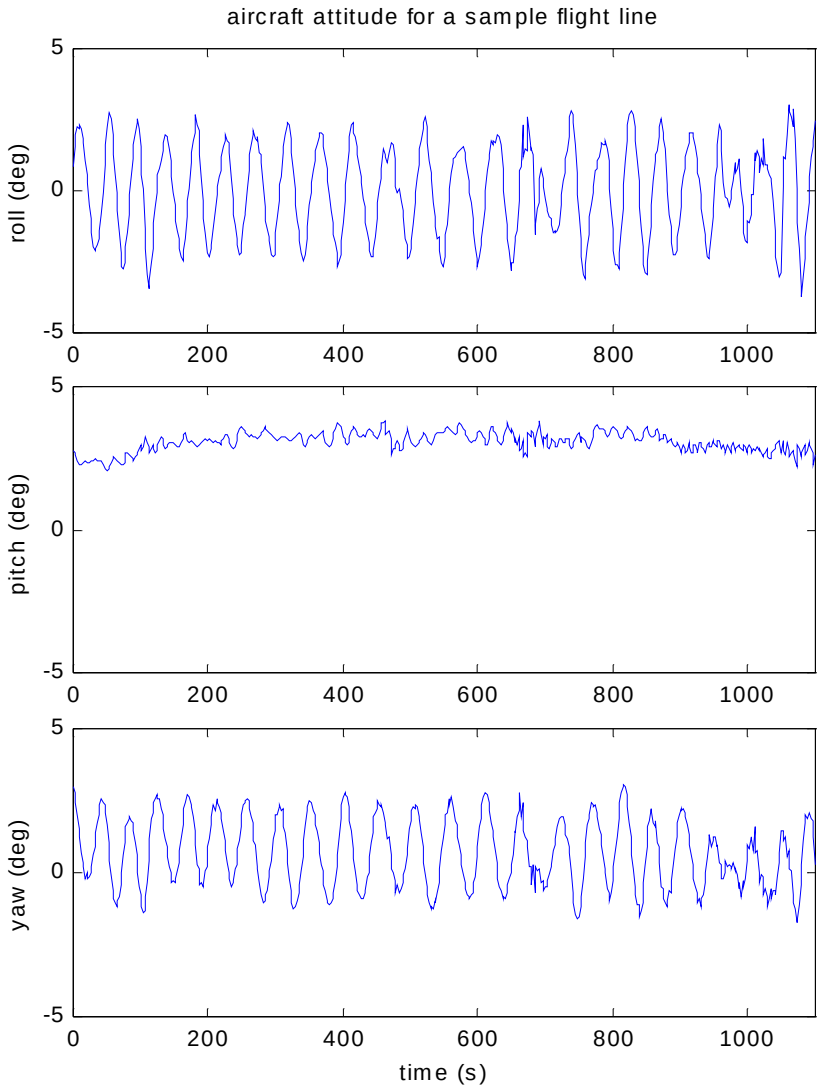

Figure D3-4 Attitude for the sample flight line of the Kananaskis campaign

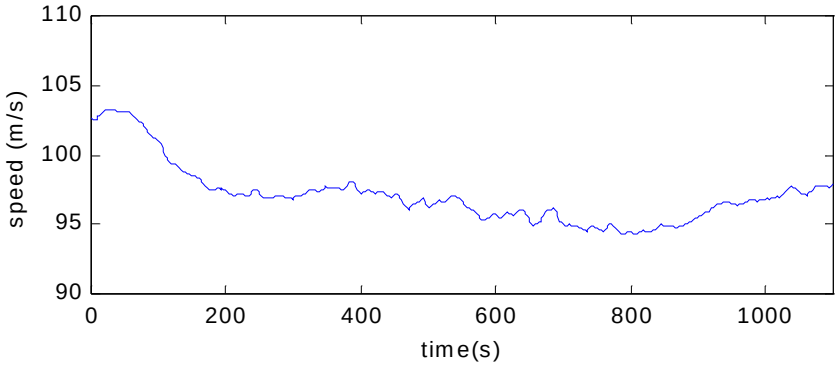

Figure D3-5 Speed for the sample flight line of the Kananaskis campaign

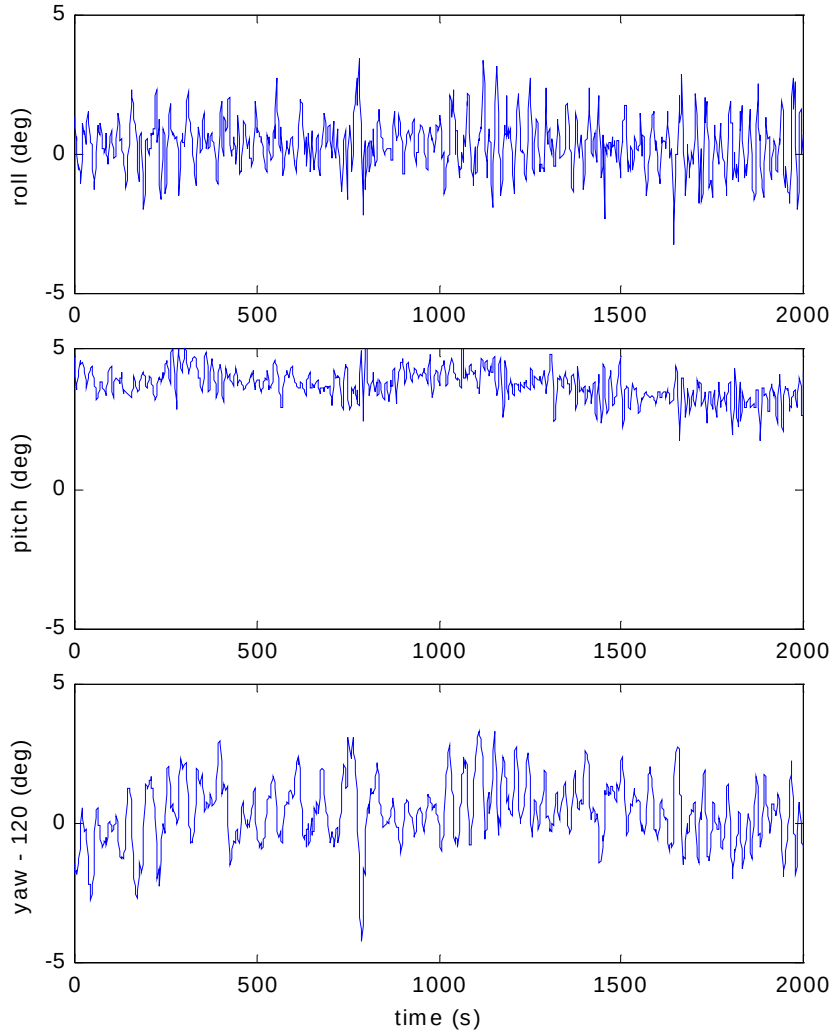

Figure D3-6 Attitude for the sample flight line for the Ottawa campaign

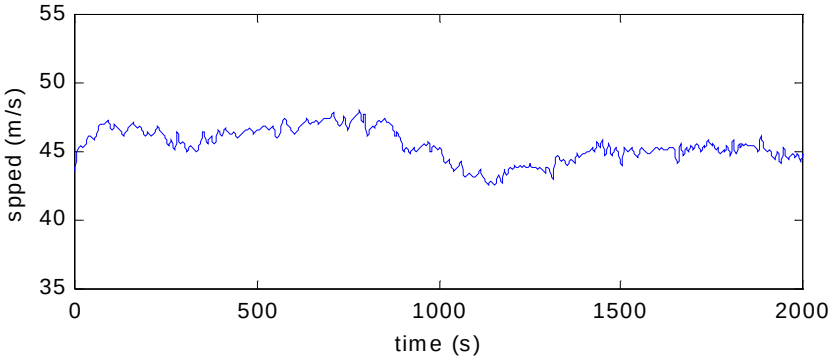

Figure D3-7 Speed for the sample flight line for the Alexandria campaign

Appendix E: Improved Results for the Greenland Campaign

This appendix updates previously published results of the SINS data processing for the Greenland campaign. As in Glennie (1999), the goal of this appendix is to process data using the GREATGUN package and to analyze the results.

See Appendix A.2 for a description of the Greenland campaign and recall that its objective was to compare results obtained with the LaCoste and Romberg to those obtained with the LRF-III SINS for the same flights. First results of the processing for the Greenland campaign can be found in Glennie et al. (1999). The performance of the SINS was not as good as expected in the initial analysis, due mainly to unexplained irregularities in the data, as mentioned in Glennie (1999).

This appendix describes work done by the author to explain the irregularities mentioned above and improve the performance of the LRF-III for that campaign. The irregularities are removed and the results are significantly improved. The performance of the LRF-III is demonstrated to be slightly better than that of the LaCoste and Romberg for this data set. In order to remain consistent with the data provided by KMS for the LaCoste and Romberg, all processing instances described below make use of a filter that has a filtering period of 200 s. All solutions also make use of the same DGPS solution that KMS used (a dual frequency ionospheric-free solution obtained using Trimble's GPSurvey). As in Glennie et al. (1999), a line of best-fit was removed from the LRF-III and LaCoste and Romberg data prior to any comparisons, meaning that only relative gravity information is being evaluated.

E.1 Analysis and repair of the irregular behaviour of the SINS

After considerable analysis, the author discovered that the observed irregularities were caused by the data acquisition system that was used for the LRF-III in the campaign, and not due to errors in either the SINS or GPS sensors. There are several problems with the

synchronization of the raw SINS data and the raw GPS data in this campaign. The errors that result can be classified as problems with the time **intervals** and the time **offset**.

The errors in time intervals are characterized by many jumps in the recorded time interval. Most of the jumps are large, often reaching values of +/- 0.8 s and often occur in complementary pairs (i.e. a forward time jump is usually followed by a backward time jump of roughly equal magnitude). In these cases, they are most likely due to the incorrect labeling of the data for epochs within the time period between jumps and fixing is relatively easy. Unfortunately, however, there are also several cases in which single, unpaired jumps occur. Consider the top portion of Figure E-1 where an example of this is shown for the data collected on June 8 during line F. Because the sampling frequency is 50 Hz, the time interval should be 0.02 s for this data. Clearly, there is a significant jump in the interval at approximately 1200 s. As shown in the bottom portion of Figure E-1, this results in a large error in the gravity field centered about the same epoch. This type of error occurs on a number of occasions on all three days of testing. They correspond to similar errors in the gravity field in all cases. By visual inspection, the time jumps were understood well enough that a general purpose algorithm could be devised. A simple program was written to do this for all of the data sets. The result for line F is also shown in the bottom portion of Figure E-1.

E.2 New results for the SINS

This section summarizes the results of the Greenland campaign after correction for the irregularities discussed above. Table E.1 compares the results obtained using the LRF-III with the shipborne gravity data. The performance is clearly significantly better after the timing problems have been repaired, reaching the level of 1.3 and 2.0 mGal for the flight lines A and G1. Also shown in Table E.1 are the same values for the LaCoste and Romberg gravimeter. After correction, the performance of the LRF-III is slightly better than the LaCoste and Romberg. Figures E-2 and E-3 show the solutions obtained for the same flight lines as well as the reference values obtained from the shipborne data (after shifting by 20 mGal for clarity). For flight line A, note that the shipborne data is only available for a portion of the line and that the data from the LaCoste and Romberg is not

available for the first five minutes due to a problem nulling it. For flight line G1, note that about half of the LaCoste and Romberg data was discarded because of poor quality.

The results obtained with the LRF-III are compared to those obtained with the LaCoste and Romberg in Table E.2. The level of agreement is considerably improved for three of the flight lines, as a result of repairing the timing problems in the acquisition of the LRF-III data

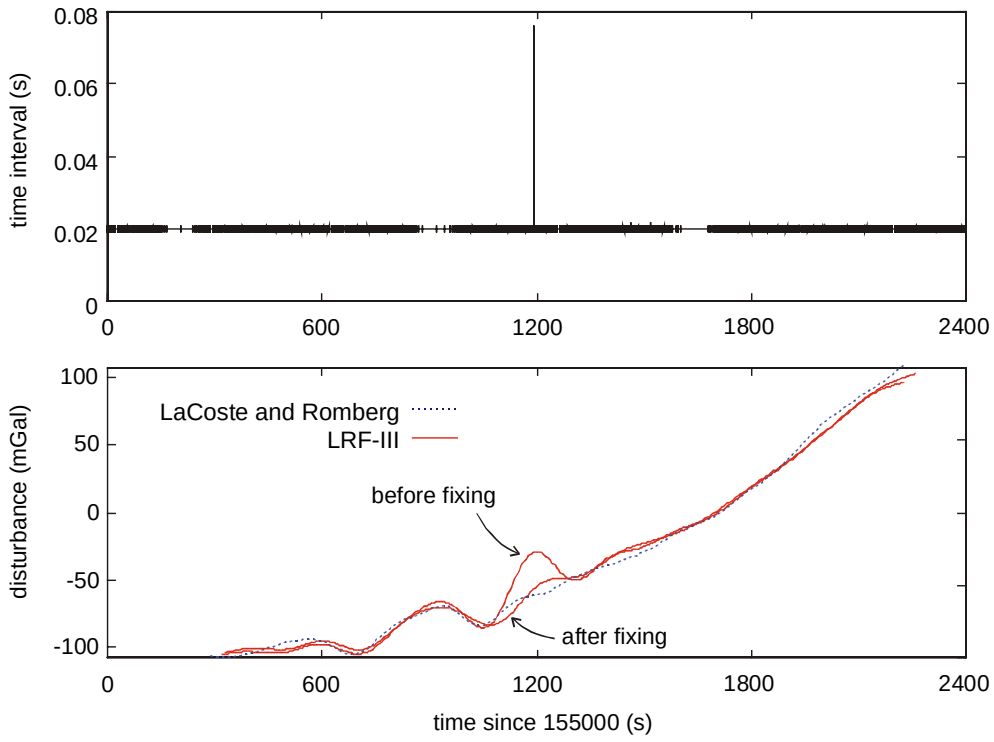

Figure E-1 A timing error and the associated error in gravity (for line F)

Table E.1 Agreement of the LaCoste and Romberg (LCR) and LRF-III with the shipborne data

	Standard deviation of the agreement		
flight line	LRF-III before repair	LRF-III after repair	LCR
A	**2.0**	**1.3**	1.7
G1	**3.8**	**2.0**	2.3

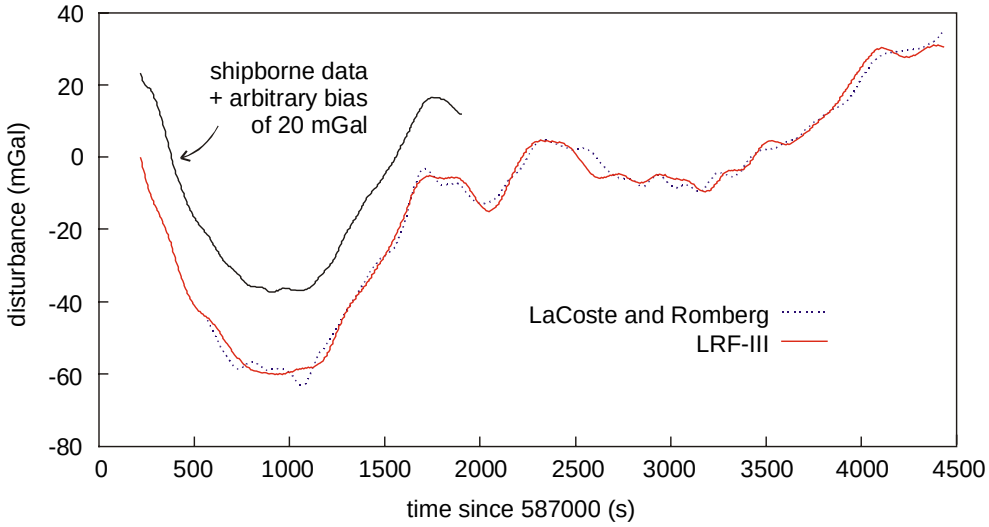

Figure E-2 Estimates of the gravity disturbance for line A on June 6

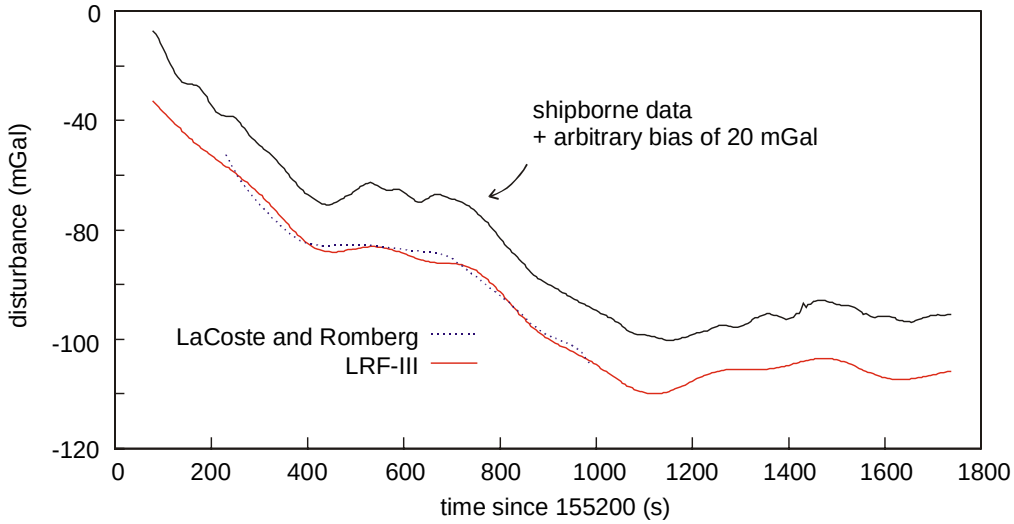

Figure E-2 Estimates of the gravity disturbance for line G1 on June 8

Table E.2 Agreement of the LCR and LRF-III with each other

	Standard deviation of the agreement	
flight line	before	after repair of time jumps
A	2.4	**2.0**
B	3.0	**3.0**
C	1.4	**1.4**
G1	4.0	**2.1**
F	7.7	**4.3**