UNIVERSITY OF CALGARY

Development of a GNSS-Based Multi-Sensor Vehicle Navigation System

By

James E. Stephen

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE

DEGREE OF MASTER OF SCIENCE

DEPARTMENT OF GEOMATICS ENGINEERING

CALGARY, ALBERTA

AUGUST, 2000

ABSTRACT

A vehicle navigation system is developed that uses a combination of GPS positioning and dead reckoning sensors. A survey of dead reckoning sensors is presented, including examples and comparisons of different technologies. The dead reckoning sensors chosen for the final implementation were a low cost piezoelectric vibrating gyro and differential odometry provided by a vehicle's anti-lock braking system. Filtering of the data is done by a centralized Kalman filter with independent calibration filters for the dead reckoning sensors. Augmentations of GPS positioning are also investigated, focusing on clock coasting using an OCXO and height aiding. Several series of tests were performed to analyze the system's performance, culminating with tests in downtown Calgary under urban canyon masking and multipath conditions. Results from the testing indicate that an accuracy on the order of 10 to 20 metres can be achieved under most circumstances, but more intelligent algorithms or map matching are needed to control the large deviations which can be caused by poor GPS solutions.

ACKNOWLEDGEMENTS

I would like to express my gratitude to Gerard Lachapelle for his continued support throughout my M.Sc. program. He provided advice, opportunities and assistance that greatly enhanced my studies and experience.

I would like to thank several of my colleagues, especially those who shared my office and thereby my problems. These include Sam Ryan, Mark Petovello, Rakesh Nayak, and Georgia Fotopoulos. Jayanta Ray also spent several hours talking me through concepts that were once foreign to me.

And last, but certainly not least, I would like to thank my friends and family. My beloved mother, Patricia, has put up with my presence many years more than she was obligated, and supports me in all that I do. I would not be completing this thesis without the many things she has done for me. My friends have listened to me talk about my research when I'm sure they would have liked to talk about absolutely anything else. I love you guys. My girlfriend, Rachel, has supported me and put up with being second to the creation of this thesis on several occasions, so I will have to make it up in the future.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

NOTATION

x	vector of unknowns
P	covariance of the unknowns
Q	dynamics noise matrix
K	Kalman gain matrix
R	covariance matrix of the observations
z	vector of observations
H	measurement design matrix
Φ	state transition matrix
q	spectral density values
σ	standard deviation
φ	latitude
λ	longitude
τ_m	correlation time for First-Order Gauss-Markov processes
Δt	prediction interval

ACRONYMS

ABS	anti-lock braking system
A/D	analog to digital
AS	anti-spoofing
C/A	coarse acquisition
DGPS	differential GPS
DLL	delay lock loop
DOP	dilution of precision
DR	dead reckoning
EM	electro-magnetic
FOG	fiber optic gyro
GNSS	global navigation satellite system
GLONASS	Global Navigation Satellite System (Russian)
GPS	Global Positioning System (U.S)
HDOP	horizontal DOP
IC	Integrated Circuit
INS	Inertial Navigation System
LHCP	left-hand circularly polarized
LPF	low pass filter
LS	Least Squares
MDB	marginally detectable blunder
MEMS	micro-eletomechanical system
MHE	Maximum Horizontal Error
OCXO	ovenized quartz crystal oscillator
PE	piezoelectric
PLL	phase locked loop
RHCP	right-hand circularly polarized
RLG	ring laser gyro
RMS	root mean square

SA	selective availability
WAAS	Wide Area Augmentation System
WADGPS	wide area DGPS

1 INTRODUCTION

Vehicle positioning and navigation has become a major focus of automobile manufacturers over the past decade, adopting technology developed by fleet management and machine control technologies. These systems have been available to the consumer public in Japan and Europe for several years, but have just recently received attention in North America. Driving forces behind this innovation are the desire for route finding information, safety concerns, and commercial opportunity. Several of the concerns that are being addressed by these systems are

How does one get to a destination in a minimum of time and how long will it take?

An owner has locked his keys in his car or can not find it in the parking lot.

Sensors indicate that a vehicle has been in an accident or has a flat tire.

Where are the nearest gas station and automated banking machine?

Another aspect that has fuelled these developments in recent years is e-commerce and the desire for uninterrupted connectivity to the Internet and its many resources. The technology that is making the development of multi-featured navigation and safety systems possible has been termed telematics. Telematics refers to the combination of positioning, wireless, and various other in-vehicle technologies to create a mobile information platform. The full-featured telematics-enabled car of the future could allow the menu of a restaurant you are approaching to be displayed, an order to be placed, and the order could be ready by the time the vehicle arrived. Warnings of approaching severe weather or accidents along the scheduled route could be relayed to the driver. E-mail could be retrieved, faxes sent, reservations sent to restaurants, etc. All forms of communication could be as accessible from the vehicle as from the home or office. This communication would have to be of minimal distraction to the driver, so voice recognition and HUD (Heads-Up Display) systems are being developed to this end.

Several factors have made the integration of these systems into the mass-produced vehicle market feasible. These include a decrease in the cost of electronics, including

processors, dead reckoning sensors, wireless devices and GPS receivers. In addition, the availability of digital information such as maps, locations of services, telephone directories, etc. is constantly increasing.

1.1 VEHICLE NAVIGATION

Vehicle navigation and location-aware systems are being developed or have already been implemented by most automobile manufacturers. Systems in North America have been developed mainly for improved access to emergency services. This includes systems such as Ford's RESCU system with console buttons for roadside assistance and emergency services, as well as automatic emergency assistance on air-bag deployment.

Upgrading of these systems to allow the vehicle operator to contact a central switchboard operator for assistance has been the industry's new push. Ford has begun replacing the RESCU system with VEMS (Vehicle Emergency Messaging System), a system that adds stolen vehicle tracking, remote door unlocking, and route assistance. The OnStar system from GM is standard on most high-end and fully upgraded vehicles and is scheduled to be standard on all its vehicles by 2002 (Lappin 1999). It offers the same services as the RESCU system, with the addition of an included on-board voice communication system. The OnStar switchboard operators can also perform concierge services such as making restaurant reservations. Some manufacturers are designing fully functional navigation systems to allow the vehicle operator to find and navigate to any destination without the need for assistance from a switchboard operator. BMW and some other manufacturers are also providing real-time traffic updates to avoid traffic jams via systems like the Etak/Metro Networks system, which covers 65 metropolitan areas and 90% of the metropolitan population of the United States (Etak 2000).

A positional accuracy of about 20 m is required for a vehicle operator to find a particular entrance to a parking lot or find a lost vehicle. Many navigation systems use GPS alone for computing the position that is used for navigation or emergency service dispatch. However, GPS alone does not provide sufficient coverage when operating the system

under signal masking conditions such as those encountered in an urban centre with tall buildings, a tunnel, or under heavy foliage. Multipath can also cause errors in the computed position, especially in urban canyon environments. The manufacturers will have to deal with this problem shortly due to the FCC (Federal Communications Commission) E911 (Enhanced 911) mandate that all cellular emergency assistance calls be accompanied by a position accurate to 125 m or better.

1.2 NAVIGATION SYSTEM REQUIREMENTS

A navigation system must meet strict requirements to be acceptable for use in a mass produced vehicle market. The system must be low cost or potentially low cost in the near future. Only inexpensive sensors are viable for mass-production automotive markets. The system accuracy must be around the 20 metre level after several minutes with limited GPS availability. Small position biases are often not critical since map-matching is used to snap the position to the correct road, but drifts must be controlled. The number of states in the Kalman filter and other computations must be kept to a minimum, since a limited amount of data and processing time would be made available to the navigation algorithm. In addition, it is not acceptable to make the driver perform any special initial or maintenance calibration procedure. The operator can not be hindered by or have the opportunity to cause errors in the navigation system.

1.3 BACKGROUND AND OBJECTIVES

The purpose of this research was to develop and evaluate a multi-sensor positioning system, addressing the requirements previously described. Only the positioning aspects of the system will be investigated. Likely, this position solution would be passed through a map-matching algorithm before it would reach the user. Other aspects of a telematics or navigation system, such as the UI (User Interface) and route finding algorithms, have not been investigated.

An integrated navigation system was described by McLellan (1992), which employed a series of low cost dead reckoning sensors developed by ETAK, including a compass with inclinometer, spinning rate gyro, odometer, and GPS. The Trimble Placer system's navigation algorithm, detailing the use of a rate gyro and odometer with GPS, was presented by Sushko (1993). Map aiding was explored by Bullock (1995), mainly to remove position bias caused by GPS selective availability, but other advantages of map assistance were discussed, as well. Hayashi (1996) used GPS with a fiber optic gyro and barometer to increase the availability and accuracy of GPS under signal masking. Zhang (1997) continued this work with the addition of a precise oscillator. Insight into the use of heading sensors was gained from Harvey (1998), who detailed the development of an integrated heading sensor. Kealy et al. (1999) addressed adaptive dynamics noise to maintain a better trajectory, termed "smart stochastic modeling".

A survey of available navigation aids was done to justify the selection of a set of sensors, and to gain insight into their costs and performance characteristics. These include GNSS technologies and augmentations, and dead reckoning sensors. From these, equipment was chosen for evaluation that covered most of the available sensor types. These sensors were

Precision Navigation TCM2 compass with inclinometer

Murata ENV-05D Gyrostar piezoelectric vibrating rate gyro

Differential odometry from the ABS (Anti-lock Braking System) on the test vehicle

Viatran model 246 pressure transducer (barometer)

NovAtel MiLLenium GPS receiver

Vectron CO-330 OCXO (ovenized quartz crystal oscillator)

The compass was removed from the system due to the difficulty experienced in calibrating out the permanent magnetic field of the car and compensation for vehicle pitch and roll with the inclinometer. The barometer failed after the first series of tests, and was deemed too expensive to be replaced. The equipment used in the final implementation was differential ABS odometry, the rate gyro, and GPS with height and

clock augmentations. Extended Kalman filtering was used to process the navigation data, while independent filters were employed to calibrate the dead reckoning sensors.

1.4 THESIS OUTLINE

Chapter 2 gives an overview of GPS and satellite positioning and its augmentations, while Chapter 3 gives an overview of the sensors and techniques used for dead reckoning. Chapter 4 deals with the formulation of the Kalman filter and math models that were used to describe the vehicle dynamics. Several adaptive features were added to the filter, and these are discussed as well. Chapter 5 gives detailed explanations of how the information from each of the sensors was used in the filter. Chapter 6 details three series of tests that were used to incrementally test the system, followed by conclusions in Chapter 7.

2 GNSS

GNSS refers to the use of one or more satellite navigation systems. The U.S. GPS (Global Positioning System) is one such system, as are the Russian GLONASS (Global Navigation Satellite System) and the proposed GALILEO system, a project funded by the EU (European Union). Recently developed augmentation satellite systems such as the FAA (Federal Aviation Administration) WAAS (Wide Area Augmentation System) and EGNOS (European Geostationary Navigation Overlay System) also provide supplementary capability. Several of these systems are described, beginning with GPS since it is the core of the system developed in this research, as are some of the factors that affect satellite positioning performance. Clock and height augmentations of normal satellite positioning are also described, and the methodology presented for several additions that were made to the University of Calgary's C3NAVG2 GPS/GLONASS code processing software.

2.1 GPS

GPS is a satellite based radio-navigation system designed by the U.S. DoD (Department of Defense). It is a spread spectrum CDMA (Code Division Multiple Access) system that broadcasts bi-phase modulated RHCP (Right Hand Circularly Polarized) signals centred at about 1.2 (L2) and 1.6 GHz (L1). This makes it an all weather system since there is little atmospheric absorption at these frequencies, and because it is spread spectrum, it is also resistant to jamming (Spilker 1996). Due mainly to geometrical spreading, the signal strength at the Earth's surface is only around –160 dbW, so the system operates basically through LOS (Line Of Sight).

The satellite constellation nominally consists of 24 satellites on 6 orbital planes with an inclination of 55° to the equator (U.S. JPO 1995). The satellites are in near circular orbit at a height of about 20,000 km above the surface, with a nominal period of just under 12 hours. The specifications call for a minimum of 21 operational satellites and three spares

to provide a minimum of four visible satellites at any place and any time on the planet. At present, the constellation has 27 satellites and unobstructed visibility is always seven or more satellites in the Calgary region.

Each satellite broadcasts a unique C/A (Coarse Acquisition) PRN (Pseudo-Random Noise) code modulated on to the L1 carrier. SA (Selective Availability) was used until May 1, 2000, to degrade the performance of civilian users in single-point mode through dithering of the satellite clock offset. A 50 bps navigation message, which gives the time of transmission of the PRN sequence and other necessary information, and a second code (P) are modulated on both the L1 and L2 carriers. These P code is intended primarily for military users and provides better resolution (less receiver noise) and immunity from SA, but is encrypted through a procedure known as AS (Anti-Spoofing). Civilians can still make use of the P code, but a loss is suffered through correlation techniques used to sidestep the encryption, and SA is not removed.

The MCS (Master Control Station) processes range measurements taken at the five monitoring stations and develops predictions for the orbits and satellite clock behavior. The MCS then sends this data to the monitor stations for upload of the navigation message to the satellites. The navigation message also includes the health of the satellites and an almanac that can be used to predict the visibility of satellites for any time and place.

2.1.1 GPS POSITIONING

Four or more satellites are normally required to compute a GPS position due to the use of low cost oscillators in commercial GPS receivers (Leick 1995). If the receiver clock were synchronized with GPS time, only three range observations would be required to compute the receiver coordinates in 3-D space. The fourth unknown, referred to as receiver clock bias, is the difference between the time estimated by the receiver and the GPS time. A range measurement with an error in time synchronization is referred to as a

pseudorange. Pseudoranges are most often used in a LS (Least Squares) parametric model to solve for the 4 unknowns in code positioning.

Three observations are normally made from the signal tracking procedures of a GPS receiver. These are the pseudorange, carrier phase, and Doppler measurements (Kaplan, 1996). The code pseudorange measurement is derived from a DLL (Delay Lock Loop) which correlates the incoming signal with locally generated versions of the signal. The discriminator uses the output of the correlators to make a pseudorange measurement. The carrier phase measurement is an accurate measure of the phase angle and accumulated cycles of the incoming signal. The carrier signal is normally tracked using a Costas PLL (Phase Locked Loop). A standard digital PLL is not used since the navigation message is usually left modulated on to the carrier, making the correlation positive or negative depending on the bits of the navigation message. When using Costas discriminators, the tracking error must be kept below 90°, unlike a standard PLL, which can tolerate errors of up to 180°. The Doppler is an estimate of the frequency difference between the locally generated carrier and the Doppler shifted incoming signal. It is noted the ionosphere causes a delay of the code and an equivalent advance of the carrier phase. The ranges are subject to several error sources, which are summarized in Table 2.1 for a typical C/A code receiver (Cannon 1997).

Table 2.1 Typical Errors for a C/A Code Receiver (Leick 1995)

Error Source	Typical Values	
	Differential (ppm[1])	Single-point (m)
Troposphere Delay	0.1-1	2-30
Ionosphere Delay	0.5-2[2]	2-50
Orbital Error	0.1-0.5	5-10
Clock Error	-	3-150[3]
Carrier Noise	-	0.001-0.006
Carrier Multipath	-	0.001-0.02
Code Noise	-	0.1-3
Code Multipath	-	0.1-100

The single point positioning observation equations can be given by (Cannon 1997)

$$P = \rho + d\rho + dT + dt + dTrop + dIono + \varepsilon_P \tag{2.1}$$

$$\Phi_P = \rho + d\rho + dT + \lambda_G N + dt + dTrop - dIono + \varepsilon_\Phi \tag{2.2}$$

and for differential

$$\Delta P = \Delta\rho + \Delta d\rho + \Delta dT + \Delta dt + \Delta dTrop + \Delta dIono + \varepsilon_{\Delta P} \tag{2.3}$$

$$\Delta\Phi_P = \Delta\rho + \Delta d\rho + \Delta dT + \lambda_G \Delta N + \Delta dt + \Delta dTrop + \Delta dIono + \varepsilon_{\Delta\Phi} \tag{2.4}$$

where

P is the pseudorange measurement

---

[1] Differential GPS errors are usually given in pats per million of baseline length

[2] The ionosphere can cause higher delays as well as scintillation (nearly instantaneous changes in phase and frequency of the signal) during geo-magnetic storms.

[3] This term is specified 100 m 2DRMS due to SA. With SA off, the value will be 0.1-3 m

ρ is the geometric range from the satellite to the receiver

dT is the satellite clock error which includes SA dithering when applicable

dt is the receiver clock offset to GPS time

dTrop is the tropospheric delay

dIono is the ionospheric delay

ε is the measurement noise and multipath

Φ_P is the carrier phase measurement

λ_G is the GPS carrier wavelength

N is the integer ambiguity

Δ is an operator to denote the difference between receivers on corresponding satellites

DGPS (Differential GPS) can be used to minimize the errors of single-point GPS by canceling the parts of the error that are common to receivers in close proximity. Differential GPS is normally implemented by differencing the ranges to common satellites from two receivers. If the coordinates of one station are known, an accurate position of the second station can be determined. Alternatively, a coordinate difference between stations can be computed using approximate coordinates for one of the stations. Differential GPS reduces or eliminates errors caused by satellite clock and orbital errors, and atmospheric propagation. It does not reduce multipath, and the noise of a differenced observation is larger than that of an individual measurement by a factor of $\sqrt{2}$.

2.1.2 MULTIPATH

Multipath is a highly variable error source that refers to the reception of a signal through any path other than the minimum travel time path (Braasch 1996). Most often multipath is caused by reflections, but in some circumstances, it may also involve diffraction or refraction. Multipath is weakly spatially correlated, but sometimes correlated for periods up to several minutes in time. It is broadly categorized into two types which are shown in Figure 2.1. Specular multipath is characterized by parallel incident rays remaining parallel after reflection. This type is clearly correlated and results from reflection from

smooth, regular surfaces. Diffuse multipath usually results from reflection from rough surfaces such as rough seas or a gravel road. Little correlation is seen in this type of multipath since the reflected power is distributed over a large spherical angle.

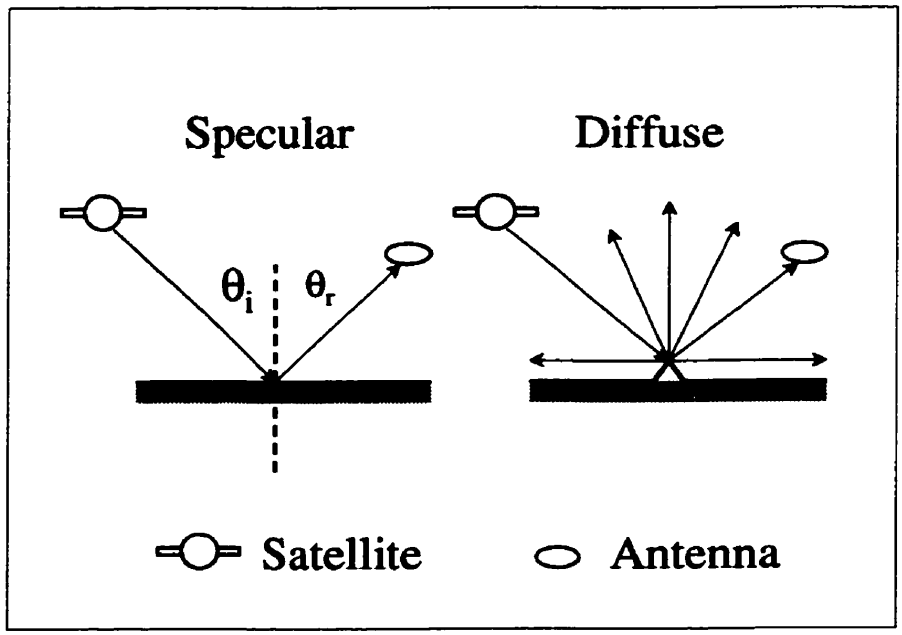

Figure 2.1 Specular and Diffuse Multipath (Cannon 1997)

The character of a multipath signal depends on many factors. Firstly, the geometry is necessary for a signal to reflect toward the antenna. Second, EM properties of the medium control the amount of absorption, reflection and refraction of the signal from a surface. The polarization of the signal usually becomes elliptical after reflection, and also changes to LH polarized if its incidence angle is less than the Brewster angle. These properties essentially control the power, polarization and phase angle of the reflected signal. The size of the reflecting surface is also a consideration since very little signal would be reflected toward the antenna from a small surface. Stationary receivers experience clear correlation of multipath with correlation times of several seconds to several minutes, while moving vehicles experience some multipath correlation depending on geometry.

Antennas are designed to minimize the effect of multipath signals by having a lower gain near the horizon, since most multipath signals are reflected off the ground. Antennas are also designed with a gain of around 10 dB for RHCP signal over LHCP at higher elevations, but this gain (axial ratio) usually drops toward the horizon. It should be noted that conductive materials have a large Brewster angle, while non-conductive materials tend to have a smaller Brewster angle. For example, it is much more likely to have the reflected GPS signal be RH polarized off of a sheet of glass than a door made of steel. However, the proportion of the signal that is reflected is much higher for conductive materials. Back lobes are minimized to prevent reflections from just below the antenna bouncing up to the back of the antenna.

The range error resulting from multipath signals entering a GPS antenna is therefore a combination of the geometry, reflection sources, and antenna properties. This error usually does not exceed 5-10 m, but has been shown to exceed 100 m (van Nee 1992). The sum of the direct and all multipath signals from a satellite can be given by (Braasch 1998).

$$s(t) = -Ap(t)\sin(\omega_0 t) - \sum_{i=1}^{n} \alpha_{mi} Ap(t+\delta_i)\sin(\omega_0 t + \theta_i) \tag{2.5}$$

where

s(t) is the total incoming signal

A is the amplitude of the direct signal

p(t) is the PRN sequence of –1 or 1

ωis the frequency of the direct signal

n is the number of multipath signals

α_m is the relative power of the multipath signals

δ is the delay of the multipath signals relative to the direct signal

θ is the phase of the multipath signal relative to the direct signal

The presence of multipath in the correlation causes asymmetry in the correlation function. This can be seen in Figure 2.2. The performance of a typical 1 chip spacing discriminator in the presence of multipath is also shown in Figure 2.3.

Multipath can cause errors in ranging which are both positive and negative, depending on the strength, phase and delay of the multipath signal. An approximate envelope that contains the entire range of errors for a 1 chip spacing early-late C/A code discriminator with a multipath to signal ratio of -20 dB can be found in Figure 2.4. If the multipath and direct signals are in phase, a maximum positive range error will occur, while if the signals are 180° out of phase, a maximum negative range error will occur. If the two are separated by a phase angle of 90°, no error is induced. Improved discriminators can reduce the error envelope, but receivers with better discriminators are generally more expensive to purchase.

Unlike code multipath, carrier phase multipath can only have values up to one quarter of a wavelength, due to the nature of the PLL. A derivation for carrier phase error due to multipath and experimental results can be found in Ray (2000). Typical values for carrier phase multipath seldom exceed 1-2 cm.

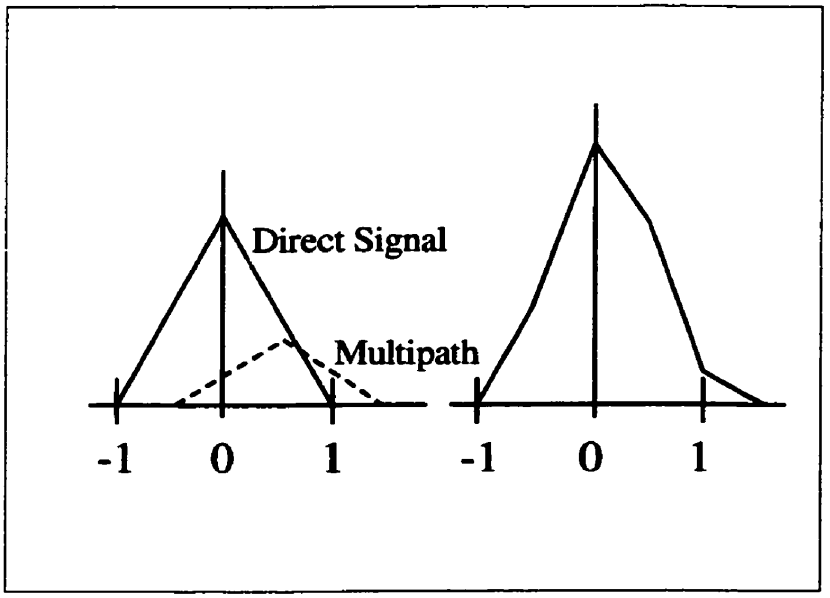

Figure 2.2 Correlation Function in the Presence of Multipath (Braasch 1998)

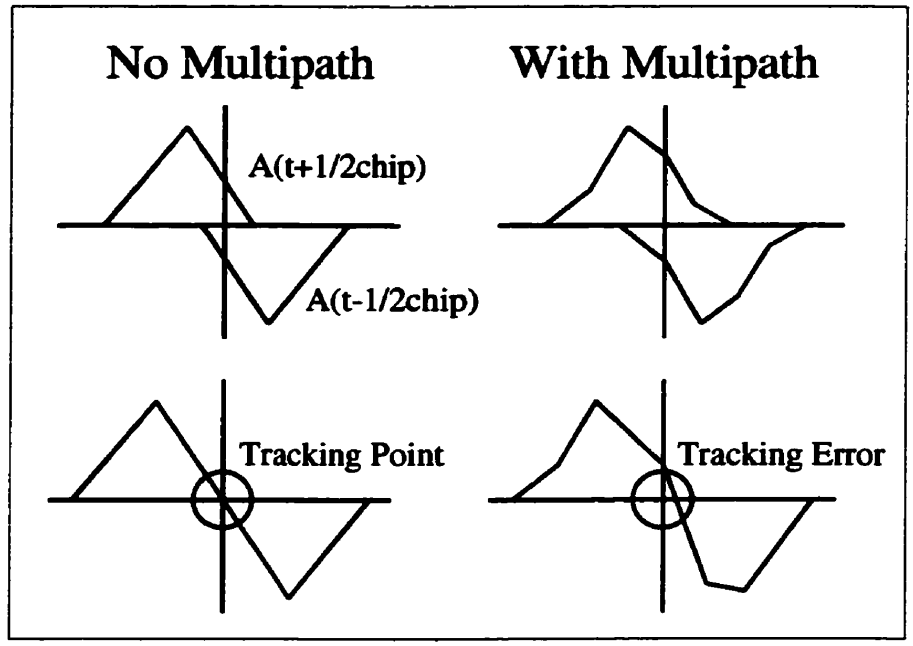

Figure 2.3 Discriminator Function in the Presence of Multipath (Braasch 1998)

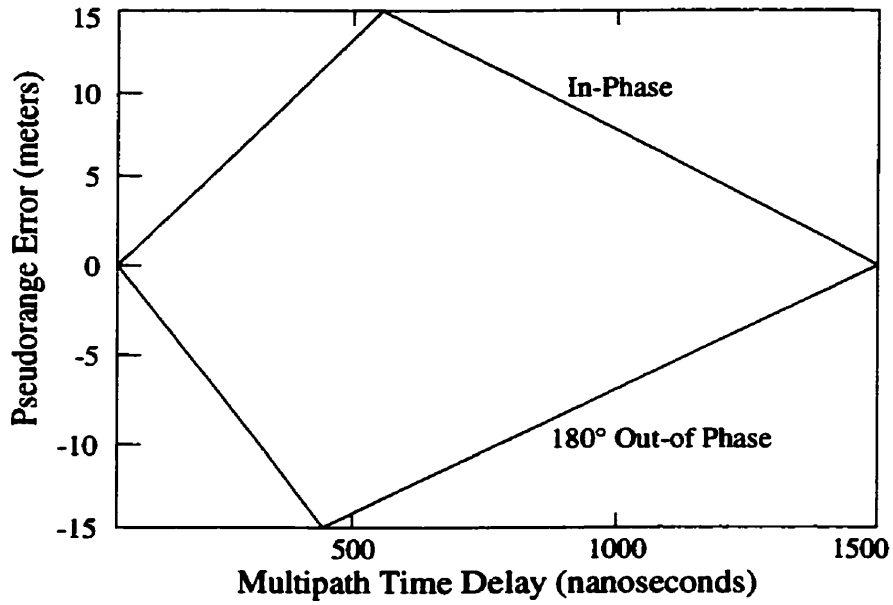

Figure 2.4 Error Envelope for a 1 Chip Spacing Early-Late C/A DLL (M/D –20 dB) (Braasch 1997)

2.1.3 CARRIER SMOOTHING

A common processing technique that combines the absolute pseudorange measurement with the high accuracy of the carrier phase measurement is carrier smoothing. The algorithm essentially estimates the ambiguity term of the carrier phase observation through combining the two observations over time. It thereby smoothes the noisy pseudorange measurements and reduces the effect of code multipath errors, which are much greater than multipath on the carrier phase. A carrier smoothing algorithm is used in the C3NAVG2 software based on the formulation of Hatch (1982)

$$Q_j = \alpha_j \left(Q_{j-1} + \delta\Phi_{j-1,j} \right) + (1-\alpha_j) P_j \tag{2.6}$$

where

Q_j is the smoothed pseudorange

α_j is the weight which varies from 0 to ≤1.0 by small increments (≤0.1)

$\delta\Phi_{j-1,j}$ is the change in phase over the epoch

P_j is the measured pseudorange

The carrier smoothing algorithm has a fault in the fact that the signs are opposite on the ionospheric effect term in the two observation equations. This means that the smoothed range will diverge from the true range by a value of double the change in the delay. The problem is usually avoided by initializing a new ramp periodically, or leaving the weight at slightly less than one. Another limitation of carrier smoothing over ambiguity resolution techniques is the danger of initializing the smoothed pseudorange to a biased value. If the pseudorange is biased by multipath during the initialization period, the bias will remain in the smoothed pseudorange until a new ramp is initialized. This error can be somewhat limited also by keeping the weight at a value less than one.

2.2 RELIABILITY

Reliability measures were investigated to see the benefit of adding augmentations to GPS satellite positioning. However, reliability criteria were not implemented in the software used in this project. The data snooping technique to detect blunders is developed concurrently to the internal reliability concept.

There are two aspects to reliability, external and internal (Leick 1995). Internal reliability refers to the ability to detect blunders in the measurements. There must be at least one more measurement than unknowns to make blunder detection (data snooping) possible. External reliability refers to the effect that undetected blunders could have on a solution. The MDB (Marginally Detectable Blunder), which is the greatest blunder which could pass statistical scrutiny, is often used to characterize internal reliability. The largest horizontal component of the error caused by missing one of these blunders is referred to as the MHE (Maximum Horizontal Error). The HDOP (Horizontal Dilution of Precision) is also used for evaluation of the geometry used in formulation of the LS. Only horizontal components are discussed here since they are of primary concern to land navigation users. It should be noted that statistical testing done in this project only tested for a single blunder; the possibility of multiple blunders was not considered.

The redundancy number, v_i, for a particular observation is given by

$$v_i = \left(P_r R^{-1}\right)_{ii} \tag{2.7}$$

$$P_r = R - H\left(H^T R^{-1} H\right)^{-1} H^T \tag{2.8}$$

where Pr is the covariance matrix of the residuals, R is the covariance matrix of the observations and H is the design matrix. The redundancy numbers can take on values from zero to one and are a function of the accuracy and geometry of the observations only.

The test value for comparison can be interpreted from Figure 2.5. This figure shows the probability of committing each of the two possible statistical errors. Type I errors refer to the exclusion of an observation which does not have a bias or blunder. Type II errors refer to the inclusion of an observation that has a bias or blunder. The probability of committing a type I error is α_0, while the probability of committing a type II error is β_0. By choosing values for the two error probabilities, a non-centrality parameter, δ_0, can be generated (Krakiwsky 1987). Using values of 0.1% and 10% for α_0 and β_0, respectively, the non-centrality error becomes 4.57, which is a common choice used for blunder detection.

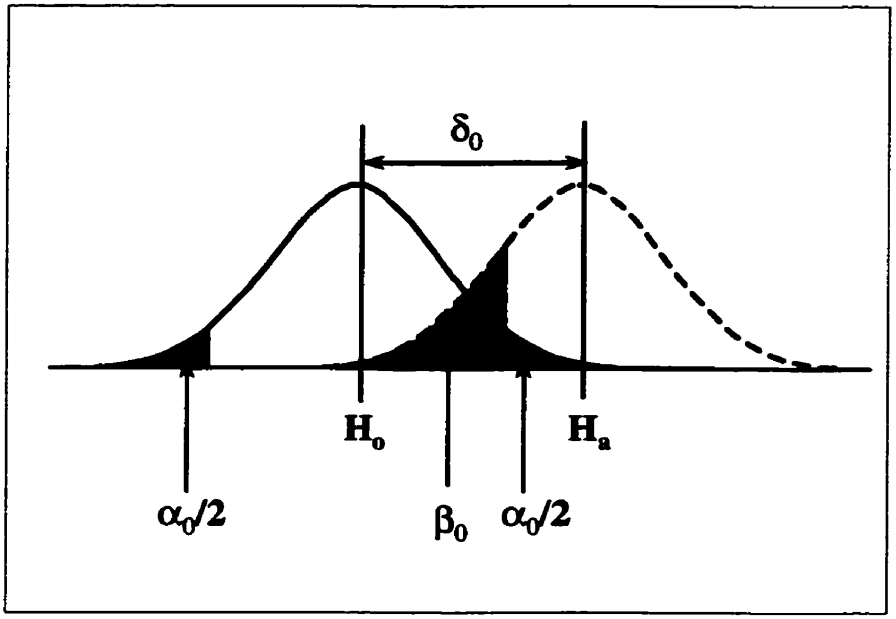

Figure 2.5 Probability of Type I and II Errors and the Non-Centrality Parameter (Krakiwsky 1987)

The measurement residuals must be statistically tested to detect blunders through

$$\frac{r_i}{\sqrt{v_i}} \leq \delta_0 \tag{2.9}$$

A similar equation gives the MDB for each observation

$$\nabla_i = \delta_0 \sqrt{\frac{R_{i,i}}{v_i}} \tag{2.10}$$

And the effect of this blunder on the unknowns is given by

$$\Delta X_i = -PA^T R^{-1} \nabla_0 \tag{2.11}$$

where ∇_0 is a column vector of zeros except for the i^{th} MDB, and P is the covariance matrix of the unknowns. The MHE would simply be the largest horizontal component of these values computed for each observation. With $\Delta\varphi$ being the latitude component of ΔX, and $\Delta\lambda$ the longitude component

$$MHE = Max\left\{\sqrt{\Delta\varphi_i^2 + \Delta\lambda_i^2}\right\} \tag{2.12}$$

GPS by itself is often not sufficient to prevent or detect massive blunders in the position solution. Several figures follow which show the effect of adding augmentations to GPS on reliability (Lachapelle et al. 1997). Figure 2.6 shows cumulative probability distributions of HDOP and Figure 2.7 cumulative probability distributions of MHE for a 24-hour DGPS simulation. A 20° isotropic mask angle was used with the user located in Calgary. Augmentation of DGPS positioning with a clock model observation (σ = 3 m), height observation (σ = 3 m), and GLONASS are shown. The GLONASS constellation is the constellation as of August 1, 1997. The augmentations improve the HDOP somewhat, but the MHE drastically. For example, without any augmentation, there would be a 90% chance that a blunder causing more than 20 m of horizontal error would be allowed into the solution. With all of these augmentations, the probability of this happening is less than 5%. These augmentations will be discussed in the following sections.

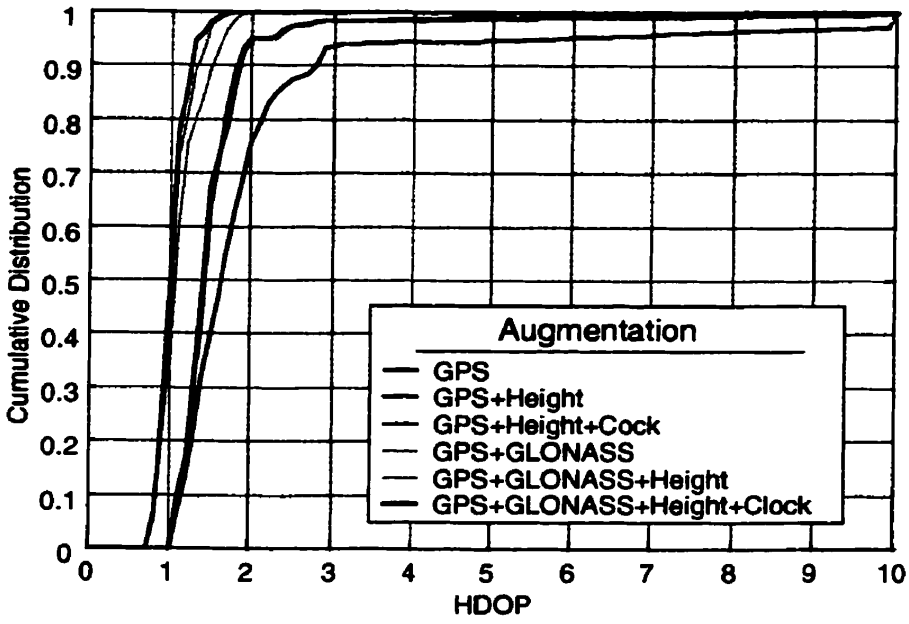

Figure 2.6 Cumulative HDOP Distribution for GPS Augmentations (Ryan et al. 1999)

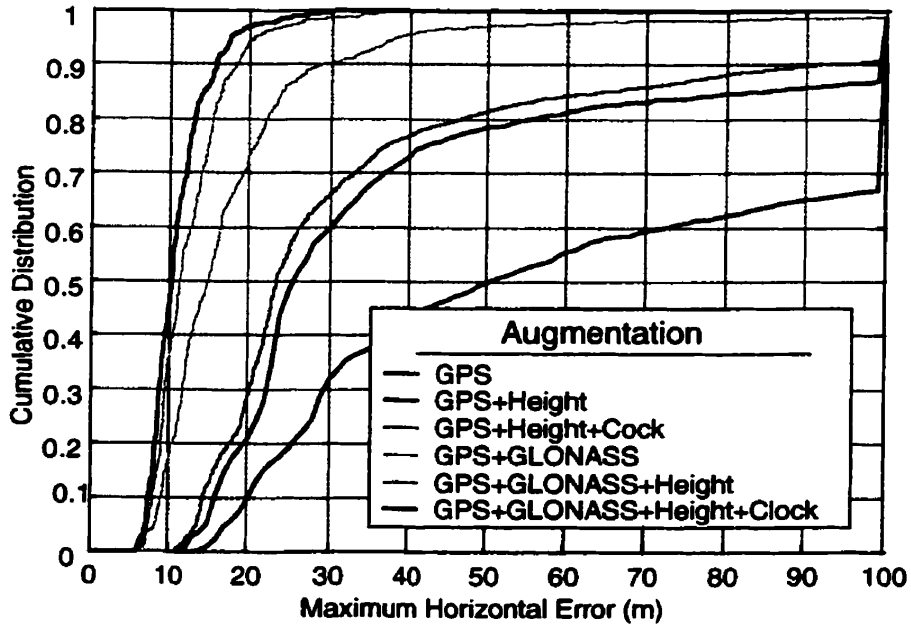

Figure 2.7 Cumulative MHE Distribution for GPS Augmentations (Ryan et al. 1999)

2.3 SATELLITE AUGMENTATION

Augmentation is the use of other available information to increase the availability and reliability of GPS positioning. Availability refers to the amount of time that a position fix can be computed. There must be at least as many observations available as unknowns, i.e. 4 for a 3-D fix with an unknown clock offset. Satellite augmentation is the use of other satellite systems to improve availability and reliability. Several satellite augmentations are presently possible including the ranging capability of WAAS geostationary satellites, and use of the GLONASS system.

2.3.1 GLONASS

A great deal of research into the benefits of using GLONASS has been carried out over recent years (Hall et al. 1997), and results have shown substantial improvements in accuracy, availability, and reliability of GLONASS augmented positioning (Lachapelle et al. 1997). Some questions as to the future of the system have arisen since the constellation has been failing and is not being maintained at the level specified in the Interface Control Document (Russian Space Command 1995). The system was recently moved to the control of a different branch of government, so additional funds might be procured to refurbish the constellation. Commercial GPS/GLONASS receivers are now common, but are much more expensive than the low cost receivers employed in most automotive navigation applications.

The GLONASS system uses different frequencies for each pair of satellites (2 satellites on the same orbital plane but opposite sides of the Earth use the same frequency), as opposed to the GPS system which uses unique codes chipped on to the same carrier frequency. The procedures for monitoring and uploading to the satellites are similar. There is an offset of 3 hours with a wander of several milliseconds between GLONASS time and UTC (Universal Time Coordinated). GPS time is virtually synchronized with UTC, but has an offset (13 seconds at present) due to UTC leap second corrections meant

to synchronize UTC with celestial phenomenon. The offset between GPS and GLONASS is therefore close to 3 hours 13 seconds (Kaplan 1996).

GLONASS augmentation requires the addition of one extra unknown to the four parameters used in GPS pseudorange positioning. This extra GLONASS clock bias parameter is necessary even in differential operation because of the variable hardware delays in GPS/GLONASS receivers. The GLONASS signals are at different frequencies from GPS, so they behave differently through the analog front-end of the receiver. Manufacturers often try to minimize these differences through a calibration procedure, but experience with two different receivers (NovAtel GG and Ashtech GG24) have shown that GLONASS ranges appear biased by values of several metres. A second clock offset or bias term is therefore needed to use GLONASS ranges accurately. The bias changes with temperature, power, and other factors that affect analog circuitry, but can be considered constant for a period of several minutes or more under normal conditions.

2.3.2 WIDE AREA AUGMENTATION SYSTEM

Using WAAS satellites is a simple and effective augmentation to GPS since they provide a ranging signal, have fixed positions (geostationary), and are synchronized to GPS time. They operate at the same frequency as GPS so the analog front-end of a GPS receiver does not have to be redesigned to accommodate the signals. Their main limitations are the small number of them (two at this time), and their low elevation at higher latitudes. This limitation can be seen in Figure 2.8.

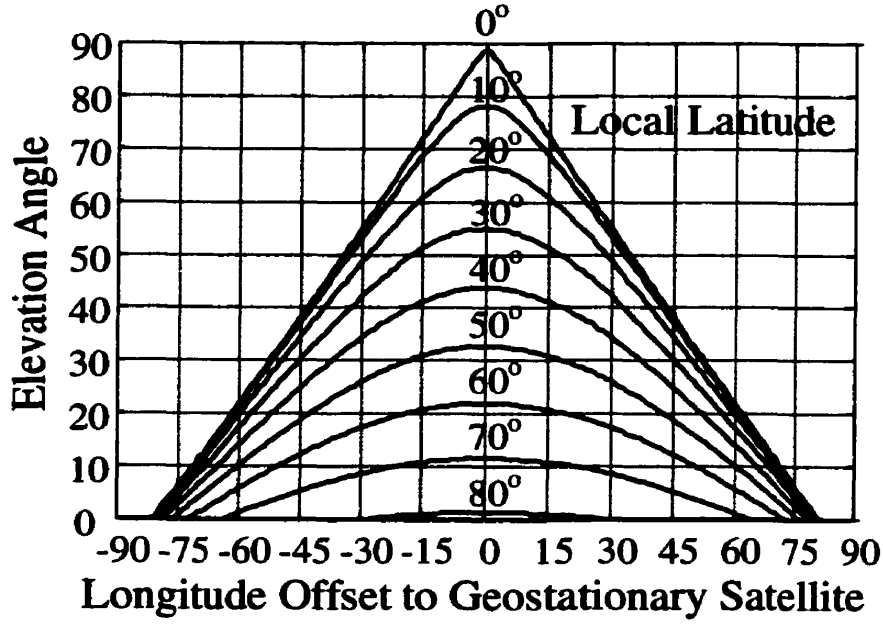

Figure 2.8 Elevation Angle for Geostationary Satellites (Ryan et al. 1999)

Prior to the termination of SA, many low cost receiver manufacturers were adding WAAS capabilities to their receivers due to the relative ease of adding this functionality and the benefit to users. With improved single point positioning, it is unlikely that most consumer markets will need the differential corrections broadcast on the WAAS signal. However, the signal does also have a ranging capability that may be desirable. The corrections are made possible by a network of 25 monitoring stations that measure the ranging errors of the satellites visible over most of North America. This technique is called WADGPS (Wide Area DGPS). Corrections are computed at master stations and uploaded to the WAAS satellites for broadcast. The ranging errors are separated by the processing software into their constituent error sources (ionosphere, troposphere, and satellite clock errors) and broadcast in parameters such that corrections can be derived for any place near or inside of the network. Similar systems are under construction that will operate over different regions of the Earth.

2.4 HEIGHT AUGMENTATION

Height augmentation is a more elegant way of accounting for a height measurement or estimate than removing the coordinate from the LS solution. Fixing to the last known height and solving for 2-D coordinates is a feature of most GPS receivers. Several problems can arise from this technique, such as introducing biases in the latitude and longitude while fixed to an erroneous height. Often before a GPS receiver switches into 2-D mode due to signal masking, the position is already degraded by multipath. The last 3-D position may have had an error in height of 100 m or more under high multipath or single point positioning with SA. Filtering of the position solution can minimize blunders and reduce some of the short-term multipath, but vehicle dynamics tend to be erratic and difficult to model. Many situations also arise where it is undesirable to have a filtered solution because the results are to be filtered by another algorithm. If the height is simply fixed, an assumption is made that the height is perfectly known. This results in an overestimate of the accuracy of the 2-D solution.

Height augmentation adds the height estimate as an LS observation. The estimate can be based on the history of heights, measured by a barometric sensor, or determined from a digital source. If an appropriate variance is used with the observation, sound statistics can be computed for the position parameters. Height information can also serve a purpose in locations with no GPS coverage such as tracking the level of a car in a parkade. A method of deriving height information from the sampled output of an analog barometer is discussed below.

2.4.1 BAROMETRIC AUGMENTATION

Barometric pressures can be used to aid in height determination in two ways, i.e. in a relative sense or through a barometric height model. Relative height changes can be computed by estimating the density of air and coming up with the height/pressure gradient (Kyle, 1991).

$$dP = -g\rho dH \tag{2.13}$$

$$\rho = MM_0 \frac{P}{KT} \tag{2.14}$$

where

dP is pressure change

g is gravity (9.81m/s)

ρ is the density of air

dH is the height difference

M is the molecular weight of air (-28.96 amu)

M_0 is the mass of 1 amu (1.66e-27 kg)

K is Boltzman's constant (1.38e-23)

T is the temperature in K

One can also use the standard volume of a mole of gas (22.141 L at 101.325 Pa and 273.15°K) to arrive at a rough value of 1.292kg/m^3.

The second method of determining heights is using a standard atmosphere model. The U.S. Standard Atmosphere Model (Lutgens et al. 1982) gives the constants for a fifth order polynomial in pressure to determine the pressure altitude, or barometric height.

$$h_b = a_0 + a_1 p + a_2 p^2 + a_3 p^3 + a_4 p^4 + a_5 p^5 \tag{2.15}$$

where

h_b is barometric height (m)

p is the observed pressure (kpa)

$a_5 = 5.846393362917872e\text{-}8$

$a_4 = -2.814440431082330e\text{-}5$

$a_3 = 5.643398665717297e\text{-}3$

$a_2 = -6.07406590143939ge\text{-}1$

$a_1 = 3.776791975661281e+1$

$a_0 = -1.424633727563486e+3$

This model gives height above sea level assuming average atmospheric conditions and standard temperature and pressure at sea level. It does not take into consideration temperature or humidity variations, and is therefore inaccurate in an absolute sense. It can be used if a bias, and possibly a scale factor, can be calculated between the GPS ellipsoidal height and the model height.

A hybrid method of filtering the barometer data was implemented in C^3NAVG^2. The software was already capable of using a last known or fixed height as an observation, so a Kalman filter was added to generate height estimates from the barometric. Following the notation of Chapter 4, the barometric filter state vector is

$$X = \begin{bmatrix} h_g \\ b_b \\ e_b \end{bmatrix} = \begin{bmatrix} GPS\ height \\ Barometric\ \ height\ \ bias \\ GM\ \ barometric\ \ height\ \ error \end{bmatrix} \tag{2.16}$$

with state transition matrix

$$\Phi = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 - e^{-\Delta t / \tau_m} \end{bmatrix} \tag{2.17}$$

and process noise matrix

$$Q=\begin{bmatrix} \sigma_h^2\Delta t & 0 & 0 \\ 0 & \sigma_b^2\Delta t & 0 \\ 0 & 0 & \sigma_e^2\frac{\tau_m}{2}(1-e^{-2\Delta t/\tau_m}) \end{bmatrix} \tag{2.18}$$

The observation is an integrated pressure height with observation model

$$h_b = h_g + b_b + e_b \tag{2.19}$$

The barometric pressure model was chosen to account for the slowly changing atmospheric pressure using a random walk for the barometric height bias. Local variations usually associated with wind and temperature are handled with the Gauss-Markov height error.

The Viatran Model 246 pressure transducer was originally used to collect barometric pressure data to aid in height determination. Although it is too expensive to be used for a low-cost automotive application (about $700 in small quantities), lower quality sensors are available at a reduced cost. The instrument uses a silicon diaphragm with Wheatstone bridge arrangement of piezoresistors mounted on the surface as shown in **Figure 2.9.** A Wheatstone bridge is a circuit with four balanced resistance outputs. When pressure causes the diaphragm to bulge, the inner sensors are put under tension and the outer sensors are put under compression, thereby changing the resistance of the piezoresistors. The differences in resistance are measured and converted to a pressure reading. Specifications of this unit are given in **Table 2.2.**

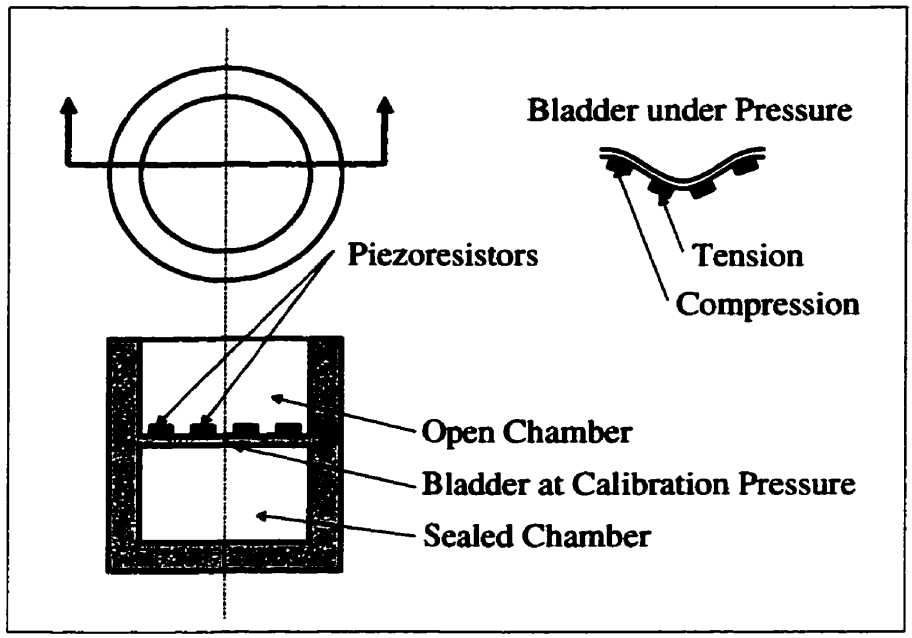

Figure 2.9 Piezoresistive Pressure Sensor (Viatran 1998)

Table 2.2 Selected Viatran Model 246 Specifications (Viatran 1997)

Characteristic	Min	Max	Units
Pressure Range	846.6	1083.6	mb
Total Error	-	0.3	mb
Response Time	-	5	ms
Supply Voltage	8.5	40	VDC
Output Signal	0	5	VDC
Operating Temperature	-17.8	93.3	°C
Supply Dependence	-	2.3e-4	mb/V

The barometer was shipped with factory calibration numbers that were used to generate a pressure estimate from the output voltage. The pressure change over an interval is scaled by the estimated pressure-height scale factor to give an integrated height change. This integrated approach is necessary to avoid the large changes in model height associated with multiplying the raw output by a changing scale factor. This scale factor is determined by running a low-pass filter on the ratio between GPS height changes and barometric pressure changes. The scale factor does not usually cause large height errors since it normally varies by less than 5%, or one metre on every 20 metres of vertical variation.

To minimize errors, an appropriate enclosure is desirable so that motion of the vehicle or wind passing over the vehicle does not create high or low pressure in the enclosure. In addition, the pressure must be measured without extreme temperature variations, i.e. away from radiant sunlight, the engine, or interior climate control.

Some preliminary testing was done with the barometer fixed to a roof rack on the vehicle with no enclosure. Spikes in the barometric pressure corresponding to elevation changes of 50 m were found. Some results of this initial testing are shown in Figure 2.10. Note that the GPS height trajectory should have been accurate to several metres since a good constellation of satellites was observed throughout the test. The results of this testing were substantially worse than those previously reported for this sensor (Hayashi 1996), but this is attributed to the poor placement of the barometer. Testing with the barometer located inside the vehicle or an enclosure could not be performed because the barometer failed before subsequent tests could be performed, and could not be repaired.

Other research using barometry with kinematic carrier-phase DGPS (Schleppe et al. 1995) has shown that meter level height accuracy is possible using barometry. Testing was performed in an area with heavy canopy, which is notorious for poor GPS vertical positioning, and differential barometry was shown to give repeatability better than 0.5 metres.

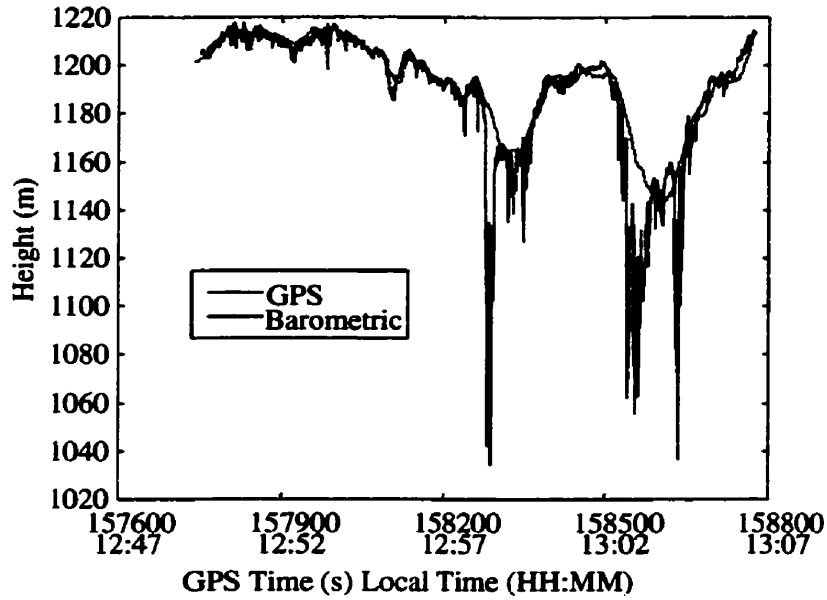

Figure 2.10 Raw GPS and Barometric Height Trajectories

2.4.2 FIXED VALUE AUGMENTATION

The approximate height of an area can be used as an observation for the LS when height augmentation is needed. Data for this type of augmentation can be acquired from a DTM (Digital Terrain Model) or some other data source with height information. Data of this type can be found for most urban areas with an accuracy of 5-20 m. Due to the failure of the barometer, this type of height augmentation was used in processing the results found in later results. Tests that indicate height augmentation were done in areas with small vertical variations so that an approximate height could be used throughout the test.

2.5 CLOCK AUGMENTATION

Clock augmentation refers to the use of a stable oscillator to be able to model the behavior of the receiver clock offset. The clock offset of a receiver using a stable oscillator can usually be modeled as a linear or quadratic function over a period of several seconds to several hours, depending on the quality of the oscillator. When the

number of range measurements is insufficient for a solution, or the geometry of the solution is poor (high DOP), the model can be used to predict the value of the clock offset.

Much of the research into the use of precision oscillators was done for RAIM (Receiver Autonomous Integrity Monitoring), which is a variation of the reliability testing described above (Lee 1993). Several methods have been used for estimating the parameters of a function to describe the clock offset. The first involves using snapshots of the clock offset from several minutes as observations in a LS adjustment (Misra et al. 1995). This model would then be used for several minutes while new observations of clock offset were accumulated. A Kalman filter method was chosen for this implementation because it seemed to outperform the LS method. This was likely due to the fact that the oscillator does not have sufficient time to settle after power-on, and the Kalman filter reacts continuously to the drift changes of the oscillator.

Both remote and reference receivers employed Vectron CO-330 10 MHz OCXOs as external clock references to allow for clock coasting in times of reduced satellite visibility. This clock has excellent stability for periods up to several minutes (Petovello and Lachapelle, 2000). Some specifications of the oscillator are given in Table 2.3.

Table 2.3 Vectron CO-330 OCXO Specifications

Characteristic	Specification
Temperature Range	0-50 °C
Electrical Tuning	0-6V $1x10^{-7}$
Mechanical Tuning	$1x10^{-9}$
Aging	$3x10^{-7}$/20 years
Supply	12 V ± 5%
Allan Variance	$5x10^{-12}$/s
Power Consumption	<10 W startup, <5 W continuous
Output sine	>0.5 VRMS into 50,0. min +7dBm

If the monitor receiver uses a precise oscillator, or with SA turned off, the user receiver can effectively use a model to estimate the clock bias and remove one unknown from the

GPS solution (Zhang, 1997). This allows position solutions with as few as two satellites, if the heights are also constrained. Aiding the determination of the clock offset from a weakly determined GPS solution is also advantageous, since measurement errors will not be absorbed into the clock offset. Height determination is dramatically improved through the use of clock aiding since there is a high correlation (which is always positive) between height and clock offset estimates, usually exceeding 0.8. Although the correlation between horizontal position and clock offset estimates takes on both positive and negative values depending on geometry, the correlation coefficients can reach near ± 1 with DOPs as low as 1.5 (Zhang, 1997).

Another advantage of using a precise oscillator is that the receiver can develop a stable clock model that can be used to accurately predict the Doppler frequency for a lost or rising satellite. This means that only code-shift bins have to be tested. The NovAtel receiver uses this technique by remaining in fast reacquisition mode for several seconds after a signal is lost; it can reacquire a signal almost instantaneously if the signal is restored.

The C^3NAVG^2 software was modified to take advantage of the use of a stable oscillator. Essentially, a filter is trained with clock offset solutions when they are available, and is used to predict the clock offset when geometry becomes poor or there are too few satellites to compute a position. The clock offset is modeled as a linear drift in a four-state Kalman Filter with state vector

$$X = \begin{bmatrix} dt \\ di \\ e_t \\ b_{gg} \end{bmatrix} = \begin{bmatrix} GPS\ clock\ bias \\ Clock\ drift\ rate \\ G-M\ drift\ rate\ error \\ GPS-Glonass\ time\ bias \end{bmatrix} \tag{2.20}$$

The need to model a time bias between GPS and GLONASS time was previously explained in the section on GLONASS. The state transition matrix is

$$\Phi = \begin{bmatrix} 1 & \Delta t & \tau_m\left(1-e^{-\Delta t/\tau_m}\right) & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1-e^{-\Delta t/\tau_m} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{2.21}$$

and the process noise matrix is

$$Q = \begin{bmatrix} Q(1,1) & \sigma_i^2 \frac{\Delta t}{2} & Q(1,3) & 0 \\ \sigma_i^2 \frac{\Delta t}{2} & \sigma_i^2 \Delta t & 0 & 0 \\ Q(1,3) & 0 & \sigma_e^2 \frac{\tau_m}{2}\left(1-e^{-2\Delta t/\tau_m}\right) & 0 \\ 0 & 0 & 0 & \sigma_b^2 \Delta t \end{bmatrix} \tag{2.22}$$

$$Q(1,3) = \sigma_e^2 \tau_m^2 \tfrac{1}{2} e^{-2\Delta t/\tau_m} + \tfrac{1}{2} - e^{-\Delta t/\tau_m} \tag{2.23}$$

$$Q(1,1) = \sigma_t^2 \Delta t + \sigma_i^2 \Delta t^3 + \sigma_e^2 \tau_m^3 \left(\tfrac{\Delta t}{\tau_m} + 2e^{-\Delta t/\tau_m} - \tfrac{1}{2} e^{-2\Delta t/\tau_m} - \tfrac{3}{2} \right) \tag{2.24}$$

The observations are the clock offsets and drift rate computed when sufficient satellites are available to compute a position solution unaided. The drift rate observation comes from the velocity solution, and is based on the Doppler measurements. GPS and GLONASS share the same drift rate, since both systems use a very accurate time base. The variances of these observations are simply the output variances from the LS position and velocity computation. The observation model is

$$Z = \begin{bmatrix} dt \\ \dot{dt} \\ b_{gg} \end{bmatrix} = \begin{bmatrix} GPS\ clock\ offset \\ Clock\ drift \\ GPS - GLONASS\ time\ bias \end{bmatrix} \tag{2.25}$$

The clock filter was chosen to suit the conditions encountered when using a precision oscillator in a moving vehicle. Short-term frequency variations occur due to vehicle dynamics putting stress on the crystal; the Gauss-Markov drift state would account for this. Long term drift change of the oscillator is accounted for with a random walk. The oscillator experiences rapid changes in frequency over the first hour and will continue to stabilize for several days, but it is many times more stable than the usual TCXO (temperature compensated quartz crystal oscillator), even on startup. If power consumption was low enough, the oscillator could be left operating at all times. The values chosen for the process noise are summarized in Table 2.4. They are conservative for an oscillator of this nature, but need to be since the receiver is solving for both position and velocity state concurrently to the clock offset and drift rate. In addition, the oscillator will likely not be operating at steady state, and there is a possibility of high noise and multipath while the filter is stabilizing. The values must therefore be allowed to vary quite easily.

Table 2.4 Spectral Densities for the Kalman Filter Clock Model

Process Noise	Value
GPS Clock Bias	1.0E-10 m^2/s
Clock Drift Rate	1.0E-8 m^2/s^3
G-M Drift Rate Error	1.0E-6 m^2/s^3
GLONASS Clock Bias	1.0E-20 m^2/s

The performance of the filter is demonstrated in Figure 2.11. The filter was allowed to initialize with increased amounts of C^3NAVG^2 clock offset and drift data and then made to predict the future clock offsets through the dynamics model. These predictions are compared to the independent epoch to epoch C^3NAVG^2 offset values for the entire data set and the differences are plotted. The batch LS solution using a quadratic model is also presented for comparison in Figure 2.12. The data set contains a spike due to multipath during the two minute static initialization. Other data sets were examined in making the decision to use the Kalman Filter method over LS, only a few of which showed such correlated spikes while stationary, and none showed them while in motion. Only one initialization time with increasing amounts of data is shown for both methods. The start

time has a greater effect on the LS method than the Kalman Filter since all of the states have process noises high enough to allow them to adapt quite quickly. Maximum errors for the two Strathcona tests, which are described in Chapter 4, are found in Table 2.5.

Table 2.5 Maximum Clock Offset Prediction Error (m)

Initialization Time (s)	Quadratic LS		Kalman Filter	
	Test 1	Test 2	Test 1	Test 2
100	473.9	8.5	29.9	21.7
200	51.8	34.1	28.5	17.3
300	20.4	17.6	25.3	17.6
400	6.1	19.6	18.7	11.2
500	16.4	10.9	9.7	8.2
600	9.7	5.8	5.2	5.6

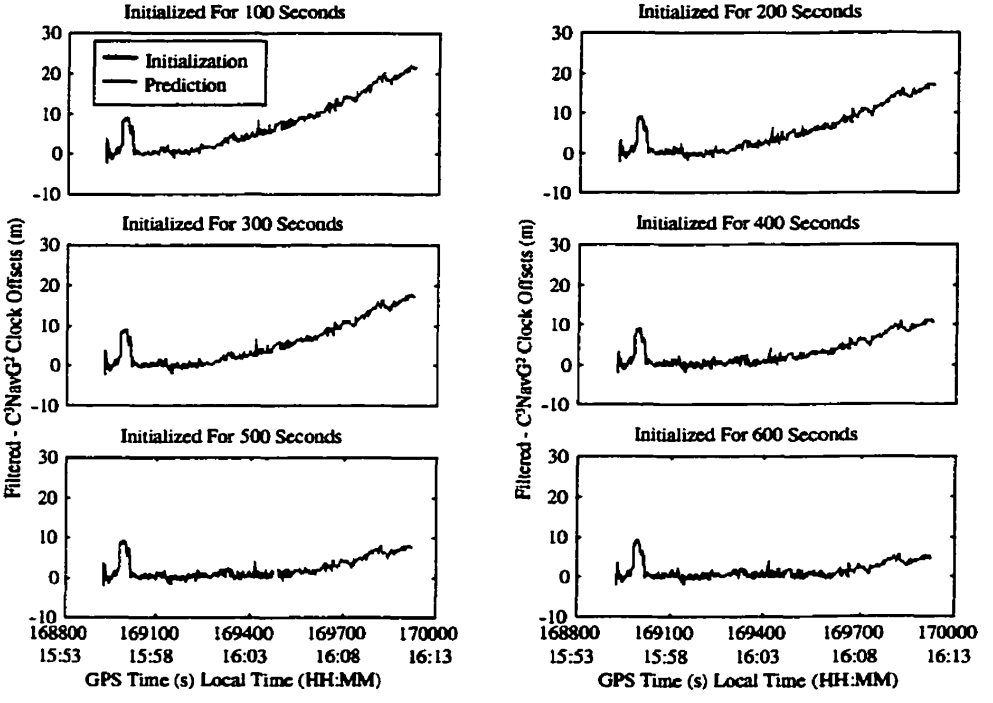

Figure 2.11 Strathcona Test 2 - Kalman Filtered Clock Prediction Errors

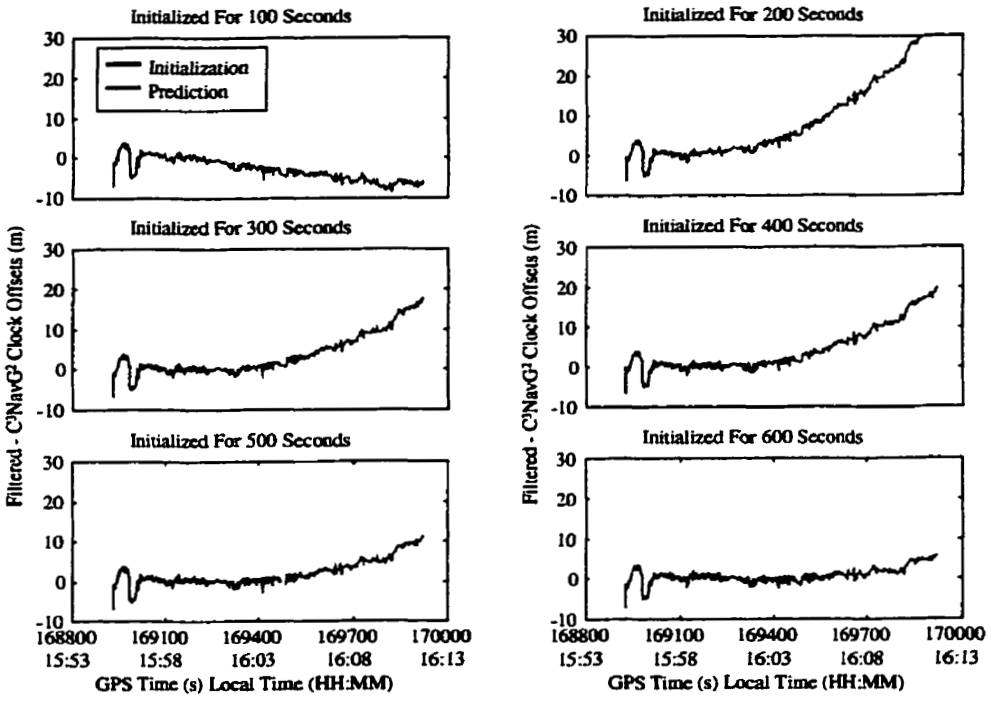

Figure 2.12 Strathcona Test 2 - LS Clock Prediction Errors

3 DEAD RECKONING

Dead reckoning is the use of sensors that provide relative positioning information (Bullock 1995). GPS, on the other hand, gives an absolute position fix in a known coordinate frame. Most dead reckoning sensors can be grouped into motion sensing and heading sensing categories. The barometer described previously can also be considered a dead reckoning sensor since pressure changes can be used to estimate changes in height. DR (Dead Reckoning) and GPS have excellent synergy since GPS provides an absolute position that is of moderate accuracy, and DR provides highly accurate changes in position, but needs absolute reference from which to apply the changes. It is often said that GPS provides the low frequency position information and DR provides the high frequency information in an integrated position solution. DR is not subject to signal masking or outages, but often requires strong external absolute calibration, a great deal of redundancy, or the use of very stable and expensive sensors to function effectively. A survey of several DR sensors, covering both motion sensing and heading sensing categories, was done in order to compare the performance and cost of each of the sensors. An understanding of the operation of each of the sensors gives insight into their advantages and limitations.

3.1 GYRO

A gyro (gyroscope) is an instrument used to measure the rate of rotation or integrated heading change of a platform (Savage 1978). A single gyro measures rotation on a single plane, but a triad of gyros is often mounted orthogonally in a single enclosure to monitor the three possible rotations in 3-D space. Quite a few different types of gyros are available, ranging greatly in price and stability. Gyros are classified into gimbaled or strapdown varieties, with gimbaled gyros maintaining a fixed orientation in an inertial frame (relative to distant galaxies). Low cost (or potentially low cost) gyro alternatives, which are all of the strapdown variety, will be presented with examples. Table 3.1 summarizes several different low cost gyro technologies for cost and accuracy.

Table 3.1 Comparison of low cost Gyro Technologies

Gyro Type	Principle of Operation	Cost ($)	Stability (°/h)
Rotating	Conservation of Angular Momentum	10-1000	1-100+
Fiber Optic	Sagnac Effect	50-1000	5-100
Vibrating Piezoelectric	Coriolis Effect	10-200	50-100+

Strapdown gyros measure rotation on a fixed plane with respect to the vehicle, which is not generally on the plane orthogonal to the gravity vector. Therefore, they do not sense the entire rotation in heading, but also sense rotations in pitch and roll. This situation is presented graphically in Figure 3.1. The formula that describes the relationship between the measured rotation rate and the desired rate of heading change is given in Eq. 3.1 (St. Lawrence 1993).

$$\omega_G = \omega_\psi \cos\xi + \omega_\varphi \cos\beta \sin\xi + \omega_\theta \sin\beta \sin\xi \tag{3.1}$$

Rearranged for the heading rate, the result is

$$\omega_\psi = \frac{\omega_G}{\cos\xi} - \frac{\cos\xi + \omega_\varphi \cos\beta \sin\xi + \omega_\theta \sin\beta \sin\xi}{\cos\xi} \tag{3.2}$$

where

ω_G is the measured rotation rate

ω_ψ is the desired azimuth rotation rate

ω_θ is the pitch rotation rate

ω_φ is the roll rotation rate

p is the roll angle

r is the pitch angle

β is the horizontal component of the gyro and LL plane misalignment

ξ is the misalignment between the gyro and LL vertical vectors

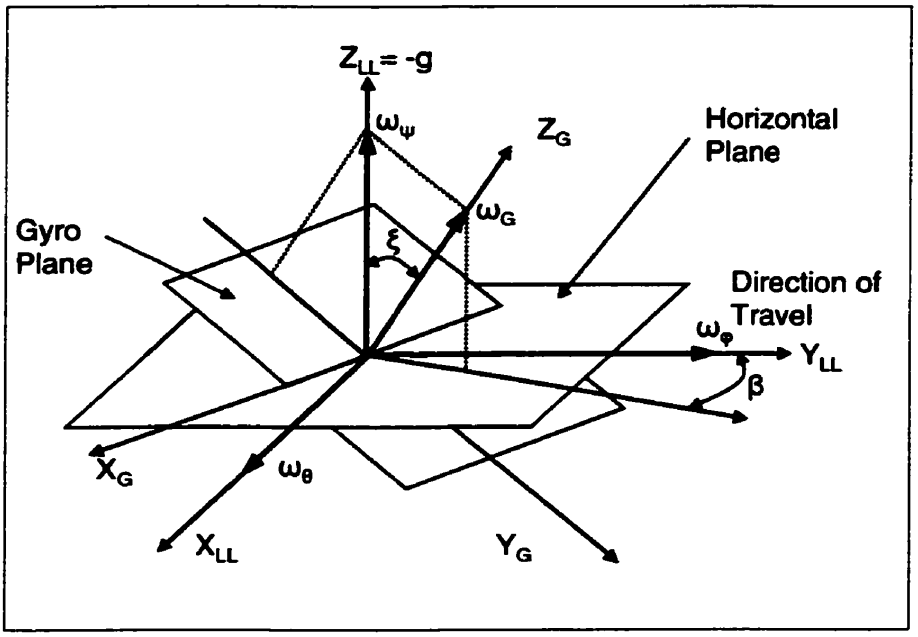

Figure 3.1 Geometry of a Tilted Gyro (St. Lawrence 1993)

Since a low cost navigation system will use a single gyro, there is no way of measuring the pitch and roll rotation rates, so their effects will be ignored. These effects are proportional to the pitch and roll of the vehicle and the rotation rates on their respective axes. The pitch and roll generally remain small and the errors caused by ignoring them tend to cancel out over the long term. Localized errors will be induced into the heading through roll, which occurs when the vehicle corners, and when vertical gradients are changing. The resulting heading errors when integrated should result in a random walk behavior (Harvey 1998). These errors are normally considered an acceptable loss in vehicle navigation applications, since adding complementary gyros would be costly and result in marginal improvement. Ignoring these rotations drops the second term from Equation 3.2. Referring to Figure 3.1, the remaining error can be derived from the direction cosines of the gyro axes on to the local level axes as

$$\sin^2 p + \sin^2 r + \cos^2 \xi = 1 \tag{3.3}$$

$$\beta = \tan^{-1}\left(\frac{p}{r}\right) \tag{3.4}$$

$$\cos\xi = \sqrt{1 - \sin p^2 - \sin r^2} \tag{3.5}$$

and using the small angle approximation sin(x) = x

$$\omega_\psi = \frac{\omega_G}{\sqrt{1-\vartheta^2}} \tag{3.6}$$

with the tilt ϑ defined as

$$\vartheta^2 = p^2 + r^2 \tag{3.7}$$

The resulting error is a function of vehicle pitch and roll, which can be derived from an inclinometer or additional gyros. Vehicle tilt has the effect of decreasing the sensed rate of rotation during a turn. A short numerical example will illustrate the error that is induced by ignoring vehicle tilt. For a sharp, slow corner, a rotation rate of 20°/s is seldom exceeded. If the vehicle was tilted 20° on a steep hill, and the unit indicated the 20°/s rotation, the actual heading rate should be 21.3°/s. This results in a heading error of 1.3° for each second of the turn, i.e. up to 6° error for a 90° turn. At a tilt of 10°, only 1.3° error on the 90° turn would result. It is difficult to equate these heading errors to positional errors since position errors can accumulate or cancel, depending on vehicle dynamics. However, small heading errors can easily lead to large position errors in the cross track direction when coasting on DR during a relatively long GPS outage.

A sample distribution of the tilts experienced in various residential and freeway driving environments is found in Figure 3.2. It can be seen in this figure that vehicle tilt seldom exceeds 10° under normal driving conditions, at which there is a 1.5% heading change error induced by ignoring tilt. This correction can be ignored if it assumed that the vehicle is normally close to level, and GPS outages seldom occur under high tilt conditions. The scale factor on the gyro output would be increased slightly to compensate

for some sort of RMS tilt value. An inclinometer might prove necessary if positions were required inside covered or underground parkades, or to maintain accuracy in areas with a great deal of height variation and signal masking.

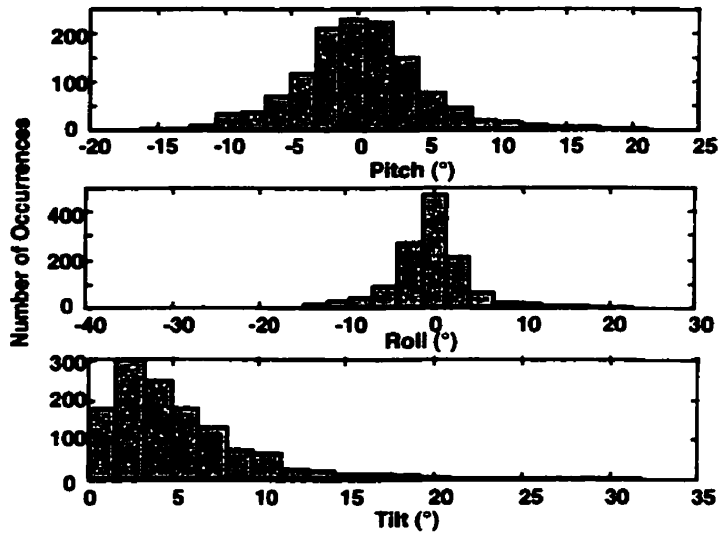

Figure 3.2 Sample Vehicle Tilts From Varied Driving

3.1.1 PIEZOELECTRIC VIBRATING GYRO

Piezoelectric materials exhibit the piezoelectric effect, in which a vibrational motion of the crystals creates an electric potential within the material. The reverse piezoelectric effect is also used, whereby application of an electric field to a PE (piezoelectric) crystal will cause it to vibrate. This process is familiar since it is used in the operation of most wristwatches. Quartz is used in many PE applications, however it tends to have temperature instability and physical limitations that have given rise to recent advances in PE ceramics. Piezoelectic applications usually make use of resonant or harmonic frequencies, which are a function of the size, shape, and dielectric properties of the piece of material. Vibrating gyroscopes are usually designated as MEMS (micro-electromechanical system) sensors, i.e. sensors that couple electrical and mechanical interactions with microchip fabrication methods.

Several manufacturers have taken to producing gyroscopes based on the use of piezoelectric technology, basing their designs on the initial work of Konno (1987). Models are being sold from $5 to $100, depending on how well they are individually tuned, and the accompanying electronics. Two basic types have been constructed, namely the free-free-bar and tuning fork varieties. Both use an excitation signal to drive the oscillation of a PE member, then sense rotation through an output PE generated signal. The tuning fork variety is simply a tuning fork that is made of piezoelectric material. The name free-free-bar comes from the use of a slender bar, with PE ceramics attached, which is fixed at the centre and has both ends free to vibrate.

Piezoelectric gyros are essentially Coriolis sensors. The Coriolis force is a fictitious force exerted on a body when it moves in a rotating reference frame (Motta 2000). It is a fictitious force because, like centrifugal force, it is a by-product of measuring coordinates with respect to a rotating coordinate system as opposed to the acceleration of mass in an inertial frame. It is given with the cross product

$$\bar{F}_c = 2m(\bar{v} \times \bar{\Omega}) \tag{3.8}$$

where

m is the mass of the object

v is the velocity vector of the object

Ω is the angular rotation rate vector

The excited vibration of the bar or fork creates an oscillating velocity vector. If this system is rotated around the sensitive axis, an oscillating force will be induced, which causes vibration of the piezoelectric crystal. This vibration can be sensed as a varying voltage, which is then processed into an output signal. The operation of the Murata Gyrostar free-free-bar implementation, which was the gyro used for this project, is described below. An example of the tuning fork type can be found in the Systron Donner GyroChip series (Systron Donner 1998).

The Murata Gyrostar piezoelectric vibrating gyro uses an equilateral triangular bar composed of elinvar (elastic invariable metal) with three PE ceramics attached to the sides. The ceramics are made from Murata's patented Piezotite materials, which claim better temperature stability electro-mechanical coupling than other PE materials (Murata 1998). The bar is excited into motion from two sides, and the third side is used to drive a feedback loop, which controls the bar oscillation. The drive ceramics are also used to detect the motion of the bar. A diagram of the bar is shown in Figure 3.3.

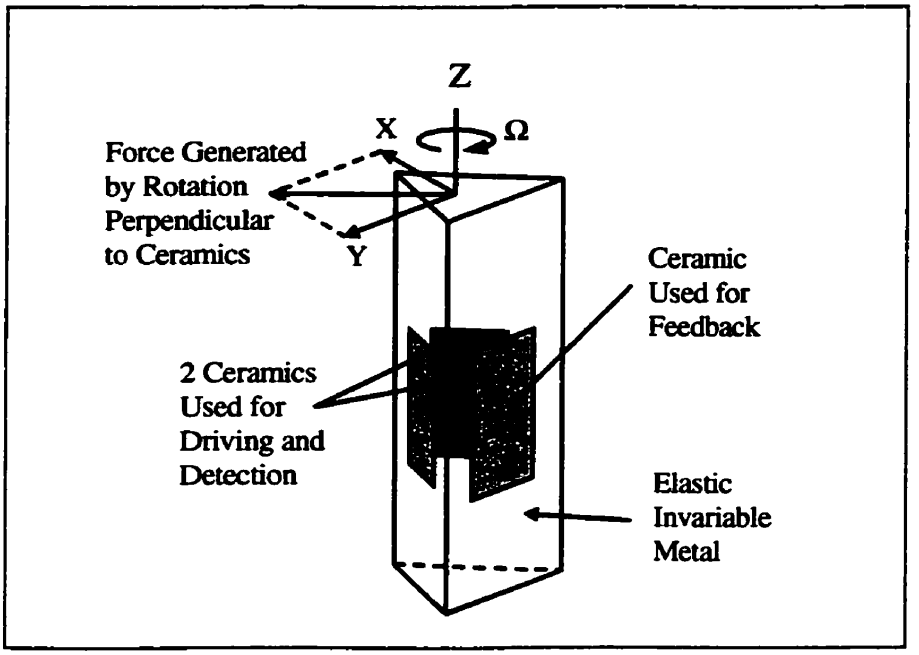

Figure 3.3 Gyrostar Free-Free-Bar and Ceramics (Hayashi 1996)

The bars can be very easily tuned to match the driving (resonant) and detecting frequency of both detectors compared to the rectangular bars previously developed. This is because adjustments of the edges of the prism only affect one flexural mode of vibration. Rotation detection is also innovative due to the process of differencing the output from the two detectors. The phase angles and magnitudes of the vibration are compared to give a direction and magnitude of rotation. The effect of differencing of the two detectors is shown in Figure 3.4. The specifications of the automotive version of the sensor (ENV-05D series) are found in Table 3.2.

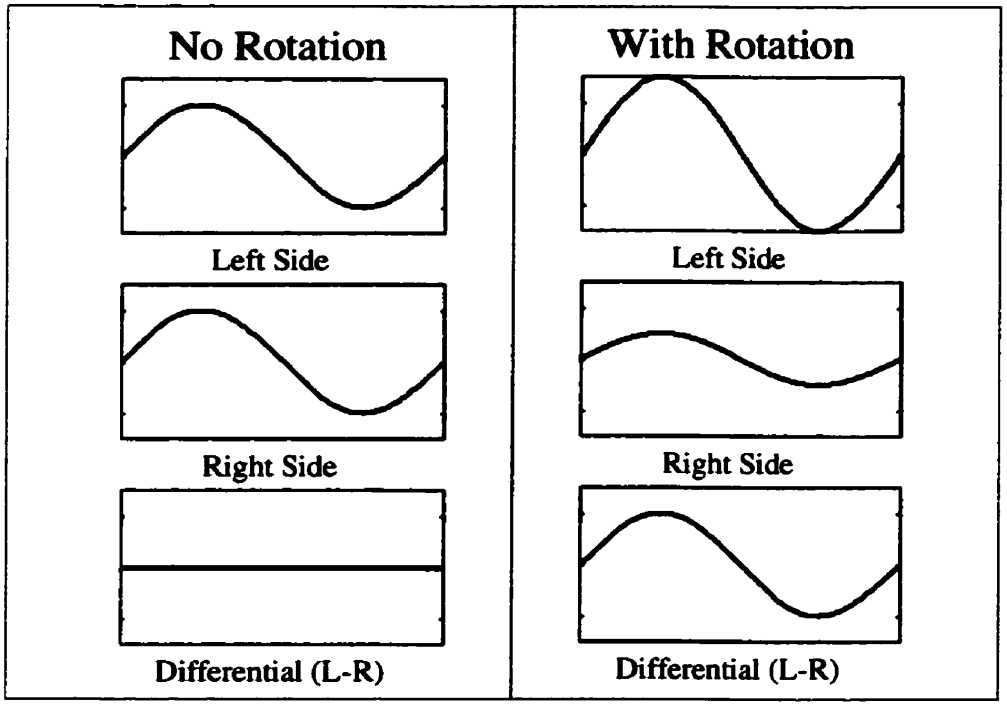

Figure 3.4 Differencing of the Left and Right Detection Signals (Nakamura 1990)

Table 3.2 Selected Murata Gyrostar Specifications (Murata 1999)

Characteristic	Min	Std	Max	Units
Supply Voltage	4.5	5.0	5.5	VDC
Current @ 5 VDC	-	-	17	mA
Max Angular Rate	-80	-	+80	°/s
Output at Zero Rate	2.2	2.5	2.8	VDC
Scale Factor	19.3	22.2	25.1	mV/°/s
Asymmetry	-	-	3	°/s
Temperature Scale	-	-	±10	%FS
Temperature Drift	-	-	9	°/s
Operating Temperature	-30	-	80	°C
Noise Level (7 kHz)	-	-	20	mV
Dimensions	18 ×30×41 mm			

This gyro was selected for its low cost (approximately $15 in large quantities), small size, and good performance relative to the others in its class. The voltage output from the gyro is about 2.5 V when stationary and extends about 1.5 V in each direction under 80°/s of rotation.

The Murata gyro has instability in both the zero-rotation output voltage and the scale mainly due to temperature variations. Since a vehicle can not be kept in a controlled environment, and the cost of using an ovenized enclosure is restrictive, regular calibration of these values is often required. The gyro has also demonstrated limited endurance to shock and some susceptibility to vibration (Geier 1998). Vibration in vehicles varies greatly, both in magnitude and frequency, depending on engine speed, temperature, and numerous other factors. Accelerations reach 0.7 g's over spectral frequencies up to 10 Hz (Skaloud 1995).

Difficulties also arise in trying to calibrate both the zero-rotation offset and scale factor simultaneously without a precise reference. The zero-rotation offset is the primary cause of heading drift and is easier to calibrate. The scale factor can only be calibrated by analyzing a series or history of turns to see if there is a tendency to overestimate or underestimate rotation. Not enough GPS heading data was available to attempt this type of calibration during the short tests presented here, however the scale factor does not seem to vary nearly as much as the zero-rotation offset. The scale factor was determined by scaling the integrated output of the gyro over circles (10) on a level surface. The resulting value was used throughout the testing.

3.1.2 RING LASER GYRO

Ring laser gyro technology is described since several manufacturers have experimented with ways of implementing this technology in a fiber optic gyro. Combining the performance of the ring laser gyro with the cost of a fiber optic gyro would be very advantageous. The RLG (Ring Laser Gyro) bases its operation on the Sagnac effect. The Sagnac effect dictates that a circuit will have different optical path lengths if it is rotated. This is a relativistic property of light, which is due to its propagation in a rotating inertial frame. RLGs operate by using mirrors to set up a closed loop for the path of a laser. Beams of light are directed both clockwise and counter-clockwise around the closed path, and position of one or more mirrors is then tuned so that a perfect standing wave pattern is created which remains stationary in inertial space. The mirrors are dielectric mirrors

that will reflect only the frequency of the laser. A fraction of the light is allowed to pass through one of the mirrors to photodiodes, which will measure the interference pattern created by rotation. They are typically very expensive, i.e. $10,000 and up, since they require a high voltage power source, a high power laser, and very fine machining of the laser cavity and mirrors. A typical RLG assembly is shown in Figure 3.5 (Savage 1978).

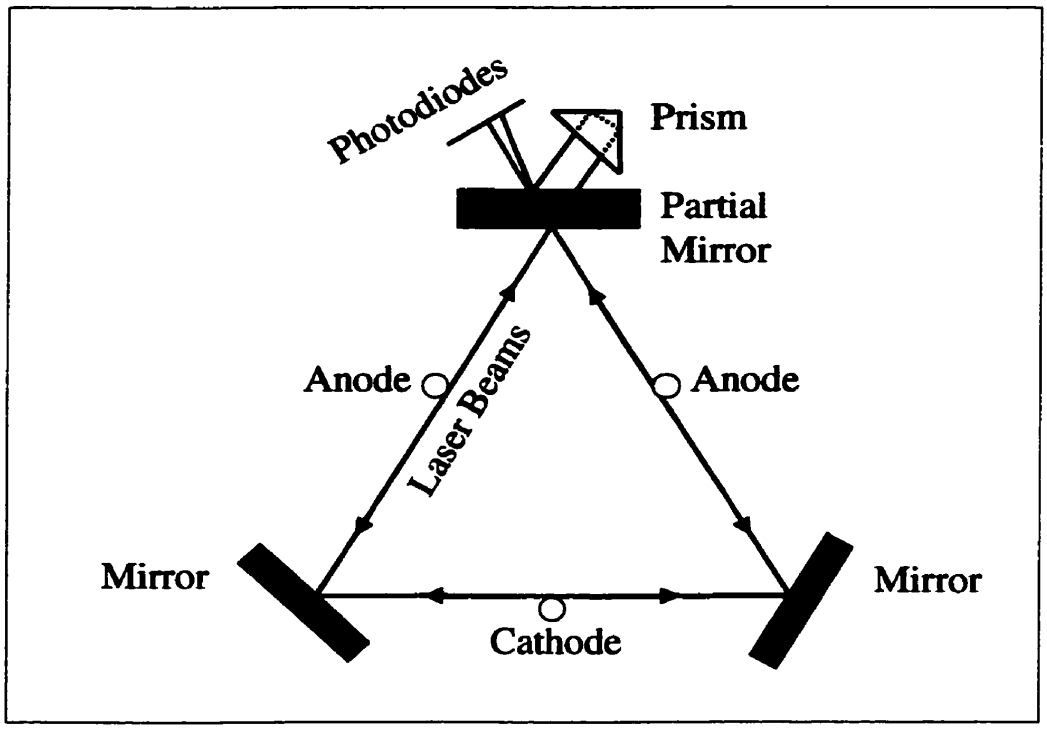

Figure 3.5 Ring Laser Gyro (Mark et al. 1991)

The figure shows an interferometric method of measuring the rotation. The recombination of the two beams by the prism would create a fringe pattern at the photodiodes that would not move if the gyro was kept stationary, but if it was rotated, the fringe pattern would travel across the photodiodes. The photodiodes could be used to count minima or maxima, and an integrated heading change is measured. This procedure is analogous to mixing a signal with a Doppler shifted return signal and counting the beats for an IDC (Integrated Doppler Count) range. An RLG is therefore a true unreferenced heading sensor, not a rate gyro. Another method of measuring the output is

simply counting nulls in the standing wave pattern with two orthogonal sensors, a procedure that simply accumulates the path length differences (Mark et al. 1991). The resolution can be estimated to be about

$$\theta_r = \frac{Path\ Length}{\lambda} \tag{3.9}$$

where λ is the wavelength of the laser, usually about 0.63 micrometres. Most RLGs have triangular legs 10-15 cm in length, and have a resolution of about two arcseconds. One of the largest errors for this type of gyro is called lock-in. Lock-in is the result of backscatter (caused by imperfections in the mirror) directing light back from the direction it came, effectively converting some of the beam between clockwise and counter-clockwise rotation. This effect is negligible at higher rotation rates, but at low rotation rates, the frequencies of the two beams tend to lock-in to the same frequency. This error starts as a nonlinearity and progresses to full lock-in at about 100°/hr (Rüger 1982), as shown in Figure 3.6. Several means have been developed to avoid lock-in, the most effective is a mechanical method, but the use of mechanical parts decreases the reliability and life of the unit. The entire laser cavity can be rotationally vibrated (dithered) using a piezoelectric mechanism that generally increases the noise, but removes the non-linearity across the lock-in region. A block control diagram including the dithering mechanism is shown in Figure 3.7. Other methods have used electrical means to cause continuous variable phase shifting of the two beams. Using either method, the induced rotation must be removed from the measured signal before it is used to generate output.

Advantages of laser technology over any of the mechanical implementations include an extremely high dynamic range, excellent reliability since there are no mechanical parts to fail, and none of the dynamics (acceleration) related errors which are found in mechanical implementations. However, the laser cavity must be precisely machined and sealed, and an expensive high voltage laser is needed to construct an RLG.

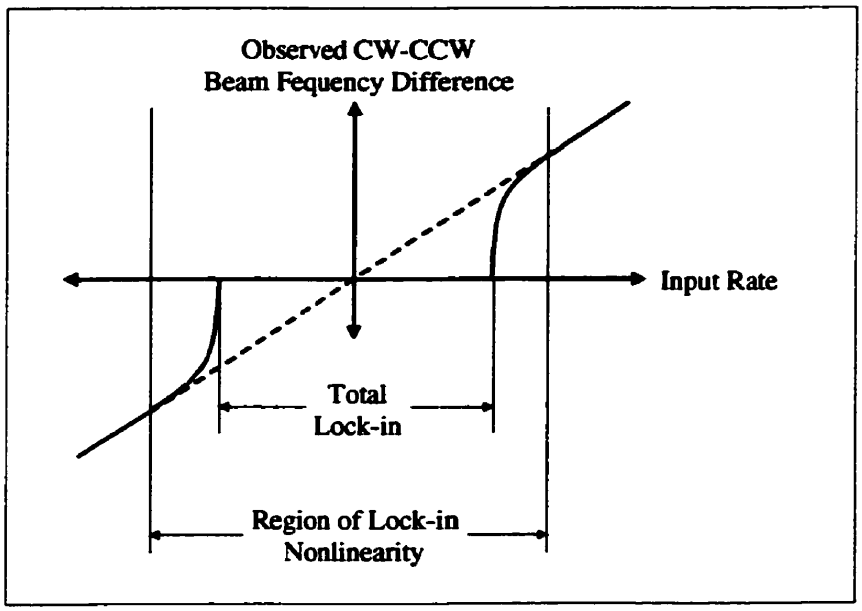

Figure 3.6 Lock-in Phenomenon (Rüger 1982)

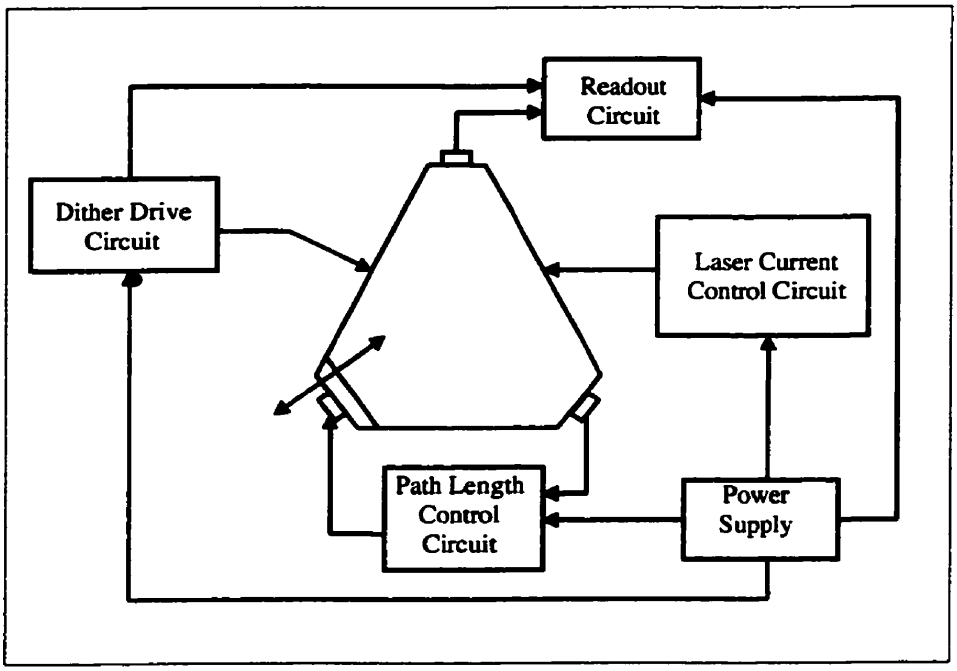

Figure 3.7 RLG Block Control Diagram Including Dithering (Rüger 1982)

3.1.3 FIBER OPTIC GYRO

Fiber optic gyros have grown in popularity due to the rapid growth of fiber optic technology and improvements in the cost and accuracy of optical electronics. FOG (Fiber Optic Gyro) technology developed from RLG systems, which were previously described. The FOG uses a similar principle, but does not require the high voltage laser and laser cavity construction that are the most cost restrictive features of the RLG. Attempts have been made to build resonant FOGs that operate like RLGs (Mark et al. 1991), but technological limitations have so far restricted their development. For this reason, only the most common implementation, the IFOG (Interferometric FOG), is described. Costs for navigation quality FOGs have dropped steadily over the last decade, but it is questionable whether they will get down to the level required for bulk automotive applications in the near future. Size is also an issue, since larger packages (increased radius) and more fiber optic cable provide better accuracy, while many other gyros are being made smaller with no loss of performance.

In a FOG, the optical path is a fiber optic cable that is normally coiled many times around a cylinder. The beam is emitted from a laser diode or SLD (superluminescent diode), passed through a polarizer, and split by a coupler. The light going to one side of the coil passes immediately through a phase modulator, while the light moving the opposite direction will not pass the phase modulator until it has almost completed the path and is about to be recombined by the coupler. Various phase modulation implementations have been used, including alternating orthogonal phase shifts (± 90°), Serrodyne step functions, and sinusoidal modulation (Bennett et al. 1998). The light is combined at the coupler, cleaned again by the polarizer, and passed on to the photodetector. KVH has also developed a system that uses the photodetector that is already part of the laser diode assembly to minimize fabrication costs. The basic components of a FOG are pictured in Figure 3.8. Very often, the couplers, polarizer, and phase modulator are built into a single MIOC (Multi-function Integrated Optic Circuit).

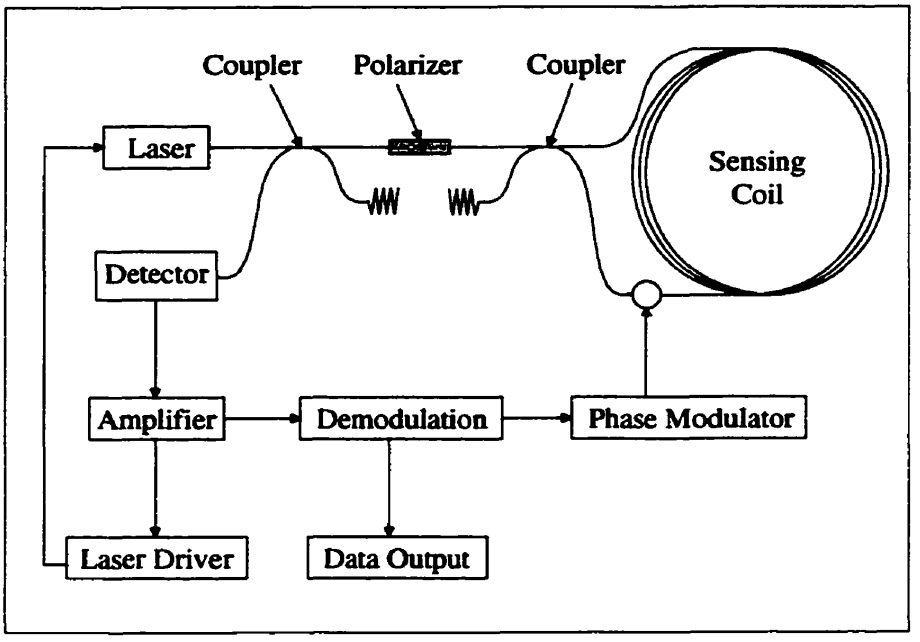

Figure 3.8 FOG Optical and Electical Circuits (Bennett et al. 1990)

The optical path length with N loops can be given by

$$L = 2N\pi \tag{3.10}$$

The Sagnac phase shift, i.e. the difference in optical path lengths in a rotating circuit, can be given by

$$\Delta L = \frac{4N\pi r^2\Omega}{c^2 - r^2\Omega^2} \tag{3.11}$$

where

r is the effective radius of the circuit

Ω is the rate of angular rotation

c is the speed of light in the medium

and since $c^2 \gg r^2\Omega^2$

$$\Delta L = \frac{4\pi r^2 \Omega}{c^2} \tag{3.12}$$

The induced Sagnac phase change on monochromatic light of wavelength λ is therefore

$$\Psi_S = \frac{8N\pi^2 r^2}{c\lambda}\Omega = K_S\Omega \tag{3.13}$$

where K_S is called the Sagnac scale factor.

A photodetector is used to measure the incident light intensity variations, which would be related to the Sagnac phase change, and thereby the angular rate of rotation. However, a sensor operating in this manner would have zero gain at low rotation rates and a high degree of nonlinearity. To overcome these limitations, the phase modulator creates bias rotations that can be used to accurately estimate the rotation rate. An example of how the modulation permits this is found in Figure 3.9 (Matthews 1990).

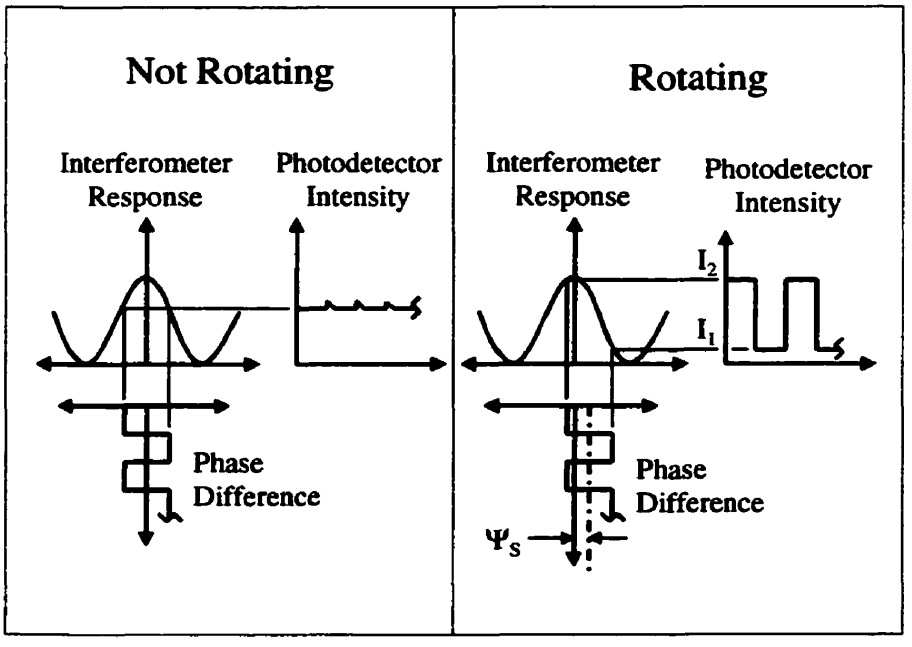

Figure 3.9 Use of Phase Modulation to Detect Roatation in a FOG (Matthews 1990)

The figure demonstrates a binary modulation scheme, whereby the phase is varied between ± 90° by the phase modulator, seen as the square wave in the phase difference. The interferometer response is sinusoidal, but this relationship is now well known from other control problems. The phase changes are usually timed to occur once for each transit of light through the coil. This type of gyro is referred to as an open-loop gyro, and has limited range and high scale factor nonlinearity. The rotation rate for this type of gyro can be determined using

$$\Omega = K_S(I_1 - I_2) \tag{3.14}$$

The closed loop formulation adds feedback to the oscillator from the discriminator to drive the observed Sagnac effect to zero, a process similar to VCXO (voltage controlled oscillator) control in a phase locked loop. Closed loop operation allows extremely high dynamics and good linearity. A FOG can be made to behave as an unreferenced heading sensor, or rate integrating gyro, by operating the measurement, update and control functions at an interval less than or equal to the transit time of light through the coil (Mark et al. 1991). There is also a great deal of flexibility in the quality and cost of components, the length of coil, and signal processing techniques that allow for construction of FOGs for applications from navigation to inertial units. Very high resolutions are possible using a FOG, e.g. using a 12-bit A/D conversion, a resolution of 0.004 arcseconds is possible for 200 metres of fiber around a cylinder of 10 cm diameter (Matthews 1990).

3.1.4 SPINNING GYROS

Many different types of spinning gyros are available, but the physical principles employed to measure rotation are similar. A spinning gyro operates on the principle of conservation of angular momentum. Angular momentum is a vector quantity that will not change direction or magnitude unless a torque is applied. Forces that are applied without causing a torque will not change the angular momentum, i.e. linear acceleration. The application of a torque on the plane of rotation will cause the gyro to simply increase or

decrease speed. A torque perpendicular to the plane of rotation will not change the magnitude of the angular momentum, but will change the direction of the angular momentum vector. A simple example is that of a rolling bicycle wheel. If a torque is applied in the plane of rotation, the angular momentum and rate of rotation will change. This occurs when you push the top of the wheel to make it roll faster. If a torque is applied perpendicular to the rotation, the direction of the angular momentum will change via the RHR (Right Hand Rule). A push to the top of the tire from the side will cause the tire to change direction and turn toward the push; a constant force applied in this fashion would cause the tire to roll out a circular path. The change in direction of the angular momentum is called precession, and a graphical representation of it can be found Figure 3.10.

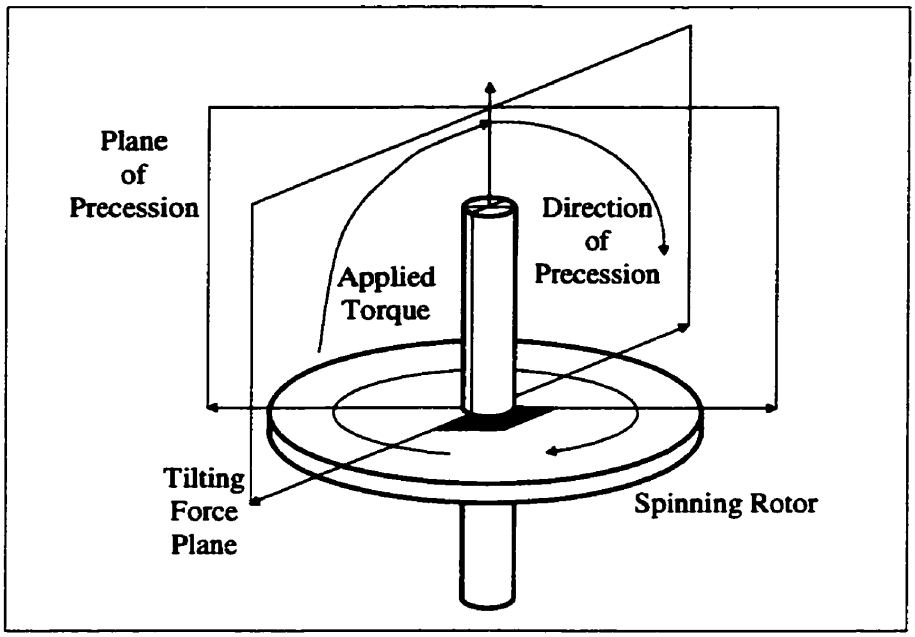

Figure 3.10 Floating Gyro

The phenomenon is clearly demonstrated through examining the riding of a bicycle. A rider finds it very difficult to maintain balance at low speeds, and it is nearly impossible to keep balance when stopped. Balance is maintained at high speeds due to the gyroscopic nature of the wheels. Any error in balance causes the bike to lean, which

tends to make the bike turn, which rights the rider. This phenomenon happens continuously and instantaneously, so the rider does not sense it.

A gimbaled gyroscope resists rotation of its axis, but rotation does occur while a torque is applied. The torque does not cause momentum of the axis, however, since the energy is dissipated to friction by the gimbals. A gimbaled gyro is pictured in Figure 3.11.

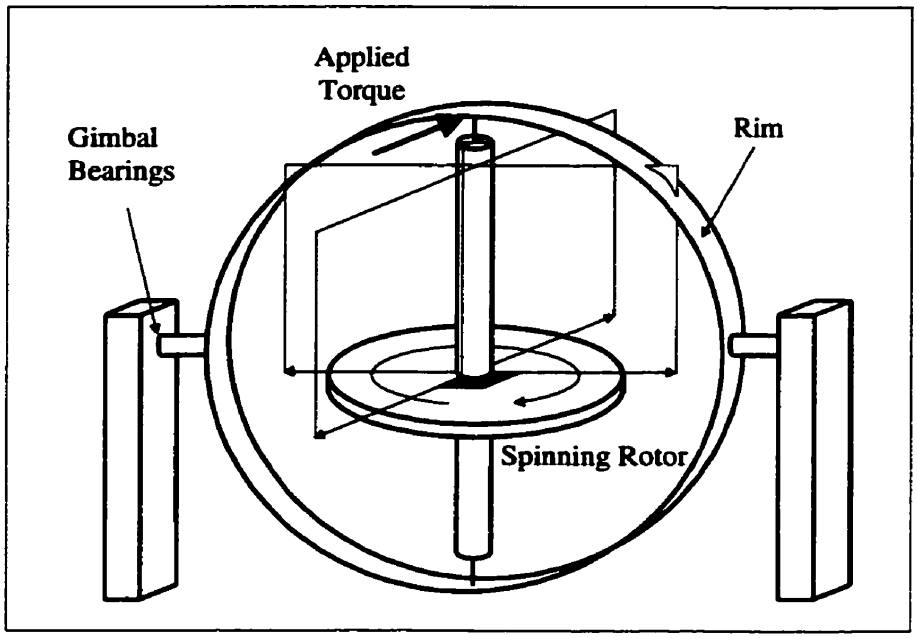

Figure 3.11 Gimbaled Gyro

The rotor will spin at the rate where losses (mostly frictional) account for the entire power consumption of the motor, since maintaining the angular momentum without friction requires no energy. This high speed operation is critical to gyro performance, since the greater the angular momentum, the more resistance there is to perturbing torques. If the gyro is spinning too slowly, some of the torque will tilt the gyro. This can be seen using commonly found gyroscope toys or tops, which will remain vertical until they have slowed sufficiently. After which time, they noticeably tip and precess out a large arc. A spinning gyro uses these principles to monitor the amount of rotation that has occurred or sense the torque being experienced by the rotor. Some gyros operate by returning the precessed gyro to its original orientation and keeping track of the total

amount of correction applied to each axis. Others use the deformation or movement of the rotor to estimate the rotation rate. The greatest limitation of most spinning gyros is the limited lifespan due to friction on the moving parts. Several extremely low cost alternatives are presented below.

An innovative MEMS gyro has been developed by researchers at Sheffield University, however it has not yet been commercially produced (Shearwood 1997). Their design uses electromagnetic induction to levitate a rotor and spin it at high speeds (Williams et al. 1997). The rotor will tilt orthogonally to an applied rotation, as described above. The angle of tilt is then measured by variable capacitance from the stator coils to the rotor. The stator coils are used in a cyclic manner to perform the three functions of levitation, rotation, and torque measurement. A photograph of the micromotor with and without the rotor is found in Figure 3.12.

Initial testing indicates that the present system, which is excited by a 10 kHz AC current, can spin the rotor at 1040 rpm in air. Given a four phase motor, the rotor could be made to spin two orders of magnitude faster, greatly increasing the gyro's sensitivity. There is minimal sensitivity of the parts to temperature; only the accompanying electronics would be affected, so the temperature sensitivity should be better than that of a vibrating gyroscope. If a feedback loop is used to keep the rotor level even while a turn is torquing the rotor, the dynamic range and linearity can be improved, as well.

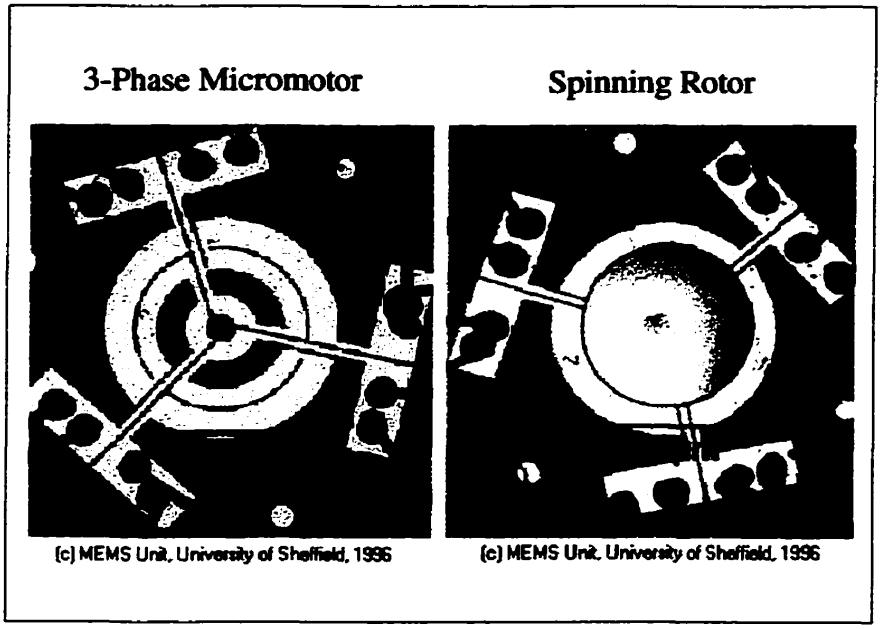

Figure 3.12 MEMS Spinning Micromotor Gyro (Shearwood 1997)

Another low cost spinning gyro, which is commercially available, is the Etak turn sensor (Phillips 1993). It is similar in operation to the tuned gyros used in many INS (Inertial Navigation System) units. The sensor has been made very inexpensive by the use of a simple low-cost motor and novel signal processing to avoid costly analog circuitry. The Etak gyro is shown Figure 3.13.

The parts on the exploded view are

1. Housing
2. Motor
3. Flexible Disk
4. Ground Plate

5,6,7. Capacitive Plates

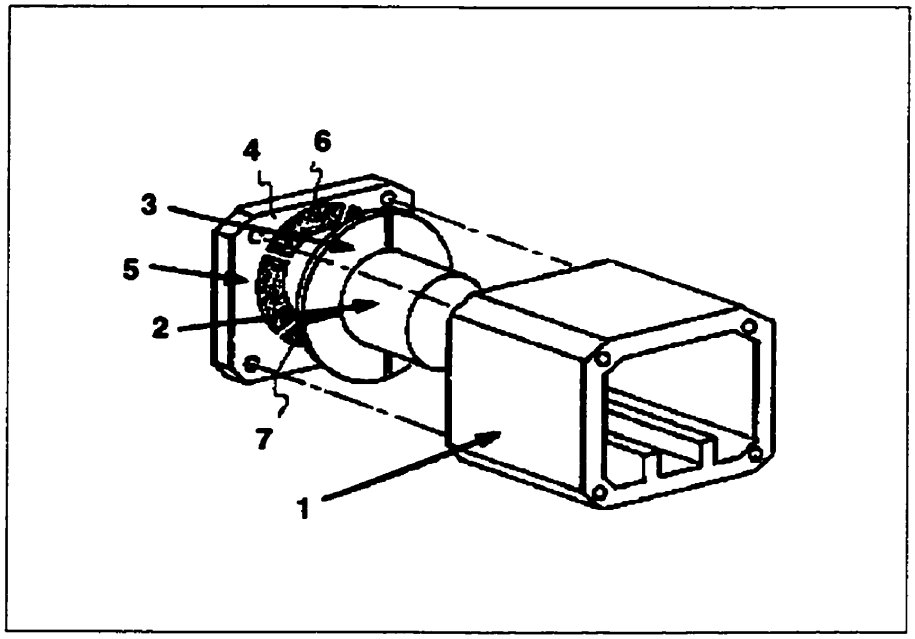

Figure 3.13 Etak Gyro Exploded View (Phillips et al. 1989)

The beryllium-copper flexible disk is attached to a spring, which is in turn attached to the motor shaft. The spring acts as a damper for vibrations, acceleration, and motor noise. Rotation on either of the two sensitive axes of the gyro causes the flexible disk to twist into a plane oblique to the capacitive sensing plates. The same is true of the tuned rotor gyro, but in the tuned rotor gyro a torque is applied to the rotor to keep it perpendicular to the shaft and improve linearity. The sensing circuit is shown in Figure 3.14. A balanced current source is used to charge opposite capacitive plates, while a crystal oscillator discharges the capacitors at a fixed rate, creating a sawtooth signal. The difference of the two voltages is passed through a differential amplifier and into an ASIC (Application Specific Integrated Circuit) that uses this signal to control a digital pulse train generation. The output is about 39 kHz when there is no rotation, and varies down to zero and up to 78 kHz under full scale rotation. Etak estimated that the sensor could be manufactured for as little as $10 in large quantities. Some specifications are found in Table 3.3.

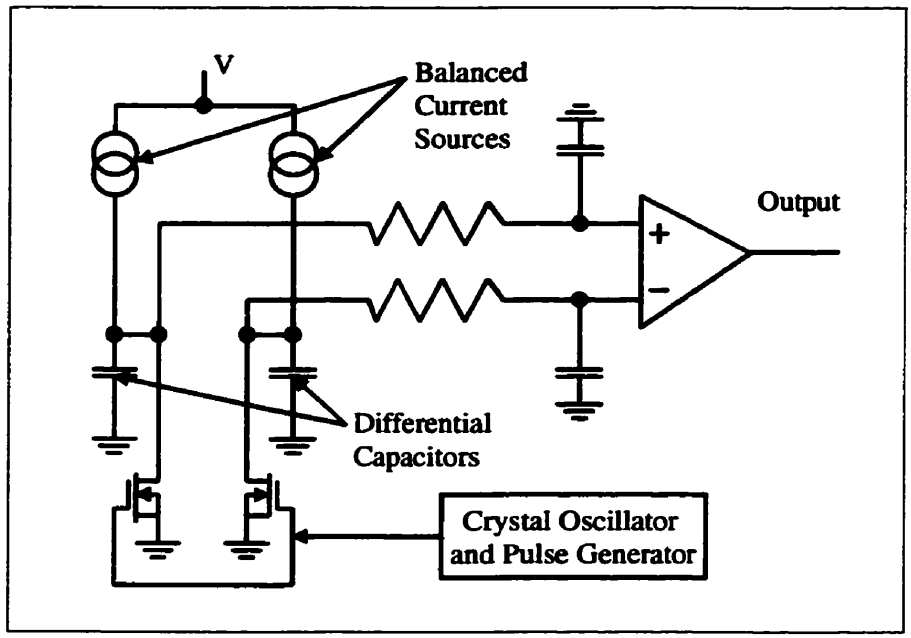

Figure 3.14 Etak Gyro Capacitance Measurement Circuit (Phillips et al. 1993)

Table 3.3 Etak Gyro Specifications (Phillips et al. 1989)

Characteristic	Specification	Units
Supply Voltage	5	V
Supply Current	50	mA
Operating Temperature	-20 to +80	°C
Stability	1	°/min
Gain Temperature Drift	0.1	%/°C
Linearity	1	% at 60°/s
Gain g-Sensitivity	6	% at 1 g
Offset g-Sensitivity	10	°/min at 1 g
Motor Life	>5000	hours

3.2 COMPASS

The magnetic compass is the only low cost absolute heading reference presently available for automotive markets. Other absolute references, such as north-seeking gyros, are far too expensive. The serious drawback for the use of a magnetic compass on a vehicle is the hostile magnetic environment of an automobile. Several manufacturers, including KVH, have built gyro stabilized compasses which blend the short term relative accuracy of gyro with the absolute sensing of a compass for an improved heading sensor.

A magnetic compass senses the magnetic field of the Earth on two or three orthogonal sensors, sometimes in conjunction with a biaxial inclinometer. Since this field should point directly North, some method can be used to estimate the heading relative to the magnetic North pole. There is a varying declination between the magnetic and geodetic North poles, but models can easily generate this difference to better than one degree. The magnetic sensors are usually flux-gate sensors; one of which is pictured in Figure 3.15.

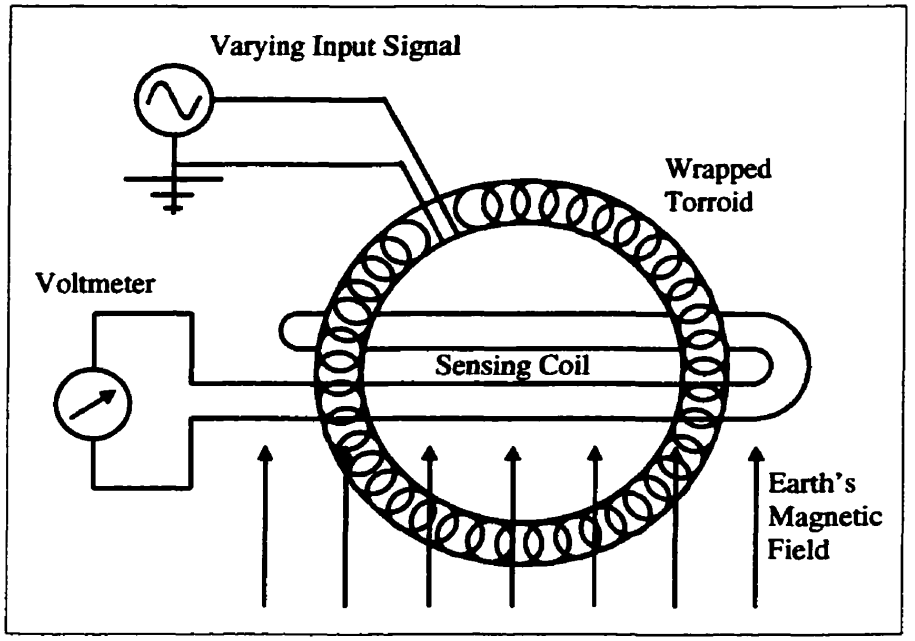

Figure 3.15 Toroidal Wound Fluxgate Compass Sensor (Ganssle 1989)

The operation of a fluxgate is based on Faraday's law, which states that a current (or voltage) is created in a loop in the presence of a changing magnetic field (Ganssle 1989). A fluxgate is composed of a saturating magnetic core, with a drive winding and a pair of sense windings on it (only one is shown in the figure). The drive winding is wrapped around the core, which is normally a toroid. These sense windings are often wound flat on the outside of the core and are arranged at precisely 90° to each other. When not energized, a fluxgate's permeability 'draws in' the Earth's magnetic field (Ripka 1992). When energized, the core saturates and ceases to be magnetic. As this switching occurs (hence the name fluxgate), the Earth magnetic field is drawn into or released from the

core, resulting in a small induced voltage that is proportional to the strength and direction of the external field. The varying permeability of the core can be seen in Figure 3.16.

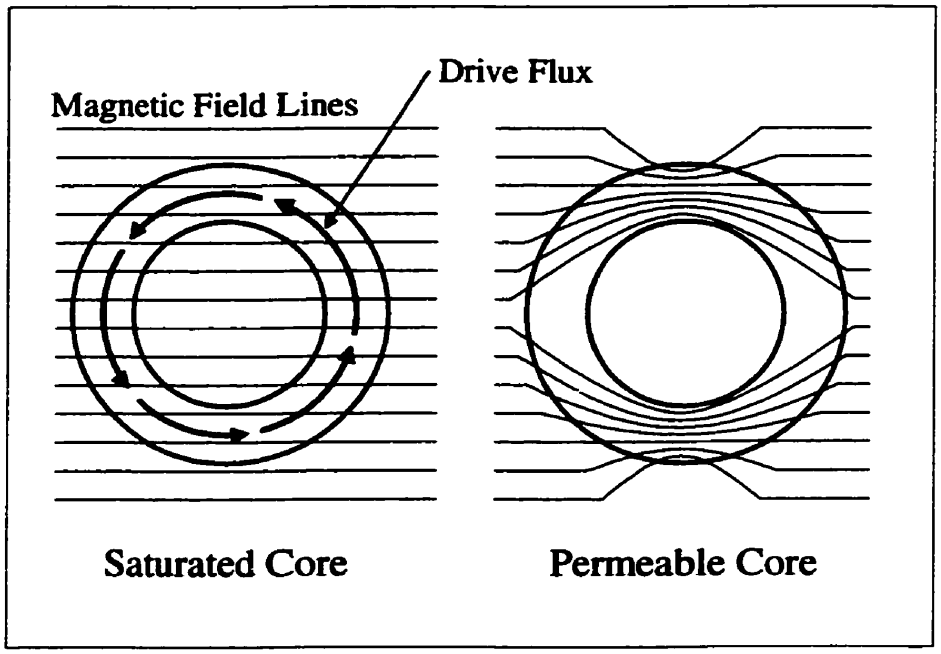

Figure 3.16 Magnetic Field Around A Fluxgate (Phillips 1993)

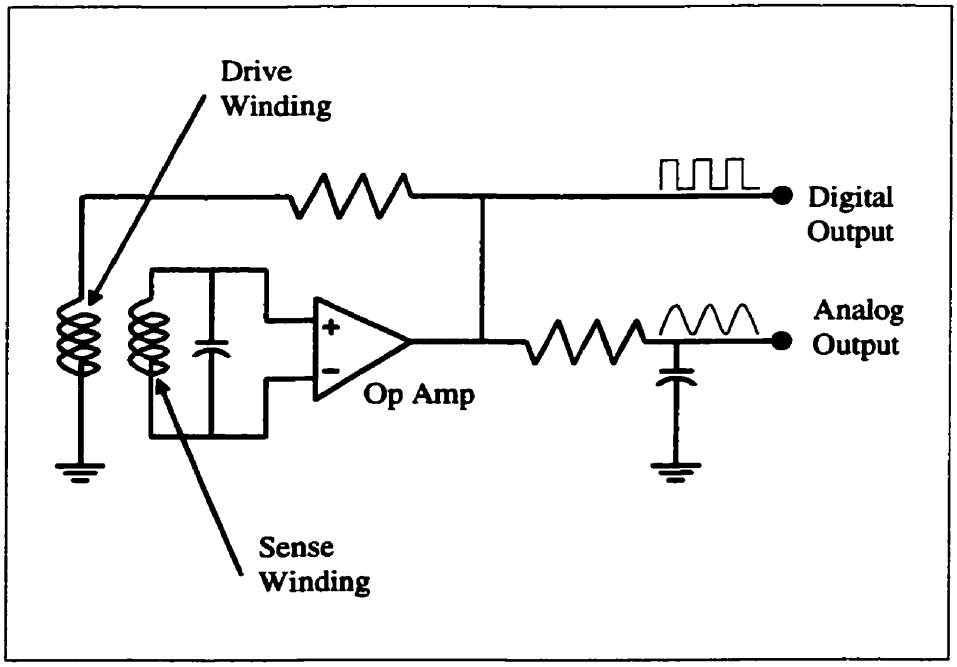

Figure 3.17 Resonant Fluxgate Magnetometer (Lancaster 1996)

The signal processor is often configured as in Figure 3.17 (Lancaster 1996). This configuration is referred to as a resonant fluxgate magnetometer. The duty cycle (amount of time the output digital signal is high) increases as the externally applied magnetic field increases. If the duty cycles times are summed, the sum can be converted to a magnetic flux measurement.

The hostile environment is the result of the magnetic field created by the car, itself. A car is made of conductive material through which current is allowed to flow freely to the battery ground. The result of current passing through a conductor is a magnetic field. The iron in the vehicle will also have some randomly occurring magnetic alignment. Speakers and the alternator in the vehicle also employ magnets. It is also important that the compass not be too heavily shielded by conductive material. A compass closed inside a conductive box will not sense any external magnetic field, including the Earth's. Ultimately, the compass is measuring the sum of the magnetic field of the vehicle, nearby vehicles and metallic objects, and the Earth.

It should be possible to calibrate out the magnetic field of the car since it is constant and keeps its orientation with the vehicle, while the Earth's magnetic field always points North. However, disturbances such as opening a door, transporting metallic objects, or turning on equipment such as a stereo and rear-window defrost can change the car's magnetic field. Temporary external disturbances such as nearby vehicles can not be calibrated out. Since the total magnetic field strength should not vary with orientation, magnetic disturbances are often identified by checking that total measured field strength variation does not exceed a threshold. The magnetic field of the Earth has a strength of about 50 μT, but this value, or certainly its horizontal component, can easily be exceeded by the magnetic field of the vehicle.

The horizontal components of the magnetic field are needed to compute the heading relative to the North magnetic pole. If the sensor is tilted with respect to the local level plane, some amount of the vertical component will be sensed by the horizontal axis sensors. For this reason, a two-axis sensing compass can not properly determine the

heading of the vehicle if the vehicle is tilted. A sensor with three magnetic sensing axes can determine the orientation of its axes only if the magnitude of the magnetic field is known or assumed. The addition of an inclinometer allows for full determination of the strength and orientation of the magnetic field with respect to the sensing apparatus. A method developed by Plessey Overseas for handling a complement of three magnetic sensing coils and a biaxial inclinometer, similar to the Precision Navigation TCM-2 used for this project, is presented below (Marchent and Foster 1985).

If a three-axes magnetometer is operated without the presence of spurious magnetic fields, the measurement locus of the magnetic field of the Earth would appear as a sphere by allowing all values of pitch, roll, and heading for the sensor. In the case of the sensor being mounted on a vehicle, the locus becomes an ellipsoid that is shifted from the origin and has arbitrary orientation and size, depending on the permanent magnetic field of the vehicle. The sphere and ellipsoid are shown in Figure 3.18. Point A and C refer to the same heading, pitch, and roll, so a transformation is desired which restores point C to point A where given pitch and roll from the inclinometer, computation of the heading is trivial.

The equation for the arbitrary ellipsoid can be given by

$$ax^2 + by^2 + cz^2 + 2fyz + 2gzx + 2hxy + 2ux - 2vy + 2wz = d^2 \qquad (3.15)$$

where

x is the along-track sensor measurement

y is the transverse sensor measurement

z is the vertical sensor measurement

a,b,c,d,f,g,h,u,v,w are constants

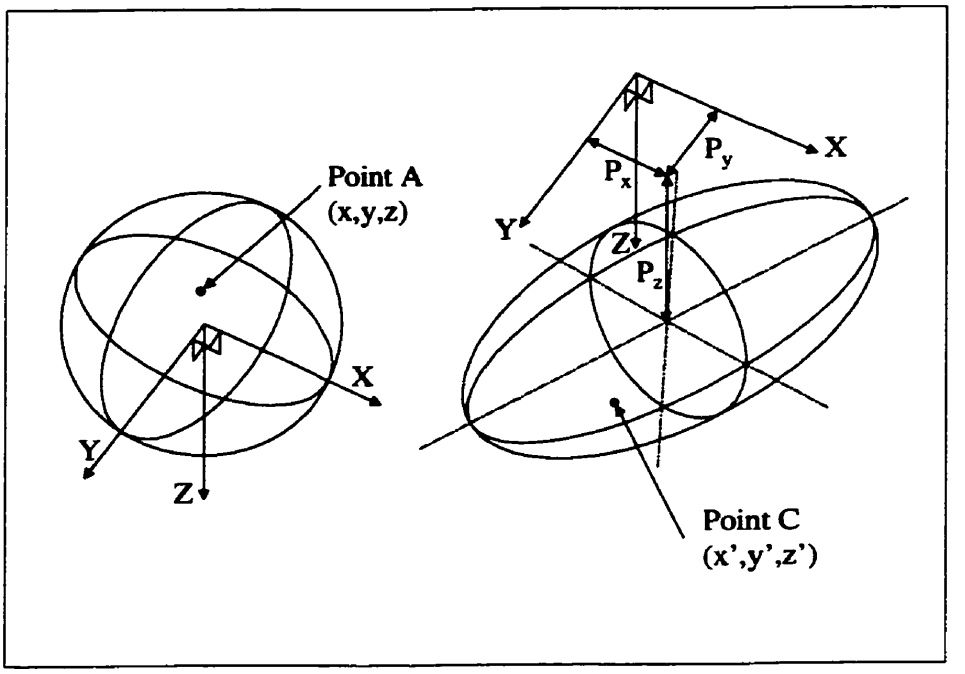

Figure 3.18 Spherical and Ellipsoidal Locus of Magnetometer Observations (Marchent and Foster 1985)

The solution to this equation requires at least nine sampled points covering the largest possible portion of the ellipsoid. This requires collecting calibration data from at least two different slopes. The d parameter is arbitrary, so choosing d=1, a solution to the constants can be given by

$$\begin{bmatrix} a \\ b \\ c \\ f \\ g \\ h \\ u \\ v \\ w \end{bmatrix} = \begin{bmatrix} x_1^2 & y_1^2 & x_1^2 & 2y_1z_1 & 2x_1z_1 & 2y_1x_1 & 2x_1 & 2y_1 & 2z_1 \\ & \ddots & & & & & & & \\ & & \ddots & & & & & & \\ & & & \ddots & & & & & \\ & & & & \ddots & & & & \\ & & & & & \ddots & & & \\ & & & & & & \ddots & & \\ & & & & & & & \ddots & \\ x_n^2 & y_n^2 & x_n^2 & 2y_nz_n & 2x_nz_n & 2y_nx_n & 2x_n & 2y_n & 2z_n \end{bmatrix}^{-1} \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} \tag{3.16}$$

The correction of measured data can then be equated to

1. Translation of the ellipsoidal coordinate system to the centre of the spherical coordinate system through (Px, Py, Pz)
2. Rotation of the spherical coordinate system such that it is coincident with the axes of the ellipse
3. Scaling of the ellipsoidal axes to force it into a sphere
4. Rotation of the sphere back to its original orientation (inverse of step 2)

Given the parameters of the ellipse, (P_x, P_y, P_z) can be given by

$$\begin{bmatrix} P_x \\ P_y \\ P_z \end{bmatrix} = \begin{bmatrix} bc-f^2 & gf-hc & hf-bg \\ fg-hc & ac-g^2 & hg-af \\ hf-bg & hg-af & ab-h^2 \end{bmatrix} \begin{bmatrix} -u \\ -v \\ -w \end{bmatrix} \times \frac{1}{|G|} \tag{3.17}$$

where G is given by

$$G = \begin{bmatrix} a & h & g \\ h & b & f \\ g & f & c \end{bmatrix} \tag{3.18}$$

The direction cosines (l_i, m_i, n_i) of the ellipsoidal axes relative to the sphere are computed from the eigenvalues of G. For the axes i = 1, 2, 3

$$l_i = R_l / R_0 \tag{3.19}$$

$$m_i = R_m / R_0 \tag{3.20}$$

$$n_i = R_n / R_0 \tag{3.21}$$

with the eigenvalues of G given by L_i

$$R_l = (b - L_i)(c - L_i) - f^2 \tag{3.22}$$

$$R_m = fg - h(c - L_i) \tag{3.23}$$

$$R_n = fh - g(b - L_i) \tag{3.24}$$

$$R_0 = R_l^2 + R_m^2 + R_n^2 \tag{3.25}$$

The scaling factors are given by

$$s_x = \sqrt{L_1} \tag{3.26}$$

$$s_y = \sqrt{L_2} \tag{3.27}$$

$$s_z = \sqrt{L_3} \tag{3.28}$$

The next step is to account for the vehicle inclination. With compass pitch, P_c, and roll, R_c

$$\begin{bmatrix} B_x \\ B_y \\ B_z \end{bmatrix} = \begin{bmatrix} \cos P_c & \sin P_c \sin R_c & \sin P_c \cos R_c \\ 0 & \cos R_c & -\sin R_c \\ -\sin P_c & \sin R_c \cos P_c & \cos R_c \cos P_c \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} \tag{3.29}$$

the heading can be computed by

$$\Psi_c = -\tan^{-1}(B_y / B_x) \tag{3.30}$$

Data was generated at a rate of five hertz by the Precision Navigation TCM2-50 digital compass. The unit specifications are given in Table 3.4. The compass provides RS232 proprietary or NMEA (National Marine Electronics Association) format output, but the proprietary format was chosen due to its increased level of detail. A pass-through log on the NovAtel receiver was used to log the data, giving the data a GPS time stamp accurate to better than 1 ms. All of the fields were output, namely temperature, three magnetometer measurements, pitch, roll, and computed heading.

Table 3.4 TCM2-50 Digital Compass Specifications (Precision Navigation 1997)

Parameter	Specification	Units
Heading Accuracy (Level)	± 1.0	° RMS
Heading Accuracy (Tilted)	± 1.5	° RMS
Tilt Accuracy	± 0.2	°
Tilt Range	± 50	°
Magnetometer Accuracy	± 0.2	μT
Magnetometer Range	± 80	μT
Supply Voltage	+5 Regulated or 6-18	VDC
Supply Current	15-20 mA	mA

The TCM2 uses a patented magneto-inductive sensing technique that makes use of a material whose inductance varies with the magnetic field strength to which it is subjected. A process similar to the method presented above was likely used in the TCM2 to remove the permanent magnetic field of the vehicle. The calibration process involves making slow turns through varied pitch and roll. A calibration score is given which indicates whether enough data with varied heading, pitch, and roll variation was collected to estimate the field of the vehicle. To avoid saturation of the sensors and difficulty in eliminating the permanent field, the contribution of the car should be less than 30 μT. Unfortunately, no mounting location above the roof of the vehicle, even raised by several decimetres, gave a field contribution below the target value.

Many attempts were made to get effective using the internal calibration algorithm, but without success. An attempt to calibrate the compass externally using the algorithm presented above and raw magnetometer and inclinometer output was made, but results achieved were no better. Extreme sensitivity was observed, e.g. the magnetic anomaly alarm would regularly go off if the brake pedal were depressed. Calibration difficulties of this severity using the same unit were not encountered by Harvey (1998), however some difficulty was documented. Precision Navigation's technical support confirmed difficulty in calibrating out a vehicle's magnetic field. The instrument is normally used on marine platforms where it is usually easy to isolate it from ferrous materials.

Table 3.5 Geomagnetic Parameters for Calgary for the Year 2000 (NOAA 2000)

Parameter	Value	Annual Change	Units
Total Magnetic Field	58030	-61	nT
Horizontal Magnetic Field	15953	+6	nT
Vertical Magnetic Field	55794	-66	nT
Declination	18.1	-0.11	° E
Inclination	74.0	-0.02	° N

It was therefore decided that an external calibration should be attempted. The method detailed above was implemented, but similarly poor results were obtained. Some authors suggest, however, that the ellipse described above is often close to spherical, and numerical instability sometimes necessitates the use of a spherical model (McLellan 1992). A simple method, based on that employed by the Etak compass, was then investigated. The initial step toward modeling of the compass errors was examining the locus of horizontal measurements from the first test performed on January 10, 2000 (which will be described in detail in a later section). The locus is pictured in Figure 3.19, and is nearly circular. The radius of the circle was estimated with LS to be about 15.5 μT, which closely matches the predicted values for horizontal field strength value for the Calgary region. Table 3.5 gives a summary of the geomagnetic parameters of interest for this research estimated for the year 2000.

The offset due to the constant magnetic field of the car was then subtracted from the measurements (also computed by LS). The resulting locus is a circle centred on the origin. The variability around the circle should be due to vehicle pitch and roll, as well as measurement errors. The Etak compass had no inclinometer data, so there was no way to compensate for these errors. However, using the inclinometer and vertical axis magnetometer measurements available from the TCM2, a transformation of the measurements to the horizontal plane was possible. The locus of transformed observations, raw observations, and the best fit circle are shown in Figure 3.20.

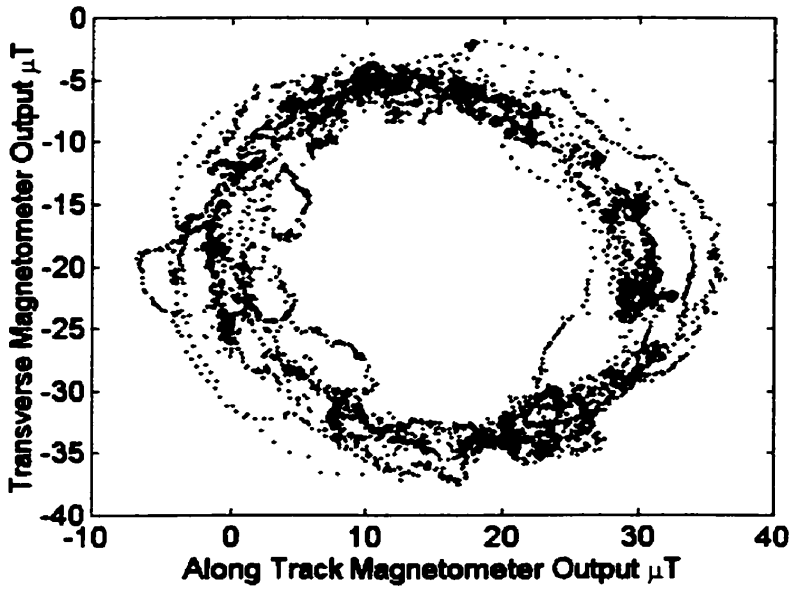

Figure 3.19 Locus of Raw Horizontal Magnetometer Output

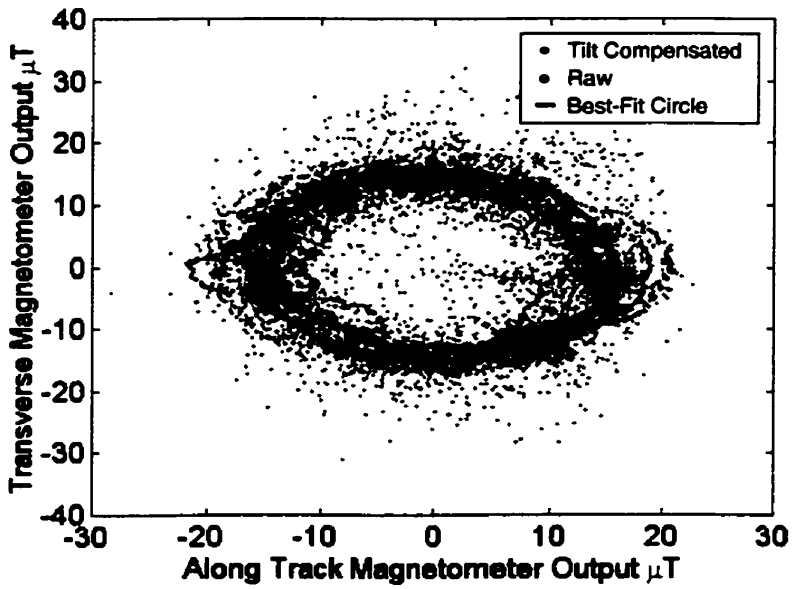

Figure 3.20 Raw and Tilt Corrected Magnetometer Output

Headings were computed for both the raw and tilt corrected observations. These were compared to GPS headings for speeds above 7 m/s, at which GPS derived headings are accurate to several degrees. The magnetic declination was used to correct magnetic headings to geodetic headings. Although the means were below 1°, the RMS error for raw headings was only 7.6° and for tilt-corrected observations it was 13.2°. From Figure 3.21, it is apparent that the inclinometer data is extremely noisy. Due to the large errors, it was decided not to continue with the use of a compass. Price was also a concern, since even the TCM2, a fairly expensive and high accuracy compass, performed poorly.

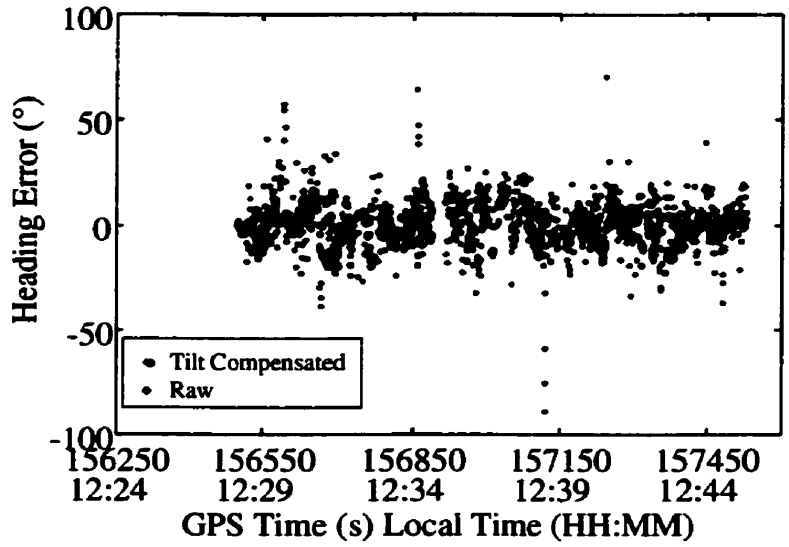

Figure 3.21 Headings Errors Derived From Raw and Inclination Corrected Measurements

3.3 INCLINOMETER

An inclinometer is a sensor used to measure the angle between the gravity vector and the platform to which it is mounted. This can be in a single direction, i.e. for sensing vehicle roll only, or two directions to estimate pitch, as well. Inclinometers suffer an error due to vehicle accelerations since it is not possible to separate them from the gravity vector. It would be possible to make corrections to the inclinometer output using the gyro and

differential odometer to isolate the vehicle accelerations, but GPS code positioning can not provide an accurate estimate of acceleration to compute this correction (Harvey, 1998).

Inclinometers for land applications must have a range of about 30° to handle normal operating conditions. Inclinometers designed specifically for low cost environments such as automotive applications would probably cost around $10. A typical 2-axis bubble type inclinometer similar to the one used in this project would be the Applied Geomechanics Model 900 Biaxial Clinometer, available at $90 in large quantities. The operation of this type of inclinometer will be described below. Several other alternatives make use of a triad of accelerometers to determine the orientation of the gravity vector.

The liquid-bubble inclinometer uses a vial, partially filled with a conductive liquid, to determine the tilt of the vehicle in one or more axes. Electrodes around the vial estimate the liquid height or height difference between sensors. These measurements are converted to pitch or roll measurements. A diagram of a typical sensor is found in Figure 3.22.

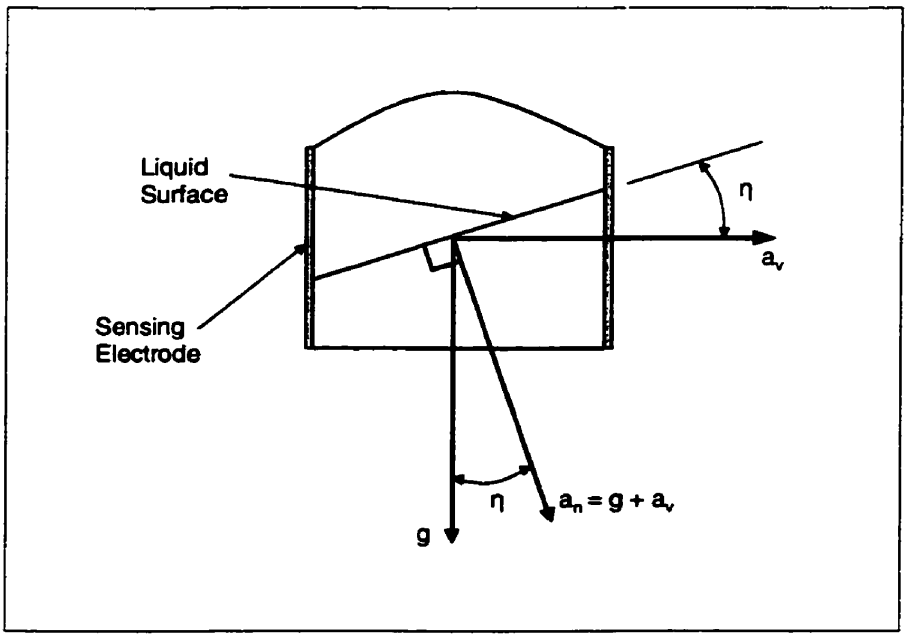

Figure 3.22 Bubble Inclinometer (Harvey 1998)

The acceleration related error suffered by bubble inclinometers is due to the fact that vehicle acceleration will make the liquid rise up one side of the vial. Also, there will be some sort of damped oscillation in the fluid even after the acceleration has finished. The error due to vehicle acceleration can be given by

$$\eta = \tan^{-1}(a_v / g) \tag{3.31}$$

where g is the magnitude of gravity and a_v is the acceleration of the vehicle.

Results obtained using the inclinometer on the TCM2 digital compass were previously presented. Errors caused by vehicle dynamics made the inclinometer observations so noisy that heading determination was also made worse. Application of vehicle acceleration corrections would likely have improved the inclinometer measurements somewhat, but due to the magnitude of the errors, it was not pursued further.

3.4 ODOMETER

An odometer is used to measure the distance traveled by the vehicle, or possibly by the individual tires. Most vehicles have a transmission based odometer to track the total distance or trip distance traveled by the car. This type of odometer counts the revolutions of the powertrain after the clutch (manual transmission) or torque converter (automatic). Wheel rotation sensors are often used for ABS and traction control systems and count the rotations of two or four wheels. ABS is becoming a standard feature on most vehicles, providing an easily interfaced dead reckoning sensor at no additional cost. The process of counting the revolutions often uses the coupling of mechanical and electrical forces or optics to detect identifiable features etched into or protruding from the shaft. There can be any number of digital pulses (discrete) or sinusoids (continuous) accumulated per rotation of the shaft; some simple systems use one per rotation, while others accumulate hundreds for demanding applications such as ABS or traction control. Both wheel and transmission odometers can be used to aid in navigation application by providing speed,

distance, and possibly heading information. Several types of sensor mechanisms are presented below and a table showing accuracy and limitations is given in Table 3.6.

Table 3.6 Comparison of Several Odometry Sensors (Geier 1998)

Sensor	Accuracy	Output	Limitations
Variable Reluctance	50% of tab spacing	Discrete	Cut-off speed
Hall Effect	5-50% of tab spacing	Discrete or Continuous	-
Optical	0.1% of a rotation	Discrete or Continuous	Often difficult to mount
Doppler Radar	1% distance traveled	Continuous (Speed)	Errors from alignment and road irregularities

3.4.1 VARIABLE RELUCTANCE SENSOR

A diagram of a variable reluctance sensor is found in Figure 3.23. According to Faraday's Law, if a coil is perpendicular to a varying magnetic field, a current (or voltage) is created in the coil. When the ferrous protruding tab passes through the gap in the magnet, the magnetic field strength will drop inside the magnet since a magnetic circuit will be created. The output signal is proportional to the rate of change of magnetic flux, and thereby the speed that the shaft is rotating. If the shaft is rotating slowly, a very weak signal is produced. This rotation speed at which the signal is no longer identified by the signal processor is usually referred to as the cut-off speed. Only a single tab is shown, but usually there would be several surrounding the shaft, giving a discrete count of fractions of a full rotation.

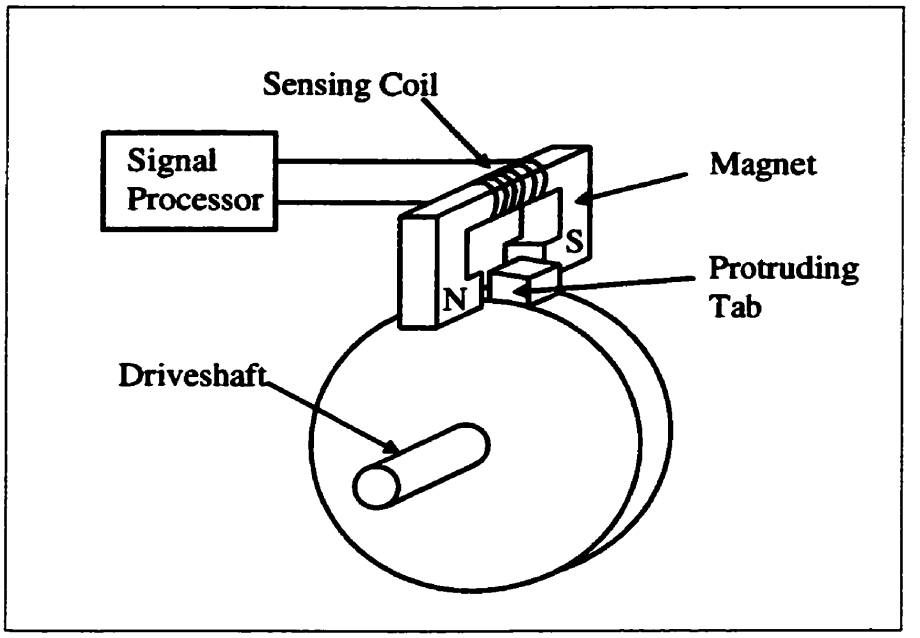

Figure 3.23 Variable Reluctance Rotation Sensor (Ribbens 1994)

3.4.2 HALL EFFECT SENSOR

The Hall effect odometer used in the test vehicle allows sensing of lower wheel speeds than many other simple mechanical or pulse detection sensors. The analog signal generated by this type of sensor is proportional to the strength of the magnetic field over the sensor, i.e. the distance from the magnetic source mounted on the shaft, and is therefore sinusoidal. This allows for measurement of the phase of the rotation, which is much more accurate than counting full cycles only. The Hall effect occurs when a charge carrier moving through a material experiences a deflection because of an applied magnetic field. This deflection results in a measurable potential difference across the side of the material that is transverse to the magnetic field and the current direction. The situation is presented graphically in Figure 3.24. A typical Hall effect digital speed sensor is shown in Figure 3.25 in which the naturally occurring sinusoidal signal would be squared up and output as a digital pulse train.

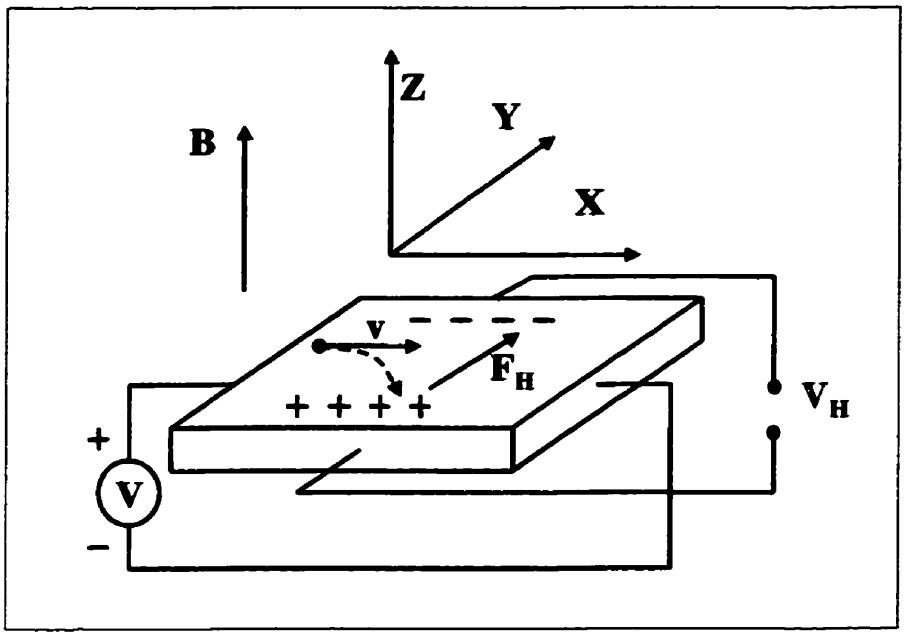

Figure 3.24 Hall Effect (Bart Van Zeghbroeck 1998)

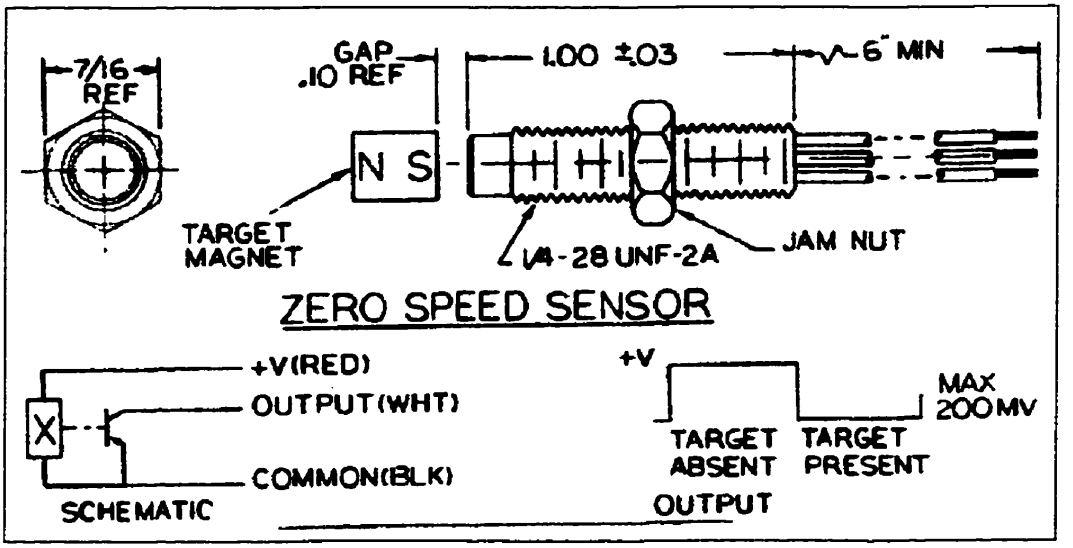

Figure 3.25 G&S Digital Hall Effect Sensor (G&S 1998)

3.4.3 OPTICAL SENSOR

There are several varieties of optical sensors available for measuring the rotation of a shaft. The simplest of which is a proximity sensor such as the one shown in Figure 3.26.

The unit emits an IR (Infra-Red) signal from a GaAs emitter, which is reflected from a nearby surface (from 1-5mm for effective return) to the collector. The collector is a silicon photodiode that will induce an output current proportional to the flux of returned light. Optical shielding blocks out all but the emitted frequency from detection. The signal would be reflected off of a shaft with a series of tabs surrounding it. A proximity sensor would measure the strength of the reflections from the shaft, and add to the count each time the detection threshold was reached, i.e. when the tabs passed the detector. Like the variable reluctance sensor, the observation is a discrete count of fractions of a full rotation. Similar systems pass a continuous signal past the edge of the shaft that is interrupted by the passing tabs. Electro-optical ICs have become increasingly affordable in recent years so sensors such as this are becoming quite common.

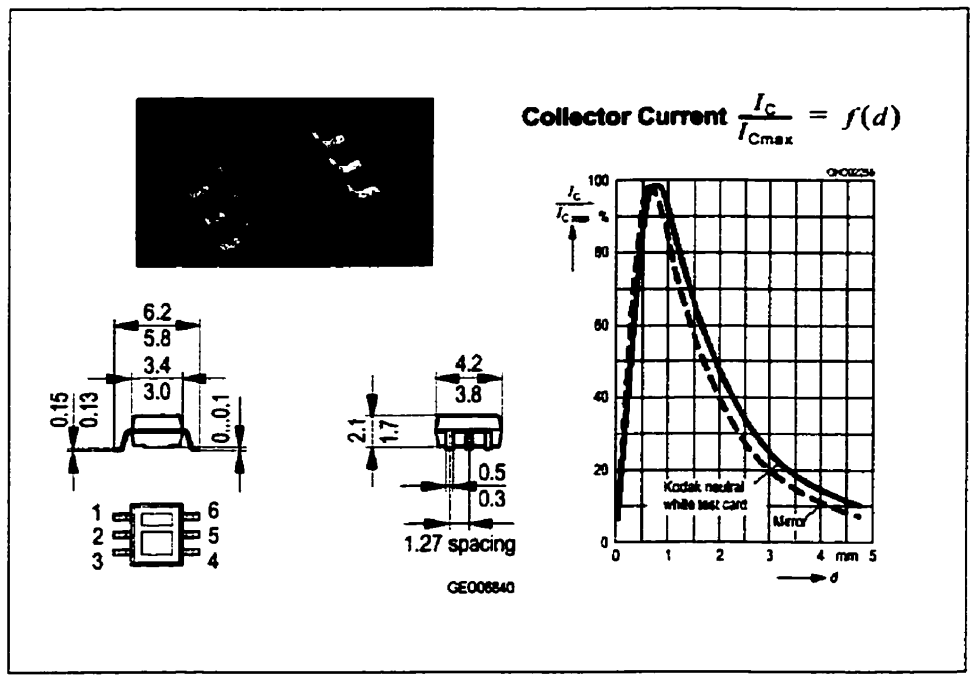

Figure 3.26 Infineon SFH9201 Reflective Optical Proximity Sensor (Infineon 2000)

Much more accurate optical sensors have been developed which make use of counting the passing of very fine lines, similar to the bar codes found of most commercial packaging. This technology is referred to as optical encoding. A sample of this type of sensor is found in Figure 3.27. The assembly surrounds the shaft, with the inner coded

wheel spinning on the shaft and the outer assembly (reader) fixed. Options allow from 50 to 1024 sinusoids given in two quadrature channels, or 4096 pulses per revolution. The cost is $47 in large quantities for most models.

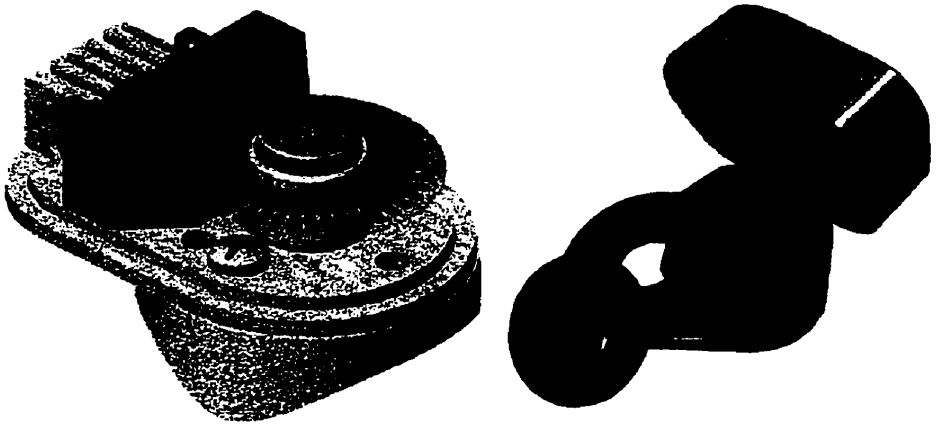

Figure 3.27 Optical Encoder Speed Sensor (U.S. Digital 2000)

3.4.4 DOPPLER RADAR

Doppler Radar systems transmit an RF (Radio Frequency) signal that is reflected from a target object. If there is relative velocity between the two, the return signal will be Doppler-shifted with respect to the transmitted frequency. This principle is used in Police radar for determining the speed of a car. The Doppler shift induced by relative motion can be given by (Young 1992)

$$f_L = \sqrt{\frac{c-v}{c+v}} f_S \tag{3.32}$$

where

f_L is the frequency received by the assumed stationary observer

f_S is the frequency transmitted by the moving source

c is the speed of light

v is the relative velocity between source and observer

The returned signal is most often mixed with the output signal. With the output signal s_1, at a frequency f, and the returned signal s_2 shifted by Δf, i.e.

$$s_1 = a \sin(2\pi f_0 t) \tag{3.33}$$

$$s_2 = -b \sin(2\pi (f_0 + \Delta f) t) \tag{3.34}$$

The sum of the output and returned signal is

$$s_1 + s_2 = ab \sin \left(2\pi \left(\tfrac{1}{2} \Delta f\right) t\right) \cos \left(2\pi \left(f + \tfrac{1}{2} \Delta f\right) t\right) \tag{3.35}$$

For relatively small Doppler shifts, the sum can be interpreted as an envelope equal to half of the Doppler shift with a high frequency component that is the average of the two frequencies. This beat frequency is estimated and used to compute the relative velocity. A typical example is shown in Figure 3.28.

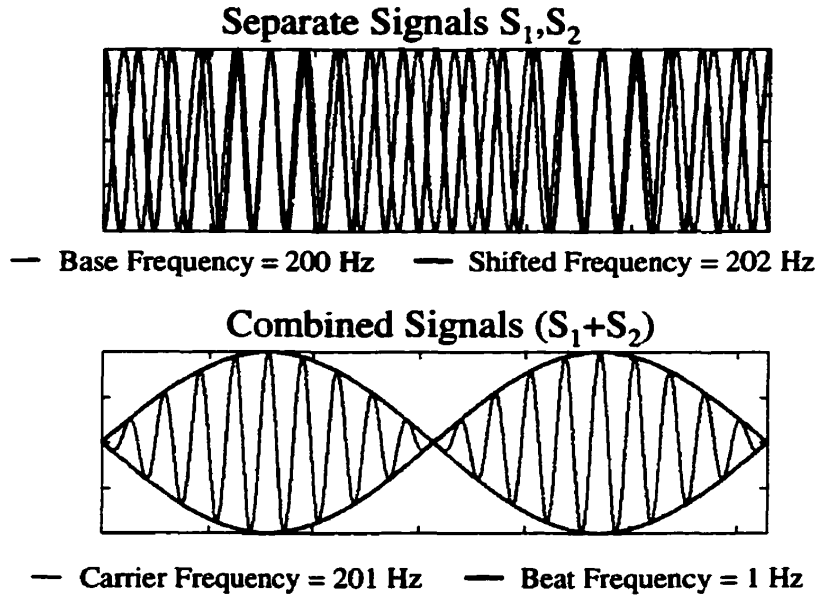

Figure 3.28 Combination of the Original and Doppler-Shifted Signals

A Doppler radar system which can be used on vehicles is shown in Figure 3.29. Although this system is quite large, low cost IC alternatives are available similar to motion detection sensors for alarm systems (Infineon 2000). The signal is usually directed angularly at the road from the underside of the car. Due to the use of microwave frequencies, the signal can pass easily through a short distance of plastic, fiberglass, or other non-conductive material. This allows the sensor to be placed in a secure location such as a bumper. One notable difference between a radar system and other systems described as odometers is that the sensor actually estimates ground speed, and not the rotation of the drivetrain or wheels. It would properly be termed a speedometer system.

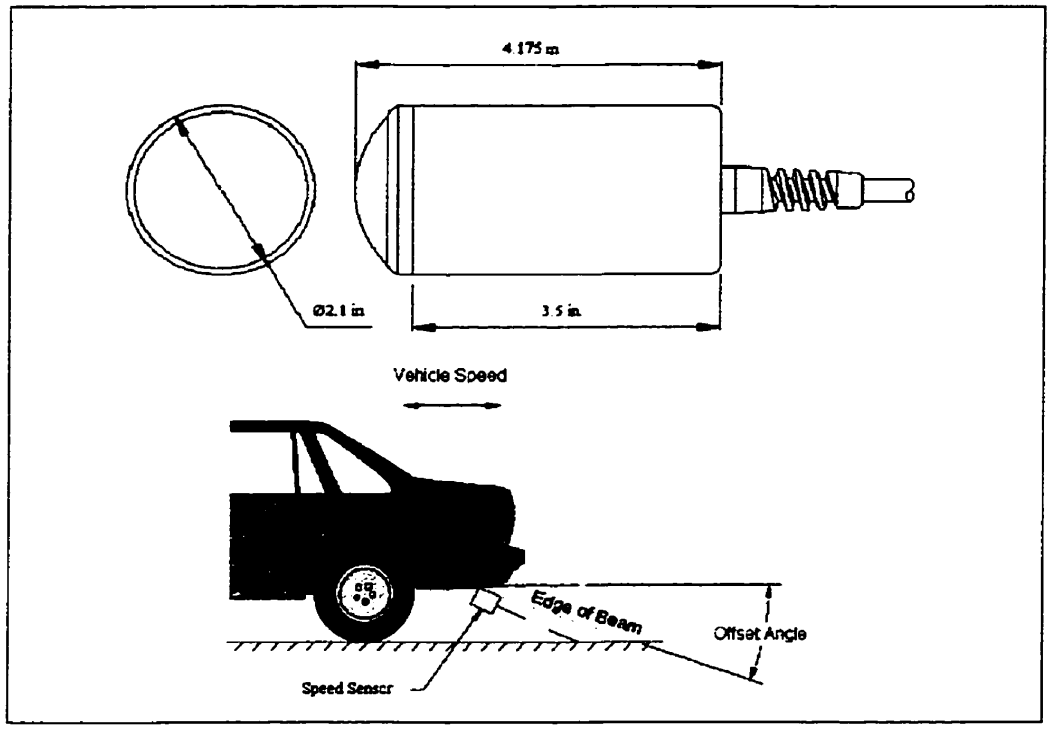

Figure 3.29 GMH Delta :Doppler Radar Speedometer System (GMH 2000)

There are several limitations of radar speed measurement, however. Firstly, the sensor is pointed at an angle to the road and therefore only senses the direction cosine of the forward movement. If there is any vibration, the angle will change and a different speed will be sensed. Road irregularities and vertical motion will cause transient variations in the measured speed, as well. If the car passed over a bump, an increase in speed would be detected on the leading edge, and a decrease on the trailing edge. A scale factor could be used to compensate for vertical axial alignment error of the beam that would be equal to the cosine of the alignment error. The sensor would also detect transverse motion if the sensor axis was not in alignment with the along track axis of the vehicle.

3.4.5 ODOMETER CALIBRATION

It is necessary with all but the Doppler speed sensor to equate the shaft rotation with the distance traveled. The determination of scale factor will be described for individual wheel

based odometers, but the same principles apply to the transmission type of sensor. The scale factor is based on the tire circumference and slippage.

Slippage is minimal on clean, dry roads. It occurs when acceleration and deceleration of the vehicle make the contact point between road and tire support a transverse load. During level driving with constant heading and speed, only a normal load is supported. The slippage will therefore vary with vertical slope, heading change, and speed change. Slippage is dependent on both type and condition of the tires and road surface. With modern tires and normal driving habits, slippage seldom occurs which exceeds 1% of the distance traveled. Wheel slippage is higher for the driven wheels (connected to the powertrain) than for undriven wheels, due to acceleration and frictional deceleration of the engine. Modern braking systems apply the braking force quite evenly, but often with slightly more pressure to the front brakes.

To minimize the error caused by slippage, the undriven wheels are normally used to aid in navigation. Slippage still occurs to some extent, but mostly during cornering and braking. The differential allows the wheels to rotate at different speeds in order to minimize the amount of slipping on corners, but this also allows the wheels to slip differently based on the surface under each tire. Cornering causes the weight of the vehicle to be shifted to the outside wheel, so the friction and slippage are different for the two wheels. The varying road surface caused by water, ice, snow, gravel, and other road contaminants, causes the friction to vary greatly between tires and changes the slippage. For example, when braking on patchy ice, it is common for one or several wheels on ice to lock and the ones on pavement to keep spinning.

Due to the complex nature and variability of slippage related errors, the only method of coping is by monitoring for potential slippage occurrences and proceeding with caution in using data which may be in error. Slippage normally only amounts to less than a decimeter for a serious occurrence, so it is not a substantial error source for distance odometry. For differential odometry, such an error can have a large impact.

The circumference of the wheels relate the rotation of the tire to a distance traveled, excluding slippage. Modern radial tires have fairly constant size and shape, unlike the bias-ply tires which preceded them. Studies have shown that several factors affect the tire radius, including temperature, speed, pressure, and weight loading. Estimated contributions of these factors is given in Table 3.7.

Table 3.7 Factors that Affect Wheel Scale Factor

Error Factor	Possible Error in Radius
Pressure	1 mm / psi
Temperature	1 mm / 5°C
Wear	5 mm
Speed	1 mm
Weight	1mm / 100kg

Temperature has very little effect, since the rubber of which tires are made is designed to dissipate large amounts of heat caused by road friction. The coefficient of linear thermal expansion for natural rubber is 0.067%/°C, which would probably be an order of magnitude higher than after vulcanization (Heidemann et al. 1987). An 80°C temperature change would be possible, since the tire could start from –30°C and rise to 50°C dissipating the heat of highway driving; this would still affect the sidewall length by less than 1 mm. Tire rotation causes a centrifugal force on the material of the tire that tends to stretch it out as in Figure 3.30. This stretching accounts for a small change in the effective radius of the tire, which increases with speed. This change is minimal with properly inflated tires, and still very small if they are under-inflated. The vehicle's weight load changes the radius by less than 1 mm / 100 kg, and air pressure about 1 mm/psi (Zickel and Nehemia 1994). Tire pressure also varies with temperature at a rate of about 5.5°C/psi (RepairWare 1998). Finally, there is tire wear, which can change the diameter by 10 mm throughout the life of the tire.

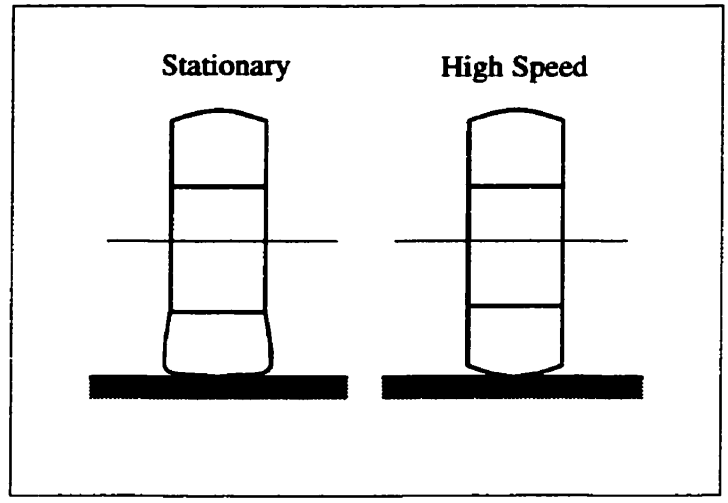

Figure 3.30 Stretching of the Tires with Speed

The distance traveled by the wheel should have an accuracy of better than 1% to meet the 10 m accuracy guideline after several minutes and several km of driving, so the scale factor must be calibrated. The alternative would be to use a fixed number based on wheel size, such as the assumed scale used by the dashboard odometer. Calibration is usually done using GPS, either by comparing GPS and odometer velocities, or interval distance traveled observations. GPS velocity errors could be worse by one or two orders of magnitude than the desired accuracy of the scale factor using a filter. To keep GPS speed accuracy above this level, a calibration cut-off speed is used. Assuming a DGPS velocity accuracy of about 0.3 m/s, a 5% error is achieved at 6 m/s. The ratio of GPS speed and wheel rotation count could be filtered above such a threshold. Another threshold is often used to detect a large change in tire size, such as loading the vehicle or changing tires, which causes a reset of the scale calibration filter. A flat tire can also be detected before further damage is done to the tire by driving on it.

3.4.6 DIFFERENTIAL ODOMETRY

When a vehicle turns, the differential allows the inside tire to travel a shorter distance than the outside without greatly increased slippage. Differential odometry uses the two wheel-based speed or distance measurements to estimate the change in heading. The

individual wheel speeds would therefore vary from the GPS speed of the vehicle while it is turning. If small gain factors were used for the calibration filters, this error would cancel out, assuming that the amount of turning in each direction should eventually even out.

Distance traveled over ground by the vehicle is computed through averaging the two accumulated distances, where the subscripts indicate left and right distance measurements.

$$\Delta d = \frac{\Delta d_L + \Delta d_R}{2} \tag{3.36}$$

Using the difference between the accumulated distances of the two tires, the heading change can be determined. A method has been suggested by Harris (1989), which is illustrated in Figure 3.31, whereby the heading change can be computed as

$$\Delta az = \frac{\Delta d_L - \Delta d_R}{Track\,Width} \tag{3.37}$$

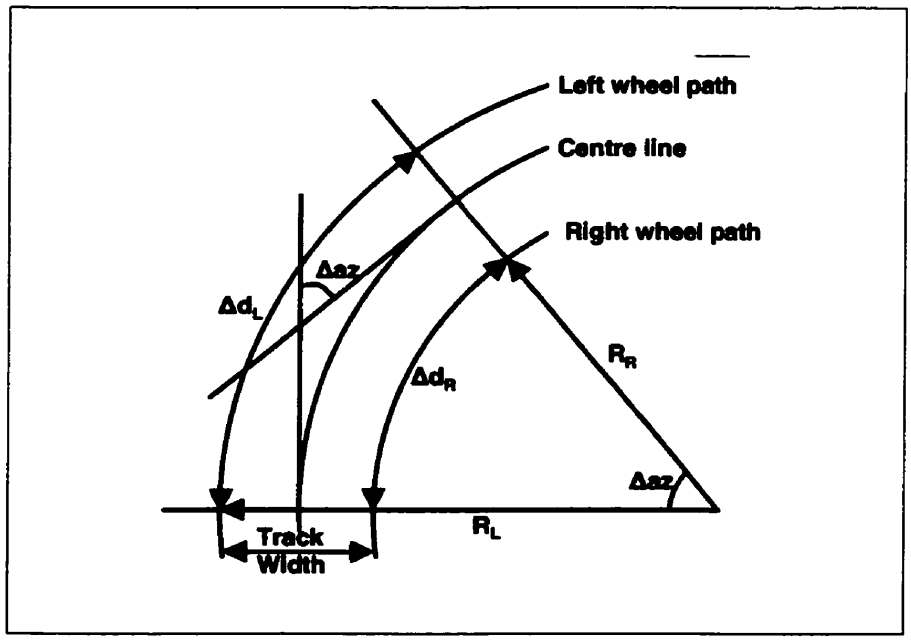

Figure 3.31 Differential Odometer Geometry (Harris 1989)

The track width is the effective turning width of the car (approximately the distance between tire centres). This formula holds for fixed wheels, which do not turn with the steering wheel. Etak has developed an algorithm for using the front wheels of a vehicle to estimate distance traveled and heading change (Zavoli et. al. 1988). Scale factor errors and slippage become a serious concern in differential odometry, however. Some numerical examples will demonstrate this. Assume that a vehicle is coming to a stop on patchy ice. If one wheel slips while the other is locked for 0.1 m, with a track width of 1.5 m, a 4° heading error is incurred. The result of a 0.5% error in the scale factor for a single tire is 0.2°/m, accumulating 2° for each 100 m block. To meet the specifications desired for the system, scale factors would need to be calibrated to an accuracy of several hundredths of a per cent, which is a difficult task given GPS speed observations with an accuracy of 1-5%. Calibration must therefore be done as accurately as possible, but some drift and slippage error is to be expected.

3.5 ACCELEROMETER

An accelerometer measures platform acceleration, which can be integrated to give velocity, and double integrated to give distance traveled. Other uses for accelerometer data include the detection of impacts for air bag deployment, and inclination measurement through sensing of the Earth's gravitation. Vibration measurements are now being investigated to detect potential engine problems before further damage or failure results. Sensor drift makes regular ZUPTs (zero velocity updates) necessary for periods without an accurate external reference, since the double integral can accumulate substantial errors. The drift is mostly temperature related, so investigation into the various compensation schemes has been done. The use of constant temperature ovens can ensure good stability after warm-up, but an oven uses a great deal of power and a much larger space. Another technique is to generate a temperature profile and count on temperature repeatability; this can be done a priori or built up over time in a real time application.

Accelerometers are generally based on observing the displacement of a suspended mass caused by inertia. Two common implementations, a damped spring and pendulum, are shown in Figure 3.32. Methods such as differential capacitance (described below), inductance, or optical methods can be used to measure the displacement. Sometimes a magnetic field or servo is employed to keep the mass in a fixed position. Purely optical methods of measuring acceleration, similar to those used in optical gyros, have been developed recently, as well.

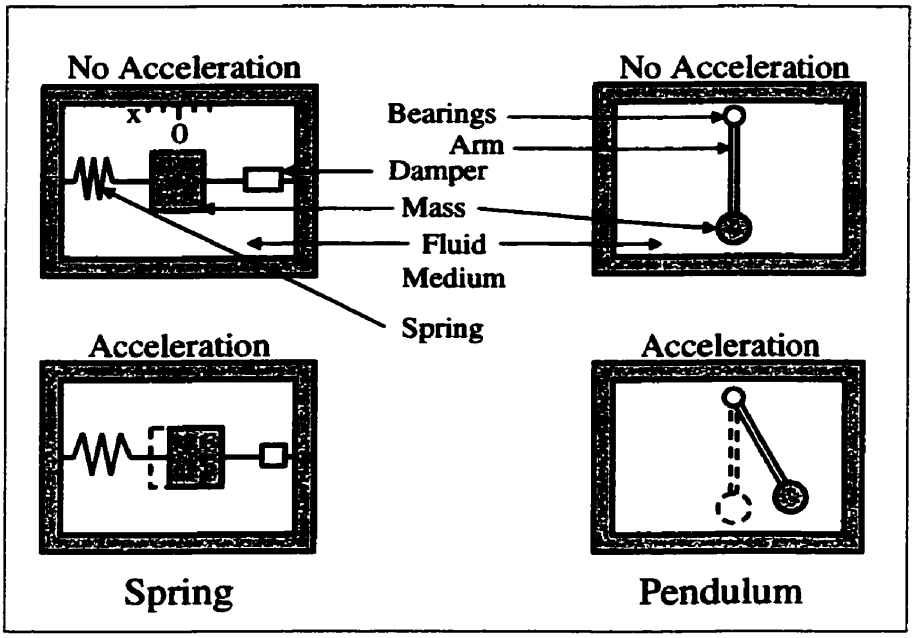

Figure 3.32 Spring and Pendulum Accelerometers (Hayashi 1996 and Savage 1978)

A damped spring will allow the suspended mass to displace under an acceleration. The movement of the mass would be sensed through capacitance, an optical method, or otherwise, while damping is usually accomplished by the use of a viscous fluid medium. The displacement can be described by

$$F = ma = m\left(\frac{d^2x}{dt^2}\right) + c\left(\frac{dx}{dt}\right) + Kx \tag{3.38}$$

where

F is the applied force

m is the mass of the suspended mass

c is the damping coefficient (a function of the medium)

K is the spring stiffness

x is the displacement of the spring relative to resting position

A resting pendulum will also displace under an acceleration, moving further up its arc with greater acceleration. The motion is again usually damped by the medium, and quality bearings are used to minimize friction. The amount of displacement depends on the weight distribution in the pendulum, but assuming the mass is all at the bottom of the arm, the arm would point in the same direction as the sum of the gravity and acceleration vectors. Angular or linear displacement of the mass could be measured using one of the methods mentioned above.

Many modern designs use MEMS technology, including the offerings from Analog Devices (2000). In them, a small proof mass is suspended from two sides with flexible coils. When the platform is accelerated, the displacement of the mass is measured by a differential capacitance as shown in Figure 3.33.

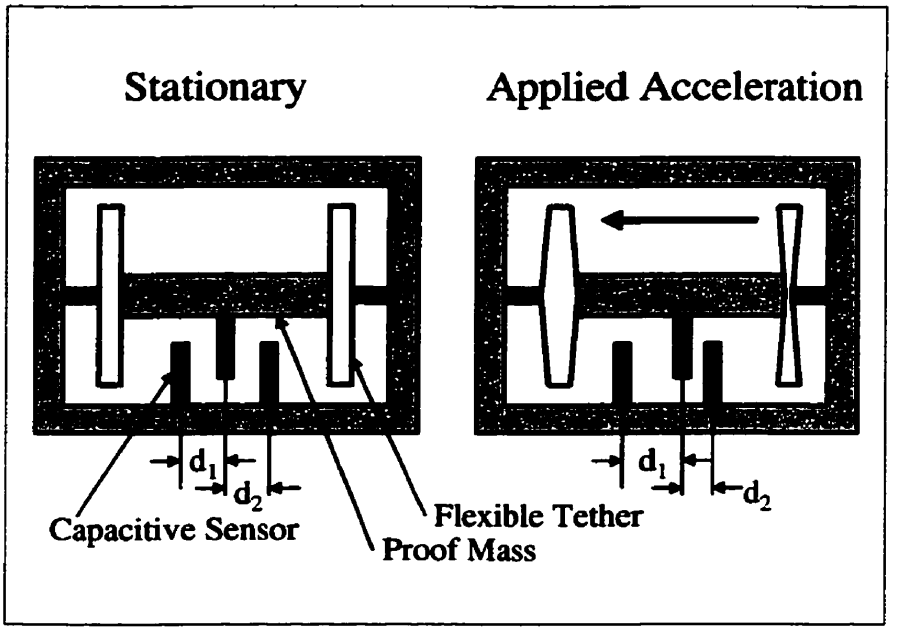

Figure 3.33 MEMS Accelerometer (Analog Devices 2000)

The sensors are surface micromachined just like a standard IC (Integrated Circuit) and are therefore available in standard IC sizes ready to be mounted on any circuit board. The moving parts are cast in place using a sacrificial layer of material that is then dissolved, making the positioning of the parts highly accurate. The small size and precision of construction are apparent from the specifications given Table 3.8

Table 3.8 Analog Devices ADXL202/210 Specifications (Analog Devices 2000)

Parameter	Specification	Units
Proof Mass	0.1	µg
Capacitor Plate Spacing	1.3	µm
Minimum Detectable Beam Deflection	0.02	nm
Minimum Detectable Capacitance Change	20	aF
Full Scale Capacitance Change	10	fF
Typical Zero Acceleration Bias	20	m/s^2
Typical Bias Temperature Stability	.02	$m/s^2°C$
Typical Scale Temperature Stability	0.5	%
Maximum Range	20/100	m/s^2

3.6 MAP MATCHING

Map matching refers to the procedure of comparing the present coordinates of a vehicle with a digital map (Bullock 1995). The position should correlate with the map, i.e. the map should indicate that the vehicle is on a road, parking lot, or some other place where a vehicle could possibly be. Also, if a car is known to be on a particular road, one can determine both possible directions for the vehicle to travel, and what adjoining roads it may move to. It is also necessary to show the operator an estimate of their location that is meaningful. The driver would not be able to make the correct navigation decisions if the apparent location of the vehicle corresponds to the middle of a building on a digital display. Most vehicle navigation systems use map matching for several of the following functions.

1. Graphical display to help the user interpret the navigation information
2. Positioning aid, including direction, snap to road, and connectivity
3. External check of the navigation solution or system initialization
4. Route finding

Map matching is also invaluable for calibrating dead reckoning sensors. For example, the odometer scale factor can be easily derived if a digital map gives road distance covered during a segment of travel. Gyro bias and scale factor may also be determined. If a road is known to be straight, the rotation rate can be assumed to be zero and any bias induced rotation can be calibrated out. Gyro scale factor estimation is possible since the angle of intersection between roads can be determined from a map and compared to the measured heading change.

The limitations of map matching stem mostly from the accuracy of the map. Accuracy includes the estimated error of lines and points due to survey or digitizing errors, but also the age of the information and whether it is out of date. New streets are regularly added which may not be on the map, so error in route finding algorithms will occur if road connectivity changes streets closed. A post-mission map matching algorithm was used to

evaluate the data collected for several tests of the proposed navigation system, and many road changes were found to have been present since the map was generated three years ago. It is not prudent to assume that a vehicle should always be somewhere that makes sense according to the map. The adoption of standards could also allow traffic and road map change information to be continually received by users.

4 FILTER METHODOLOGY

Integrated dead reckoning systems usually combine some sort of absolute positioning method with at least one heading sensor and one motion sensor. Absolute positions can be acquired from GPS, Loran, or some other radio navigation system, but many older systems allowed for the entry of street coordinates or a mouse click on a map display to give the starting coordinates. Once the system is initialized, dead reckoning may be solely used to monitor changes in the vehicle location until the next absolute position fix is entered. Using a radio navigation system, the absolute coordinates may be continually entered into the navigation solution, as they become available. This procedure usually requires some sort of filter, such as the Kalman filter, which is described in detail and was chosen for this project.

The chosen filter must be numerically stable and preferably allow a method of detecting blunders or sensor failures in the data. This includes the ability to detect large multipath errors, wheel slippage, errors due to inclination of the gyro, and magnetic anomalies sensed by the compass. Choices must be made to balance filter performance, fault tolerance, and computational load, as well. Kalman filtering is commonly chosen due to its intuitive discrete time design, and statistical optimality.

4.1 KALMAN FILTER

Kalman filtering has been used extensively in navigation data processing for several decades. Its recursive implementation is well suited to the fusion of data from different sources at different times in a statistically optimal manner. Many other filter designs can be shown to be equivalent to the Kalman filter, given several constraints. The recursive sequence involves prediction and update steps. The prediction step uses a dynamics model that describes the relationship between variables over time. A statistical model of this dynamic process is also necessary. A prediction is usually done to estimate the variables at the time of each measurement, as well as in between measurements when an

estimate is required. The measurement update combines the historical data passed through the dynamics model with the new information in an optimal fashion. The discrete-time linear measurement filter equations with no deterministic input (control) are given by (Gelb, 1974):

Given a system model with no deterministic input

$$\dot{x} = f(x,t) + u(t) \tag{4.1}$$

$$F = \frac{\partial f}{dx} \tag{4.2}$$

$$\Phi_k = L^{-1}\left\{\left[sI - F\right]^{-1}\right\}_{t=\Delta t} \tag{4.3}$$

where L^{-1} is the inverse Laplace transform. With linear measurement model

$$z_k = H_k x_k + v_k \tag{4.4}$$

and supplied with initial conditions P_0 and X_0, the prediction equations can be given by

$$x_{k+1}^- = \Phi_k x_k^+ \tag{4.5}$$

$$P_{k+1}^- = \Phi_k P_k^+ \Phi_k^T + Q_k \tag{4.6}$$

$$Q_k = \int_{t_k}^{t_{k+1}} \Phi_k(\xi) w w^T \Phi(\xi)^T d\xi \tag{4.7}$$

and the update equations by

$$x_k^+ = x_k^- + K_k[z_k - H_k x_k^-] \tag{4.8}$$

$$P_k^+ = P_k^- - K_k H_k P_k^- \tag{4.9}$$

$$K_k = P_k^- H_k^T [H_k P_k^- H_k^T + R_k]^{-1} \tag{4.10}$$

where

x is the vector of unknowns

P is the covariance of the unknowns

Q is the dynamics noise matrix

K is the Kalman gain matrix

R is the covariance matrix of the observations

Z is the vector of observations

H is the design matrix

w is the vector of spectral densities

Φ is the state transition matrix

Although an error state Kalman filter is desirable in some applications, a state space model will be discussed here as it most clearly illustrates the operation of the filter. Rigorous development of a Kalman filter requires a great deal of work in understanding the physics and electronics involved in each sensor to understand what the error sources will be (Maybeck 1994). A comprehensive (truth) model of all of the necessary variables to near perfectly match the real world phenomenon is needed. Given the specifications of the project, a sensitivity analysis would then be done to decide which variables may be ignored or lumped together. This often involves Monte Carlo simulations with several likely candidate filters and repeated tuning of the statistics. The filter must be able to operate with the allowed throughput and processing restrictions. Finally, blunder detection, adaptive filter gains, and practical limits to covariances must be set to achieve optimum performance.

There are several key factors for the proper implementation of a Kalman filter. These include

1. An adequate model for vehicle dynamics
2. An adequate model for the errors associated with each sensor
3. Adequate estimates of all statistics

Research has shown that the constant velocity model is adequate to handle most land vehicle dynamics (Lan 1996), however some implementations use a constant acceleration model. The tradeoff must be made between a higher order dynamics model, meaning more variables to estimate, and the need for a Gaussian error for the model.

A Kalman filter is assumed to be driven by Gaussian white noise. This white noise is often not an adequate description of the errors associated with the sensor measurements. Shaping filters can be used to transform the Gaussian noise into an appropriate error model. Some examples are presented in Figure 4.1. Once the modeling is complete, the variance of each of the driving noises must be determined. The variance is not necessarily constant over the data set, so adaptive noises or appropriate estimates must be selected. Each sensor's shaping filter (bias states) can be created independently and then amalgamated or kept apart in a decentralized filter as described below. Amalgamating the states allows for the filter to build a correlation between states which are related, but keeping them separate avoids corrupting others with the errors of a sensor which has blunders or is poorly modeled.

A Kalman filter allows a great deal of flexibility for informing the filter of correlations between observation errors, and between dynamics errors. The filter is able to recognize the correlations between variables as a result of the statistics and geometry. Because of this, an observation of one or more states may affect many or all of the other states. An undetected blunder can therefore have devastating consequences on the entire filter, often necessitating a total reset of the filter and the loss of all information from the historical data. Blunder detection is important, but errors such as unmodeled drifts will often escape a blunder detection scheme. These slowly varying errors can sometimes be detected through examining the history of the innovation sequence, or they can be allowed to accumulate to the point that they can be detected; the latter can be accomplished with decentralized filtering.

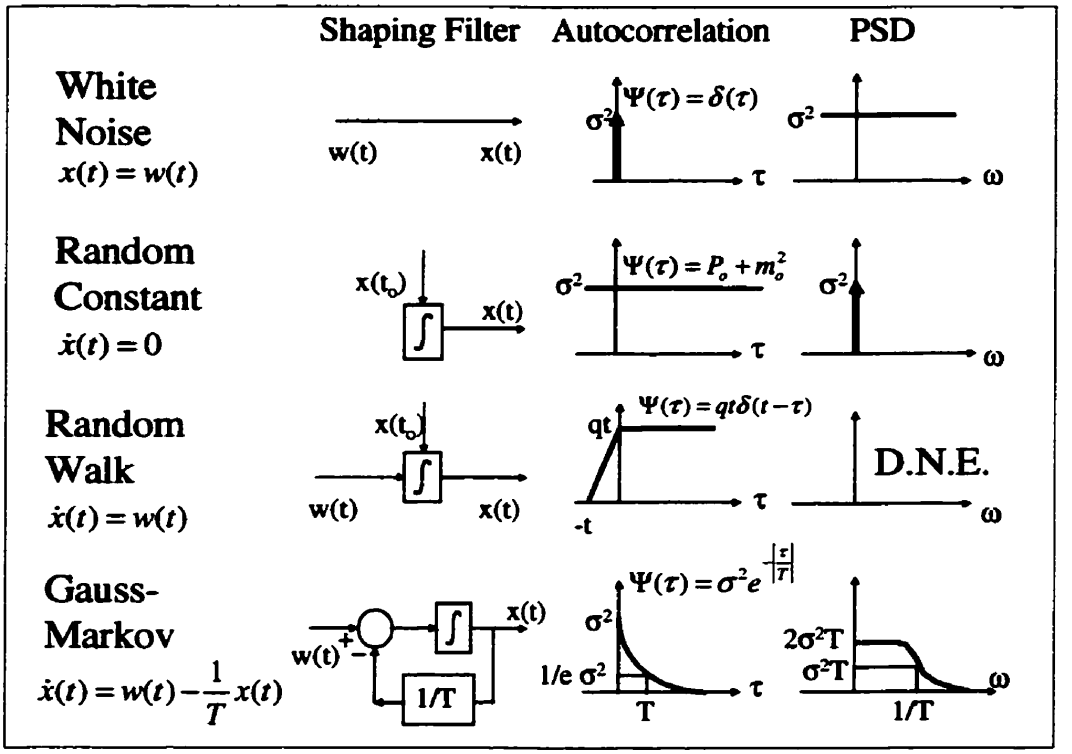

Figure 4.1 Common Shaping Functions (Gelb 1974)

4.2 DECENTRALIZED KALMAN FILTER

Decentralized filtering refers to the design of a tiered filtering system where sensor output is filtered in local filters and the results of the local filters then combined in a master filter. There are many ways to implement such a system, but many are sub-optimal, i.e. they do not conserve and use information in a mathematically rigorous fashion. An optimal filter is by no means necessary for good performance, but usually forms a solid foundation for a successful design. A simple example of a cascade (decentralized) filter could be simply low-pass filtering of a data stream before it goes to the master filter. Decentralized filtering provides several advantages over a centralized filter, including increased fault tolerance (blunder detection) and a decrease in numerical computations (Gao et al. 1993). The number of flops (floating point operation) needed for matrix multiplication decreases dramatically by reducing the sizes of the matrices. For example, if a 20 state centralized matrix could be reduced to three 10 state matrices and a

5 state master filter, a measurable savings is incurred. Each multiplication of two n by n matrices requires $2n^3$ flops, so each centralized matrix multiplication could be reduced from 16000 to 6250 flops for all of the decentralized filters combined. The fault tolerance aspect is discussed below.

A much more rigorous implementation is the optimal federated filter (Carlson 1993). In the federated filter, the information is carefully conserved in a manner that will allow proper output statistics to be derived from input statistics. Local and master filters are both Kalman filters. Tuning of most sub-optimal decentralized filters to give reasonable statistics and accurate values for the state estimates is a difficult and ad-hoc procedure. Using a federated filter, the statistics do not have to be greatly modified from the centralized formulation, but the different states and observations must be very weakly or not correlated.

The federated filter consists of a master filter and any number of local filters, often with a single filter for each sensor. The master filter contains the primary states of interest, and the local filters each include the master filter states and the necessary sensor bias states. The dynamics model is applied to each filter, but the statistical weight of the dynamics driving noise is distributed such that information is preserved. The individual LFs (local filters) are considered sub-optimal for this reason. The local filters then combine to give an optimal global estimate through the MF (master filter). The information can be returned to the local filters, or kept in the master filter, according to several schemes described below and referring to Figure 4.2. The methods are summarized in Table 4.1, describing the reset info symbol in the figure, and whether the master filter is cleared of information or needs to run a prediction cycle.

Table 4.1 Federated Filter Reset Schemes

Scheme	Reset Info	MF Reset	Information Resides
No Reset	None	Yes	Local Filters
Fusion Reset	P,X	Yes	Local Filters
Zero Reset	P=0	No	Master filter
Rescale	P=sP	No	Shared LFs/MF

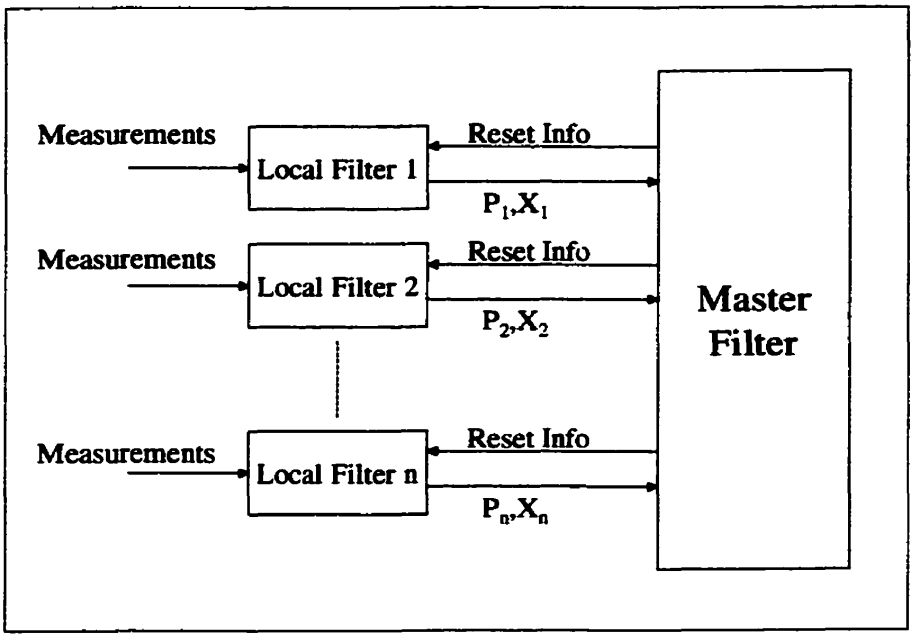

Figure 4.2 Federated Filter Architecture (Carlson 1993)

The no reset filter simply combines all of the information at a common epoch. A prediction cycle would only be used in the master filter to estimate states between master filter updates, and the master filter is cleared of all information before each update. The fusion reset filter combines the filter information in the same manner, but updates the local filters by sharing the information from the master filter solution between the local filters. The local filters therefore have the same estimates of the primary state variables after a master filter cycle. These two implementations require the local filters to be propagated to the same epoch before the master filter can combine the results. Of the four methods, only the fusion reset filter must finish its master filter update before the local filters can continue.

The zero reset and rescale implementations do not require the local filters to be propagated to the same time since a recursive filter is running in the master filter. The zero reset filter provides the master filter with all of the knowledge gained since the last master filter update cycle, while the rescale filter provides some of the information and retains some according to a sharing fraction.

The equations for the fusion reset master filter operation are presented below, but the equations for the other implementations are given in Carlson (1988). The local filters simply follow the Kalman filter equations given previously. With the subscript 'c' indicating common states and 'e' indicating sensor error states, the LF covariance matrices for sensor i can be given as

$$P_i = \begin{bmatrix} P_{c,i} & P_{ce,i} \\ P_{ec,i} & P_{e,i} \end{bmatrix} \tag{4.11}$$

$$x_i = \begin{bmatrix} x_{c,i} \\ x_{e,i} \end{bmatrix} \tag{4.12}$$

with n local filters and sharing fractions satisfying

$$\sum_{i=1}^{n} \frac{1}{\gamma_i} = 1 \tag{4.13}$$

The master filter solution is

$$P_m = \left[\sum_{i=1}^{n} P_{c,i}^{-1} \right]^{-1} \tag{4.14}$$

$$x_m = P_m \sum_{i=1}^{n} P_{c,i}^{-1} x_{c,i} \tag{4.15}$$

and the reset of the LFs can be expressed as

$$x_i = \begin{bmatrix} x_m \\ x_{e,i} + P_{ec,i} P_{c,i}^{-1} \left(x_m - x_{c,i} \right) \end{bmatrix} \tag{4.16}$$

$$P_i = \begin{bmatrix} \gamma_i P_m & \gamma_i P_m P_{c,i}^{-1} P_{ce,i} \\ \gamma_i P_{ec,i} P_{c,i}^{-1} P_m & P_{e,i} - P_{ec,i} P_{c,i}^{-1} \left(P_{c,i} - \gamma_i P_m \right) P_{c,i}^{-1} P_{ce,i} \end{bmatrix} \tag{4.17}$$

The process noises for the common states (and the initial variance) are scaled by the sharing fractions also.

As mentioned previously, the decentralized filter has FDIR (fault detection, identification and recovery) properties that often outperform a centralized filter (Gao et al. 1993). This arises from the fact that faulty information from one local filter is not allowed to corrupt other filters (except in the fusion reset mode). Also, small drifting errors can be allowed to accumulate in the local filters before the information is fused. The individual measurements may not be detected as blunders, but during the master filter update, the residual from a local filter may be recognized as a blunder. The fusion reset and zero reset formulations are the least fault tolerant since fused information is returned to the local filters, possibly corrupting them all. Every local filter is given the best estimate of the state variable each cycle of the fusion reset master filter, however, allowing optimal operation of the extended Kalman filter described below. The no reset mode is the most fault tolerant since the information is stored completely in the local filters and there is no feedback. If a fault is detected in one sensor, a solution without any influence from the faulty sensor, i.e. recovery, can be computed by simply skipping the failed sensor in the master cycle.

Centralized and federated filters with the same states are equivalent if there is no correlation between the errors in the local filters. Under this restriction the process becomes similar to the summation of normals algorithm from LS. For an extended Kalman filter, the master filter must be updated from local filters before a second local filter measurement is processed.

4.3 LINEARIZED KALMAN FILTER

Nonlinear models for system dynamics and measurements are often found in navigation, for example the satellite ranging problem. The Kalman filter is naturally a linear filter, so extensions have been developed to handle these problems. The first order Taylor series

linearized forms of the filter equations will be presented, but forms making use of nonlinear Itô calculus have also been developed (Maybeck 1994).

Nonlinear measurement models are common in satellite positioning from the satellite ranging model. A linear perturbation model is used to handle a nonlinear measurement, just as in LS. Given a nonlinear measurement model (Brown and Hwang 1992)

$$z = h(x,t) + v(t) \tag{4.18}$$

the first order Taylor series expansion gives

$$z = h(x_0,t) + \left.\frac{\partial h}{\partial x}\right|_{x=x_0} \hat{\delta} + v(t) \tag{4.19}$$

with misclosure

$$w_k = h(x_0,t_k) - z_k \tag{4.20}$$

The partial derivative matrix can be termed H, giving way to a perturbation filter described by the following prediction equations

$$\hat{\delta}_{k+1}^- = \Phi_k \hat{\delta}_k^+ \tag{4.21}$$

$$P_{k+1}^- = \Phi_k P_k^+ \Phi_k^T + Q_k \tag{4.22}$$

and update equations

$$K_k = P_k^- H_k^T \left(H_k P_k^- H_k^T + R_k\right)^{-1} \tag{4.23}$$

$$\hat{\delta}_k^+ = \hat{\delta}_k^- - K_k \left(w_k - H_k \hat{\delta}_k^-\right) \tag{4.24}$$

$$P_k^+ = P_k^- - K_k H_k P_k^- \tag{4.25}$$

A reference trajectory is needed to give the expansion point (x_0) about which to linearize. This can be from an external source, such as a predicted missile trajectory, giving rise to a linearized Kalman filter. The reference may also be generated from the history of the filter itself, giving rise to an EKF (extended Kalman filter). The choice between the two implementations depends on how predictable the dynamics are. In the case of a navigation filter, the dynamics are seldom predictable, so an EKF must be used. With an EKF, the perturbation term is applied to the state estimate, thereby reducing the predicted perturbation term to zero. It then becomes simple to write the filter in terms of state estimates instead of perturbations. The prediction equations are the same as the standard Kalman filter, as is the covariance update, and the state update equation is given by

$$\hat{x}_k^+ = \hat{x}_k^- + K_k w_k \tag{4.26}$$

Nonlinearity of the dynamics model presents another problem, and often both dynamics and measurements are nonlinear. A typical nonlinear dynamics model is the use of speed, heading, and position in a navigation filter. In a nonlinear dynamics model, the propagation of filter states (x) must be done explicitly, using any appropriate solution for the differential equation problem.

4.4 ADAPTIVE FILTERING

Adaptive filtering refers to adjusting the filter while it is operating. Most often only the scale of the driving white noises are adjusted, but it is possible to estimate errors in the state transition matrix and deterministic input functions, as well (Maybeck 1994). Statistically significant methods can be employed to estimate these parameters, but simple heuristic methods can also be used to improve filter performance. For example, the heading of a vehicle should be constant while it is stationary, but measurement errors or filter design can lead to the heading changing slowly over time. This phenomenon can be controlled through an adaptive procedure by fixing heading parameters or adjusting process noises. The adaptive elements of the filtering strategy used in this research are described after the dynamics model.

4.5 FILTER DESIGN

The first filter design attempted was a federated filter with separate filters for each sensor. It was quickly decided that this complicated design had too many shortcomings. Firstly, the heading and speed information are both important for effective prediction of the position states, so this information must be shared quickly to all of the filters doing propagation. If a fusion reset mode master filter was run after almost every measurement cycle to handle this limitation, little is saved in processing burden, and nothing is gained in fault detection. A centralized filter was therefore chosen for the final design. Another design was investigated which would operate like a GPS/INS integration scheme, whereby the dead reckoning would be corrected by GPS in a feedback loop. No particular advantage to this alternate formulation was identified, so it was not implemented.

A model with heading and speed was chosen over the standard constant velocity model (position and velocity in orthogonal directions) because the measurement models became linear. Heading information is difficult to integrate into the standard constant velocity model, since the degree of nonlinearity is high, and heading information becomes useless at low speeds. Two different models were employed throughout the development of the filter. The first was a constant heading and speed model, and the second assumed a constant heading rate and speed. Both were concerned with horizontal position only; vertical filtering could have been implemented easily, but is not required for most navigation problems.

The states for the heading rate model were

$$x = \begin{bmatrix} \textit{North Position}\,(m) \\ \textit{East Position}\,(m) \\ \textit{Heading Rate}\,(rad/s) \\ \textit{Heading}\,(rad) \\ \textit{Speed}\,(m/s) \end{bmatrix} = \begin{bmatrix} \varphi \\ \lambda \\ \dot{\Psi} \\ \Psi \\ s \end{bmatrix} \tag{4.27}$$

with the position estimates being latitude and longitude scaled into metres. The headings were defined with the origin East and counter-clockwise positive. The state transition matrix, Φ, can be derived from the dynamics model given

$$\varphi_{k+1} = \varphi_k + s_k \sin \Psi_k \Delta t \tag{4.28}$$

$$\lambda_{k+1} = \lambda_k + s_k \cos \Psi_k \Delta t \tag{4.29}$$

$$\Phi_{k,k+1} = \begin{bmatrix} 1 & 0 & 0 & a\Delta t & b\Delta t \\ 0 & 1 & 0 & c\Delta t & d\Delta t \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & \Delta t & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \tag{4.30}$$

$$a = s \cos \Psi_k \tag{4.31}$$

$$b = \sin \Psi_k \tag{4.32}$$

$$c = -s \sin \Psi_k \tag{4.33}$$

$$d = \cos(\Psi_k) \tag{4.34}$$

The constant heading model would simply drop the third column and row from the state transition matrix. This state transition matrix is used for propagating the covariance, but the dynamics model from Equations 4.28 and 4.29 is used to propagate the state vector. The heading is assumed constant over the prediction interval, ignoring the heading change that would result from a non-zero heading rate. This approximation error is depicted in Figure 4.3, but was found to be insignificant. The radius of the circle can be determined from

$$r = \frac{s}{\dot{\Psi}} \tag{4.35}$$

and the along track and across track errors over an interval Δt are

$$d_c = r(1-\cos\theta) \tag{4.36}$$

$$d_a = s\Delta t - r\sin\theta \tag{4.37}$$

Since dead reckoning data is processed at 10 Hz or more, the propagation times and errors are very small. Taking a sharp corner at 5 m/s and 20°/s rotation creates about the worst error possible, which is only 8.7 mm across track and 0.4 mm along track per 10 Hz update interval. On a 90° turn, the total error would amount to less than 20 cm when allowing for gradual entry and exit from the turn. This error computation could be applied, or the prediction could be split into smaller intervals to increase accuracy.

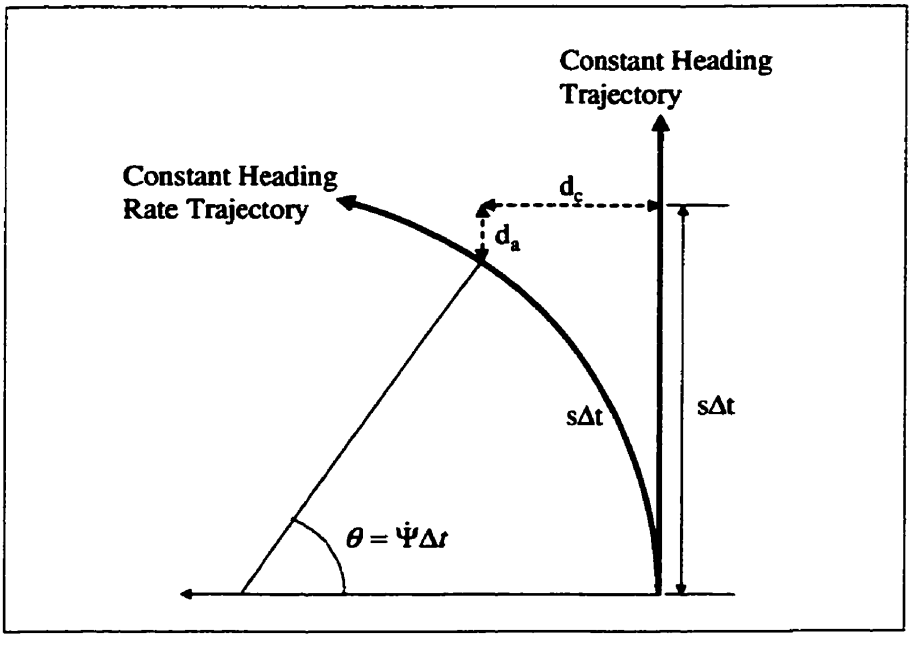

Figure 4.3 State Prediction Error

The Q matrix can be computed from (Gelb 1974)

$$Q_k = E\{w_k w_k^T\} \tag{4.38}$$

$$Q_k = \int_{t_k}^{t_{k+1}} \int_{t_k}^{t_{k+1}} \Phi(t_{k+1},\xi)\, G(\xi)\, E\{u(\xi)u^T(\eta)\} G^T(\eta)\, \Phi^T(t_{k+1},\eta)\, d\xi d\eta \tag{4.39}$$

with

$$E\{u(\xi)u^T(\eta)\} = \delta(\xi-\eta) \tag{4.40}$$

If the state transition matrix can be written in terms of Δt, the process noise matrix becomes

$$Q_k = \int_0^{\Delta t} \Phi(t_{k+1},\xi) q_k \Phi^T(t_{k+1},\eta) d\xi \tag{4.41}$$

where q_k is a diagonal matrix of spectral densities. Evaluating with a,b,c,d given above

$$Q(1,1) = q(1,1)\Delta t + \frac{a^2}{3} q(4,4)\Delta t^3 + \frac{b^2}{3} q(5,5)\Delta t^3 \tag{4.42}$$

$$Q(2,1) = Q(1,2) = \frac{ac}{3} q(4,4)\Delta t^2 + \frac{bd}{3} q(5,5)\Delta t^2 \tag{4.43}$$

$$Q(4,1) = Q(1,4) = \frac{a}{2} q(4,4)\Delta t^2 \tag{4.44}$$

$$Q(5,1) = Q(1,5) = \frac{b}{2} q(5,5)\Delta t^2 \tag{4.45}$$

$$Q(2,2) = q(2,2)\Delta t + \frac{c^2}{3} q(4,4)\Delta t^3 + \frac{d^2}{3} q(5,5)\Delta t^3 \tag{4.46}$$

$$Q(4,2) = Q(2,4) = \frac{c}{2} q(4,4)\Delta t^2 \tag{4.47}$$

$$Q(5,2) = Q(2,5) = \frac{d}{2} q(5,5)\Delta t^2 \tag{4.48}$$

$$Q(3,3) = q(4,4)\Delta t + \frac{1}{3}q(3,3)\Delta t^3 \tag{4.49}$$

$$Q(4,4) = q(5,5)\Delta t \tag{4.50}$$

Not all of the five possible spectral densities were employed. The main driving noises for the dynamics corresponded to the heading rate and speed. The effect of using several of the other densities was investigated to improve filter performance, but none were needed to represent the model errors effectively. The spectral densities used were

Heading Rate: $(5°/s)^2/s$

Speed : $(2\ m/s)^2/s$

For the constant heading filter, the heading rate spectral density was replaced with

Heading: (20°)2/s

The dynamics model was found to work well while moving, but heading and position would drift somewhat while stationary. To fix this anomaly, the filter is not run while the vehicle is stationary. With the constant heading model, however, the gyro heading bias must be allowed to continue to integrate while stationary. In this case, forcing speed to zero and the heading and speed variances to zero when the vehicle was stationary resulted in similar performance. The test for whether or not the vehicle was moving was based on whether either of the wheels indicated any motion during the last three measurements.

It was also recognized that the amount of heading change over an interval increased with the distance traveled, and not simply with time. A vehicle can change heading and heading rate faster as the speed increases, but at some point the increase levels off and the amount of possible heading change begins to decrease. This can be demonstrated by looking at a car turning circles in a parking lot. A slow moving car takes a longer time to complete a circle than a briskly moving car, but if the car gets up to highway speeds, its turning circle must be made very large, making it take longer to complete a circle. It was

decided to keep the heading spectral densities a conservative constant to ensure that the heading was not overly optimistic.

To minimize jitter about the trajectory, yet maintain the ability of the filter to follow sharp turns, it was also decided that two different process noises would be used for the heading rate if a turn is in progress. This procedure has been called smart stochastic modeling by Kealy (1999), although the threshold was implemented differently. The threshold to indicate turning action was set at a low 2°/s, as indicated by the raw measurements, to avoid underestimating sharp heading changes. The process noise was adjusted to $20°/s^2$ for the constant heading rate model. The gyro can provide an estimate of the heading rate accurate enough to decide whether a turn is starting, but in the final implementation the ABS derived heading rate is not accurate enough, so the larger process noise alone was used.

Because of the nonlinear dynamics model, strong correlations tended to develop between parameters of the dynamics model, especially the positions. Several times the cross correlations would exceed the diagonals in the covariance matrix and matrix inversions would fail. These were found to occur after periods of stationary vehicle operation or long travel times in a cardinal direction. The problem has been substantially controlled with the addition of the described adaptive features.

5 SYSTEM DESIGN

The sensors used in the final implementation were the ABS system of a 1998 Honda Civic, a Murata Gyrostar vibrating piezoelectric gyro and a NovAtel MiLLenium GPS receiver with an external OCXO. The Precision Navigation TCM2 digital compass was investigated, but was discarded due to the poor performance described previously. The Viatran model 246 barometric pressure transducer was also used, but failed before the project could be completed, so it will not be discussed further. Two laptop computers were used to log the data produced by the sensors. An A/D (analog to digital) converter was used to acquire and store measurements of the analog devices on a Pentium laptop. The second computer was used to log GPS data. A chart showing the interaction of data collection equipment is presented in Figure 5.1.

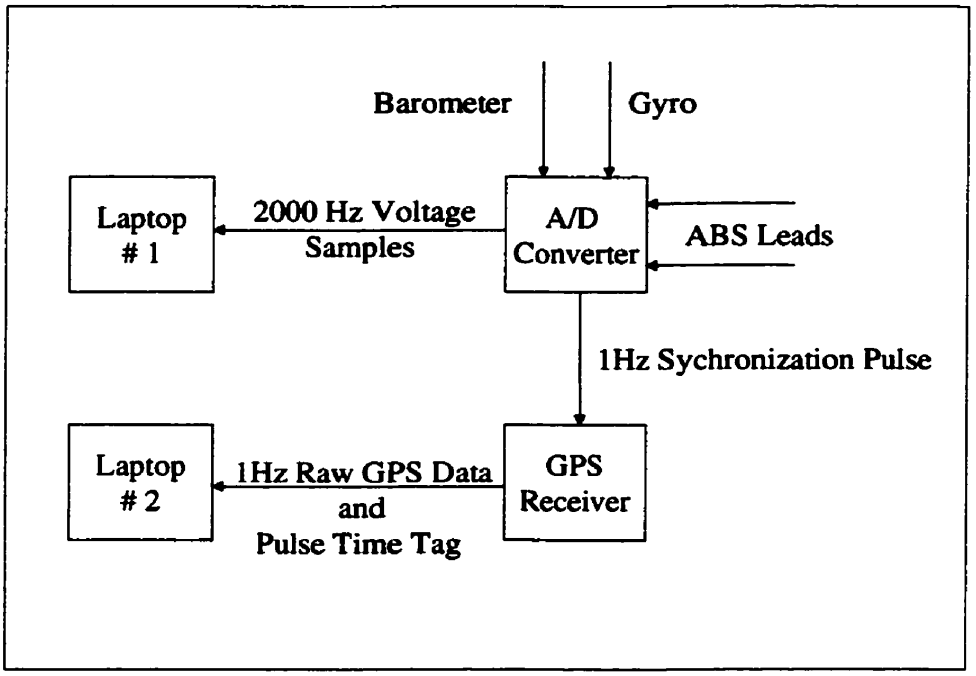

Figure 5.1 Data Collection Equipment Flow Chart

The processing was done off line, but was performed sequentially to simulate real-time performance. A flow chart depicting the flow of data to the filter is presented in Figure 5.2.

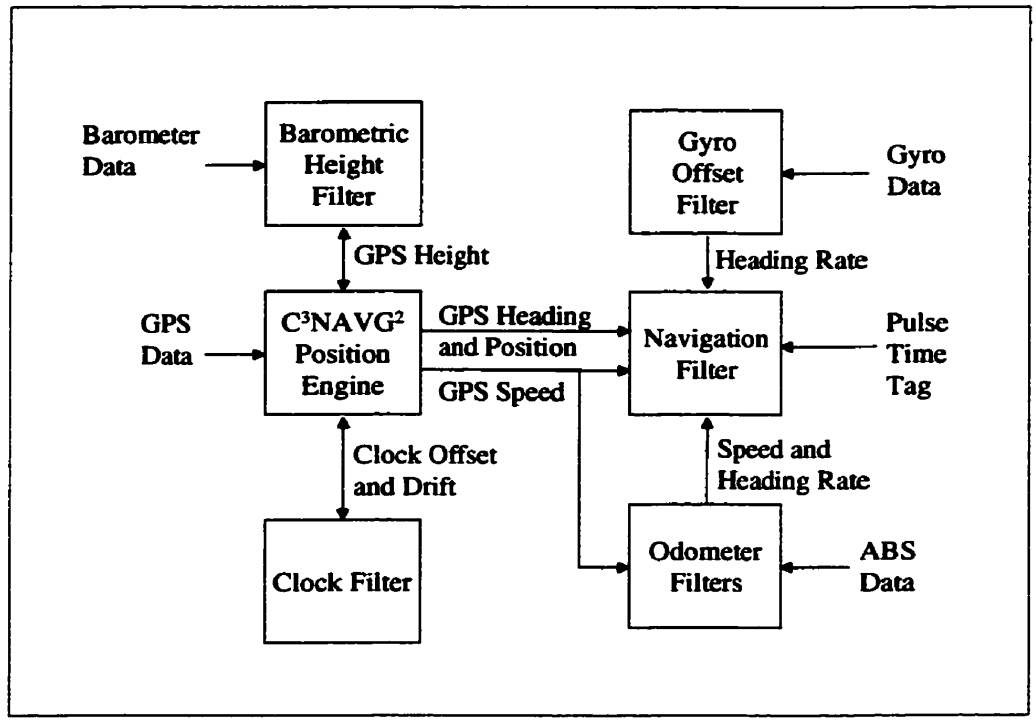

Figure 5.2 Data Processing Flow Chart

5.1 A/D CONVERTER

The analog signals from the gyro and ABS were measured using a National Instruments DAQPad MIO 16XE-50. Some specifications of the device are presented in Table 5.1. Differential input was used to minimize interference with the normal operation of the ABS. All leads and wires were formed into twisted pairs when possible, shielded with foil, and grounded at the A/D converter. All of the analog devices were low-impedance, so 100 kΩ resistors were used to ground both sides of the differential inputs (National Instruments 1992). Only the rate gyro was powered from the A/D converter 5 V supply. All other equipment was powered directly from a heavy-duty 12 V marine battery.

Table 5.1 DAQ Pad MIO 16XE-50 Specifications (National Instruments 1995)

Characteristic	Specification
Analog Input Channels	16 single or 8 differential
Signal Input Ranges	0.1 to 10 V uni- or bipolar
Resolution	16 bit
Maximum Sampling Rate	20kS/s
Settling Time	50 μs max to ± 1 LSB
Analog Outputs	2, 12 bit
Counter/Timers	2 up/down, 24 bit
Digital I/O	8 TTL/CMOS input/output
Power Supply	9-42VDC, 325 mA at 12V

The data acquisition was performed by a 'C' program, functions of which called the pre-loaded data acquisition drivers. The program called for individual samples at 16 kHz, with a scan of the channels at 2 kHz. The acquisition was driven by the internal 20 MHz oscillator, which was also used to trigger a digital pulse generator at one second intervals. This pulse was routed to the mark input pin on the NovAtel receiver. Since the two functions were driven by the same oscillator, and started by the same trigger, the pulse time-tags generated by the GPS receiver were co-incident with each 2000th sample. This was the information used to synchronize the analog data to the GPS.

5.2 ABS ODOMETRY

The factory installed ABS system of the 1998 Honda Civic is a Hall effect system, as described in the Chapter 3. As such, a sinusoidal voltage signal is generated during the rotation of the wheels. The signal varies somewhat in amplitude, but greatly in frequency. The raw voltages are sampled at two kHz using the A/D converter and stored to a file. This signal is tracked by a phase and frequency feedback digital PLL (Kaplan 1996) in post mission. This was done for the ease of sampling all of the analog output sensors near simultaneously with a single operation on the A/D converter. It also allowed for simple synchronization of the analog data and GPS data, and the ability to make measurements on the signal phase with a possible resolution of several degrees. Other signals were available at the data bus, including a brake strength indicator, reverse indicator, and park

brake indicator. These could be useful in interpreting the data, e.g. when to calibrate and whether the measured speed was in the forward or reverse direction.

The PLL was implemented as follows (Thomas 1991)

$$I_P = \sum_{t=1}^{PreSize} \sin(f * t + \theta) \tag{5.1}$$

$$Q_P = \sum_{t=1}^{PreSize} \cos(f * t + \theta) \tag{5.2}$$

$$\delta\theta = \tan^{-1}(Q_P / I_P) \tag{5.3}$$

$$\Delta\theta_{n+1} = K_1 \delta\theta_n + K_2 \sum_{j=1}^{n} \delta\theta_j \tag{5.4}$$

$$K_1 = 4 B_L T \frac{r}{r+1} \tag{5.5}$$

$$K_2 = K_{21}^2 / r \tag{5.6}$$

$$\theta_{n+1} = \theta_n + \Delta\theta_{n+1} \tag{5.7}$$

$$f_{n+1} = f_n + \Delta\theta_{n+1} / T \tag{5.8}$$

where

f is the signal frequency

θ is the signal phase

T is the integration period

B_L is the loop parameter bandwidth

K_1, K_2 are loop filter constants

r is the damping factor

I_P,Q_P are the in-phase and quadrature sums

The integration time was chosen as 0.05 s (100 samples). The noise bandwidth was set at 6 Hz, and the damping factor at 4 (critically damped). The integration values Q_P and I_P must exceed a threshold value before the loop is considered locked. The frequency is reinitialized to 120 rad/s when tracking falls below the threshold because the vehicle has come to a stop. The frequency was also forced to the positive domain, as several times the mathematically equivalent negative frequency was tracked. If one of the PLLs loses lock, the frequency from the other wheel is used to reacquire it quickly. If both wheels lose lock, a DFFT (Discrete Fast Fourier Transform) is performed to re-acquire the frequencies (Press et. al. 1992). Loss of signal lock occurs infrequently enough that this very seldom occurs. The scaled frequency estimated for each wheel is input to the navigation system. There are approximately 30 sinusoids generated per rotation of 1.5 m circumference, giving a maximum signal frequency of about 800 Hz at 40 m/s, and a wavelength of about five cm.

The ABS speed measurement is generated from an average of the scaled wheel speeds when both are available, and from only one if there is a loss of lock. The wheel scale factors were determined by estimating an average scale factor between the GPS velocity and the frequency of each wheel. GPS speeds are only used above a threshold of five m/s. A constant gain filter was used with a gain of 0.0001 for each wheel. With a gain α, the filter is given by

$$w_s^+ = a\frac{GPS\ Speed}{ABS\ Frequency} + (1-\alpha)w_s^- \tag{5.9}$$

In early filtering implementations, accumulated heading changes derived from the accumulated PLL phase were used as the observation, i.e. a relative heading. Heading changes could only be computed if both PLLs had lock, and it was found that PLL loss of lock occurred most often going into and coming out of corners when the frequencies were changing. This is problematic since the heading would not accumulate unless both wheels were in phase lock, resulting in turning errors. Reacquisition was slower on corners, as

well, because the left and right wheels are moving at different speeds and using one wheel speed to aid lock of the other is less effective.

A sensor bias model was originally implemented for the ABS system using accumulated heading observations, as opposed to the heading rate implementation used in the final implementation. Its two parameters were a heading offset and a heading rate error as a function of distance. Unlike the gyro, which would continually accumulate heading error due to biases over time, the ABS heading error accumulates with distance traveled, i.e. only when moving. Both errors were made to be random walk variables. The state vector was

$$x = \begin{bmatrix} e_{\psi_g} \\ e_{\psi_g} \end{bmatrix} \begin{bmatrix} \textit{Heading Error (rad)} \\ \textit{Heading Error Rate (rad / m)} \end{bmatrix} \tag{5.10}$$

The state transition matrix is therefore

$$\Phi = \begin{bmatrix} 1 & s\Delta t \\ 0 & 1 \end{bmatrix} \tag{5.11}$$

and the process noise matrix

$$Q = \begin{bmatrix} q_1\Delta t + \frac{s^2}{3} q_2 \Delta t^3 & \frac{s}{2} q_2 \Delta t^2 \\ \frac{s}{2} q_2 \Delta t^2 & q_2 \Delta t \end{bmatrix} \tag{5.12}$$

With a measurement model for the integrated heading of

$$\Psi_A = \Psi + e_{\psi A} \tag{5.13}$$

The spectral densities were chosen by estimating the noise necessary to create heading drift plots such as the one shown in Figure 5.3, and then tuning the values by comparing various combinations on several data sets. The ABS heading change observation is biased

from the GPS heading, and follows an apparent random walk. The power of the random walk can be determined by statistical analysis of the change in error values over time. All of the one second intervals were taken as members of a Gaussian data set, since a random walk (Brownian motion) is the sum of sequential white noise. The power, q_1, of the noise making up the Brownian motion, β, can therefore be derived from (Maybeck 1994)

$$E\left\{\left(\frac{\beta(t+\Delta t)-\beta(t)}{\Delta t}\right)^2\right\}=\frac{q_1}{\Delta t} \tag{5.14}$$

The squares of the interval changes were put into a histogram, and the centre of the peak bin was chosen as the starting point for filter tuning. An idea of how good the ABS heading will be in the absence of GPS for calibration can be found using Figure 5.3. In the absence of other sources of heading information, the heading bias will be assumed almost constant, i.e. one could draw a line forward in time from any place on the meandering offset and see how the distance to the line changes over time. Throughout the entire test, the bias changes by over 30°, so without external correction from GPS, the heading would be in error by more than 30°.

A similar procedure, but using the change of consecutive slopes, was used to estimate the power for the heading rate Brownian motion, q_2. The spectral densities which were eventually chosen to be

$q_1 = (10^{-6}\ \text{rad})^2/\text{s}$

$q_2 = (10^{-8}\ \text{rad})^2/\text{s}$

The heading observation was given a standard deviation of 0.1°. The standard deviation of speed derived from the PLL was estimated to be 0.1 m/s for each wheel, and improves by a factor of $\sqrt{2}$ if the average of both wheels is used.

If distance is measured over an interval, the variance of the speed can be computed from the variance of the distance measurement using

$$\sigma_{speed} = \frac{\sigma_{dist}}{\Delta t} \tag{5.15}$$

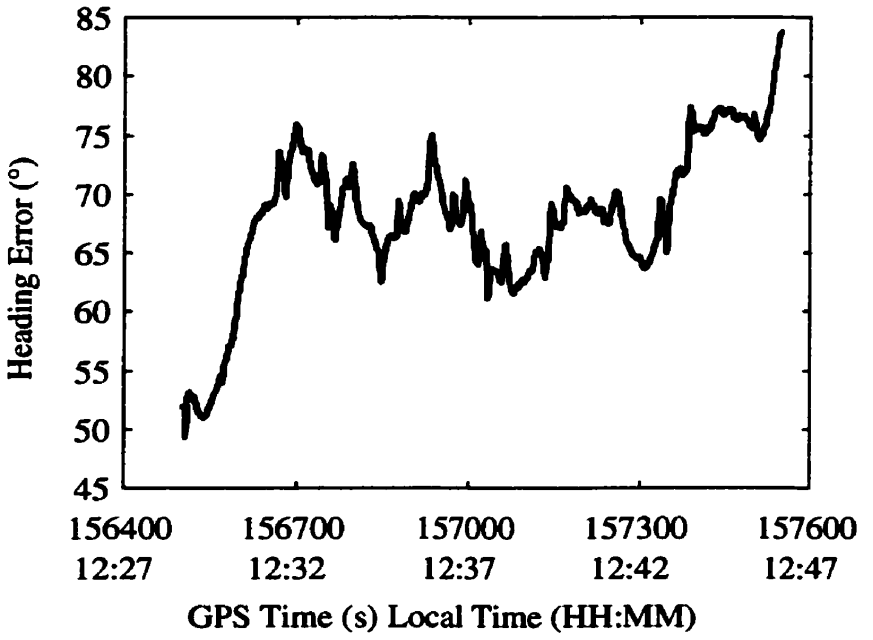

Figure 5.3 ABS Heading Error

The accumulated change model has a serious shortcoming when it comes to statistical treatment. It was found that the accumulated heading should not be treated as a relative heading since slippage or loss of lock created errors which could not be accounted for by the sensor error model. After the first few tests, it was decided that a heading rate model for the observation must be used, but the PLL was not reliable enough to provide continuous data from both wheels. Instead, heading rate observations were created by accumulating the zero crossings of the sinusoids over an interval of 0.1 s. This observation is more noisy than the one derived from the PLLs, but the observations are always available. This implementation did not require any error states, since the heading

rate is the observation, as opposed to integrated heading, and the external calibration could not be bettered by the use of centralized filter error states.

Heading rates are in discrete intervals of about 10°/s because the difference between distances traveled by the two wheels is measured in zero crossings of about 2.5 cm. The difference of the two signals creates a Doppler signal, the zero-crossings of which during the interval give the multiplier for the discrete heading rate. For example, if the vehicle was turning at 5°/s with a track width of 1.5 m, the difference in distance traveled by the two wheels is 0.13 m/s. The difference is sensed in 2.5 cm increments over the 0.1 s sampling intervals. A 10°/s heading change is observed every second interval, while no rotation is observed on the alternating intervals. This process can be thought of as breaking a turn into a series of discrete steps, as shown in Figure 5.4. The discrete nature is smoothed somewhat by the constant heading rate model, as well. A variance of 0.1°/s was chosen for the heading rate observation variance after a great deal of experimentation.

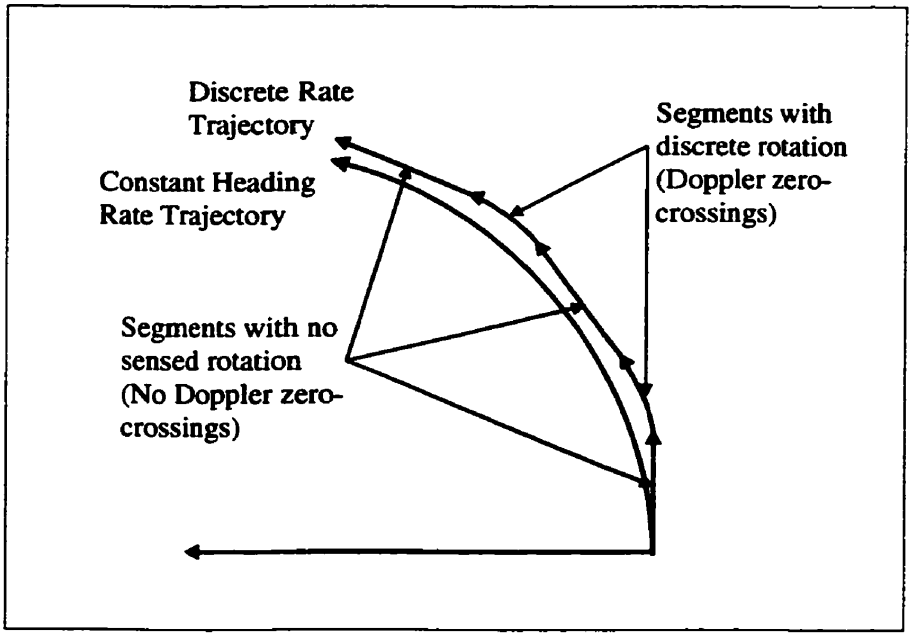

Figure 5.4 Continuous vs Discrete Heading Rate

5.3 GYRO

The Murata Gyrostar piezoelectric vibrating gyro was used for this system. The gyro was selected for its low cost (approximately $15 in large quantities), small size, and good performance relative to others in its class. The voltage output from the gyro is proportional to rotation rate, and is sampled at 2000 Hz by the A/D converter. Two different observation models were employed to use the sampled data. In the first implementation, the samples were integrated at 10 Hz to give a heading change over the interval, and accumulated to give a relative heading. In the final chosen method, the average output for 20 Hz intervals were used, i.e. average rotation rate measurements. This is a more effective way of using the information due to the dynamics correlated nature of the observation errors. The most significant errors are due to vehicle tilt and tilt rate. Two different gyro units were tested on the last test, and both were leveled with shims while the vehicle was on a level garage floor.

The gyro has instability in both the zero-rotation output voltage and the scale factor, mainly as a result of temperature variations. The specifications for the zero-rotation bias and scale factor vary by 20%, so independent calibration of each unit is required. The high temperature variations require regular calibration of these values, however difficulties arise in trying to calibrate both the zero-rotation offset and scale factor simultaneously. The zero-rotation bias can be easily determined when the unit is not rotating, i.e. when the platform is stationary, but calibration of the scale factor requires rotation of the sensor with an accurate heading reference.

The zero-rotation bias is calibrated online, but the scale factor was estimated from a short calibration procedure. The vehicle was kept stationary for sixty seconds to calibrate the zero offset, and then simply driven through 10 consecutive circles and stopped facing the same direction as the start direction. The gyro output was integrated and related to the total rotation to give a scale factor.

The zero offset is continually calibrated through a conditional filter. The filter only operates when no sensor is indicating a turn in progress, the same signal that is used to adaptively control the dynamics noise. There are then two gains depending on whether the vehicle is stationary or not. It is not sufficient to calibrate the offset only when the vehicle is stationary, since the offset may change substantially on freeway trips where there is little or no stopping. When the vehicle is moving, however, even straight travel is composed of many small adjustments to the vehicle heading, so a small gain is used on the calibration filter. An assumption is made that small left and right turns during the periods where the vehicle is not deemed to be turning will cancel out. The same style of digital LPF is used as for the ABS wheel scale factor, described above. The gain (α) used while the vehicle was static was 10^{-4}, while driving straight it was 10^{-6}, and no calibration was performed when the vehicle appeared to be turning.

One further aspect that could not be explored fully was temperature stability. If the calibration parameters could be shown to be constant at a particular temperature, a temperature profile of values could be used to correct the output. No controlled temperature environment was available to test whether this technique would be feasible.

Because the zero offset can be calibrated quite well, only one error state is used for the gyro. A first order Markov decaying state accounts for errors caused by variability of the road surface, causing vehicle tilt errors with a fast decay. It was difficult to quantify this error since it involves the variability of the error, and not the bias as in the ABS heading. A short period of data in which the vehicle heading should have remained constant is shown in Figure 5.5. An effort was made to keep the steering wheel still for several seconds to look at the result of driving on uneven road. The time constant for the exponential decay was estimated to be about 0.6 s. The first order Markov function is graphed along with the sample data autocorrelation in Figure 5.6. Some periodic noise is found throughout the data, most likely due to vibration. The observation noise was chosen to be 0.01°/s.

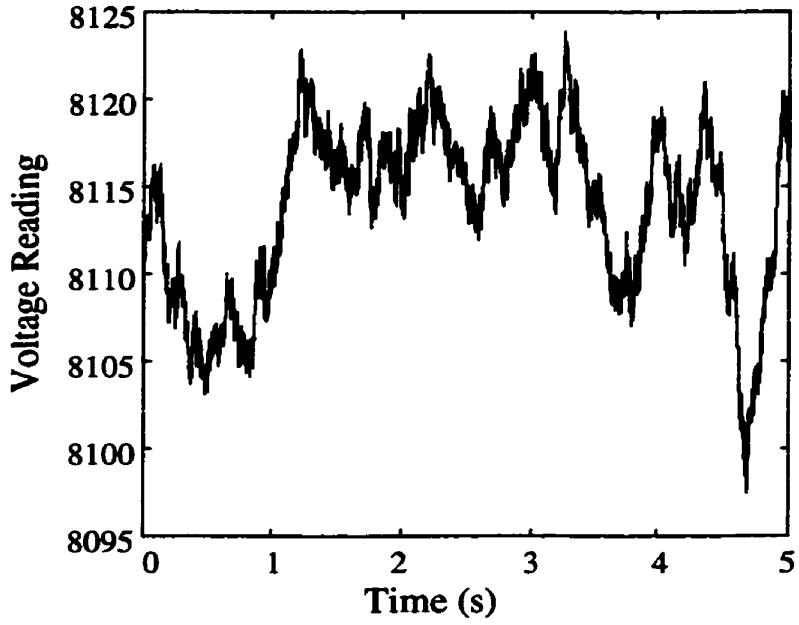

Figure 5.5 Sample Gyro Output while Driving Straight

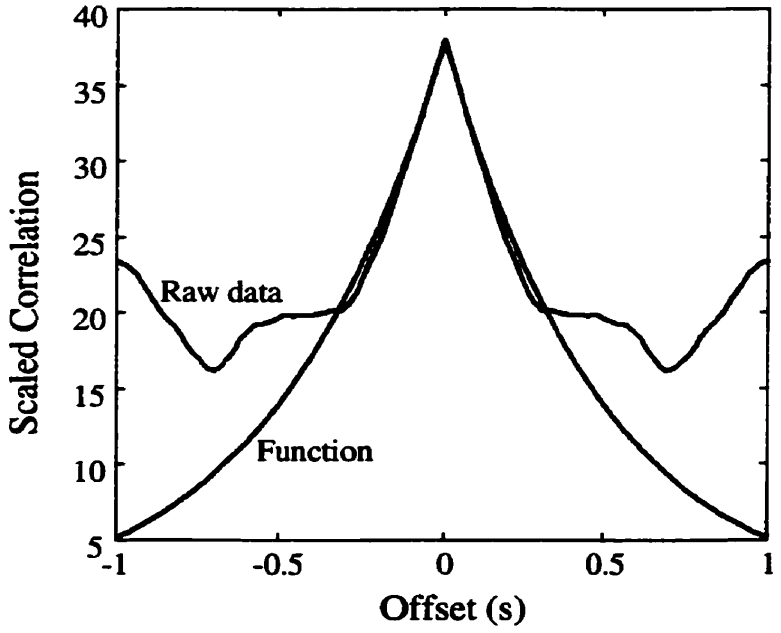

Figure 5.6 Data and First Order Markov Autocorrelation

The measurement model for the gyro is given by

$$\dot{\Psi}_G = \dot{\Psi} + e_{\dot{\psi}G} \tag{5.16}$$

and the state transition matrix for the error is given by

$$\Phi = \left[1 - e^{-\Delta t/\tau_m}\right] \tag{5.17}$$

and the process noise matrix

$$Q = \sigma^2 \frac{T_m}{2}\left(1 - e^{-2\Delta t/\tau_m}\right) \tag{5.18}$$

5.4 GPS

The system employs NovAtel MiLLenium receivers for the reference station and remote GNSS data. The remote receiver was used in a wide correlator mode to somewhat simulate low cost GPS sensor performance. Low cost equipment would no doubt be used in a mass produced automotive navigation system. In a real-time implementation, single point GPS positioning would be used, or differential corrections would be obtained from a satellite system, such as WAAS or Racal LandStar, or a surface provider, such as the 300 kHz marine navigation service, an FM sub-carrier system, or cellular service. However, with SA turned off, most users will not need differential corrections to meet the accuracy requirements needed for the pre-map matched solution. GPS solutions were post-processed using the University's C^3NAVG^2 GPS/GLONASS single-difference carrier-smoothed code processing software with clock and height aiding.

The decision was made to keep the GPS position computation external to the integration filter for several reasons. First of all, it allows flexibility in that any type of GPS receiver that can output digital positions can be used. Second, it was decided that the use of low cost dead reckoning sensors does not provide a good basis for the rejection of individual

pseudorange errors. The dead reckoning solutions are continually drifting, and sometimes errors occur which can not be modeled, such as wheel slippage. The filter statistics while coasting on dead reckoning can therefore be inaccurate, and should not be relied on to detect ranging errors.

GPS positions are therefore computed using C^3NAVG^2, and are used if the GDOP < 5. If this condition fails, constraints are added sequentially until the GDOP passes or all of the constraints have been used. If both height and clock constraints are used, height is added first, followed by clock coasting if necessary. Other reliability measures, such as MHE, could have also been given a threshold so that potentially poor GPS observations would not be used.

GPS velocity components in Northing and Easting were used to generate speed and heading observations. Vertical velocity was not included, since it tends to be noisy and is generally very small. Even on a 7% grade, the speed error would amounts to less than 0.25%. GPS speed was calculated by

$$s_g = \sqrt{v_n^2 + v_e^2} \tag{5.19}$$

where v_n and v_e are the North and East velocity component. The heading is consistent with the defined coordinate axes when computed using

$$\Psi_g = \tan^{-1}\left(v_n / v_e\right) \tag{5.20}$$

The variance of the GPS speed observation is derived using the covariance law as

$$\sigma_{sg}^2 = \sigma_{vn}^2 \frac{v_n^2}{s^2} + \sigma_{ve}^2 \frac{v_e^3}{s^2} \tag{5.21}$$

and the GPS heading observation variance is

$$\sigma_{\Psi g}^2 = \sigma_{vn}^2 \frac{1}{v_e^2\left(1+\frac{v_n^2}{v_e^2}\right)^2} + \sigma_{ve}^2 \frac{v_n^2}{s^4} \tag{5.22}$$

To save from processing observations that contain almost no information, speed and heading observations are only used if the speed is greater than three m/s. GPS heading and position observations were both checked for blunders. This process in similar to the LS data snooping described in Chapter 2, but the variance of the observation and the propagated solution (dynamics model) must both be taken into consideration when testing the residual (innovation sequence). The residual is given by

$$r = z - Hx \tag{5.23}$$

and the covariance of the residual is given by

$$P_r = HP_xH^T + R \tag{5.24}$$

Each residual is then tested against the threshold

$$r_i / P_{rii} < 4.57 \tag{5.25}$$

The first two consecutive blunders are ignored, but the next observation following two blunders is used, regardless of whether it passes the test. This allows a biased trajectory to be snapped back, but minimizes the chance of a blunder entering the solution.

6 SYSTEM TESTING

Three sequential data collection and processing campaigns were performed to evaluate the system performance under increasingly challenging conditions. The first two tests were open sky tests, and the third test was done under high masking conditions in Calgary's downtown core. All of the tests were done in the presence of SA, so differential GPS was used. The Galileo geographic information system from Geodyssey was used to plot the vehicle trajectories.

6.1 STRATHCONA/COACH HILL TESTING

Two sample data sets were collected on the afternoon of January 10th, 2000, in the residential areas of Strathcona and Coach Hill in Calgary. These areas both have vertical elevation changes of over 100 m and are located about 5 km from the University of Calgary, where the reference receiver was operated. At the time, it was about –5° C and the roads were partially covered with snow and ice. Gravel is used extensively on Calgary roads to improve traction during the winter months. Speeds during the tests reached approximately 60 km/h with few stops. Reasonable care was taken not to spin the tires, as ABS derived speed and especially heading data would be degraded. Both tests lasted about 20 minutes with several minutes between tests. Initial values for the calibration numbers were obtained from previous test data, and the values at the end of the first test were carried over to the second.

GPS positions were available throughout the test periods, although on a few occasions only four satellites were in view. GPS outages were therefore simulated by removing GPS data, alternating every 100 s. A second test with alternating 100 s availability and 200 s outages was then analyzed. Two modes of operation were tested: ABS differential odometry and ABS odometry with gyro derived heading. The tests were designed to create total masking for extended periods on slippery roads with large vertical variation, thereby testing the capabilities of the dead reckoning sensors. The gyro could suffer

errors due to the vertical variation causing pitch and roll, and differential odometry errors could result from wheel slippage due to the slippery roads. The inclinometer unit of the TCM2 digital compass was originally to be used to correct the gyro data for inclination, but after the inclinometer performed badly, it was not used. The barometer was to be used to monitor elevation changes, but it was poorly mounted so the data was not usable. There is very little data available to calibrate the dead reckoning sensors before the first simulated outage, so the test simulates a cold start of the system near an area with high masking conditions.

The GPS heading is shown with the integrated gyro output for the first test in Figure 6.1. A constant estimate of the gyro zero-rotation bias is subtracted from the gyro output, and the difference is scaled with the gyro scale factor and summed. The GPS heading is computed from velocity components when the speed is above three m/s. The difference between the two is shown in Figure 6.2. Gyro zero-rotation errors create a linear drift in the difference. The graph appears to be composed of constant slopes due to slight variations in the offset, as well as irregular errors due to pitch and roll of the vehicle and vibration and acceleration effects on the sensor. The scale factor does not appear to be a major source of error, since adjustment of the scale factor by several percent caused the heading range of bias values to increase.

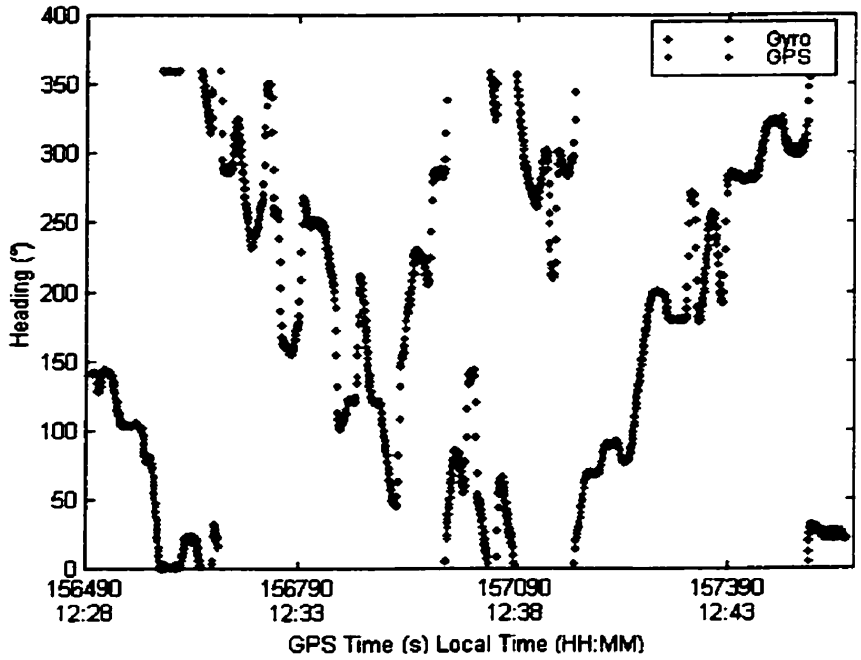

Figure 6.1 Strathcona Test 1 - GPS and Integrated Gyro Headings

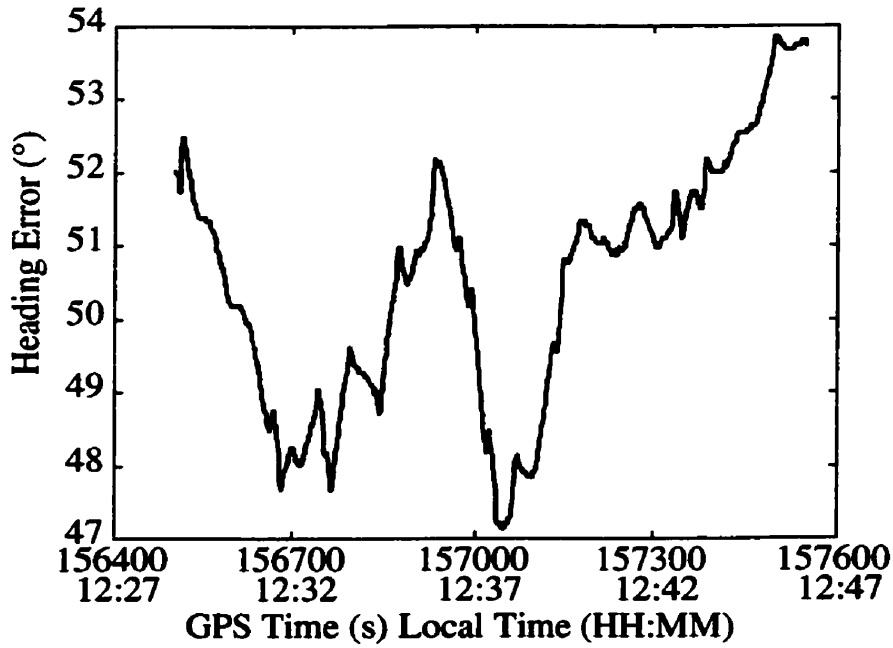

Figure 6.2 Strathcona Test 1 - Integrated Gyro Heading Bias

Heading error plots of this type were used extensively to tune both the calibration and navigation filters. It was during this tuning procedure that the shortcomings of using integrated heading observations were discovered. It was very difficult to tune the navigation filter such that a consistent level of performance could be achieved. Initial adjustment of the error state and dynamics noise spectral densities led to reasonably good filter position estimates, but optimistic statistics. These statistics occasionally led to the discarding good GPS solutions. A choice of values that better represented the statistics of the filter caused sensor biases to absorb legitimate dynamics and degraded the position results. Different choices of the number and type of filter states were adopted, but rarely could enough information be acquired to successfully settle the sensor bias states. Sensor bias states are often problematic for a Kalman filter, especially when there are long periods with no external data to improve the sensor bias state estimates. Heading rate observations were used to address this problem, as described in Chapter 5.

The ABS calibration of the wheel speeds can be verified by examining the difference between GPS speed estimates and wheel speed estimates. These errors have a zero mean, so there does not appear to be substantial errors in the wheel scale factors. Some fluctuation is to be expected, since GPS velocity observations from Doppler measurements have a noise of about 5 cm/s, and wheel speeds vary while turning. The errors and corresponding statistics are found in Figure 6.3 and Table 6.1.

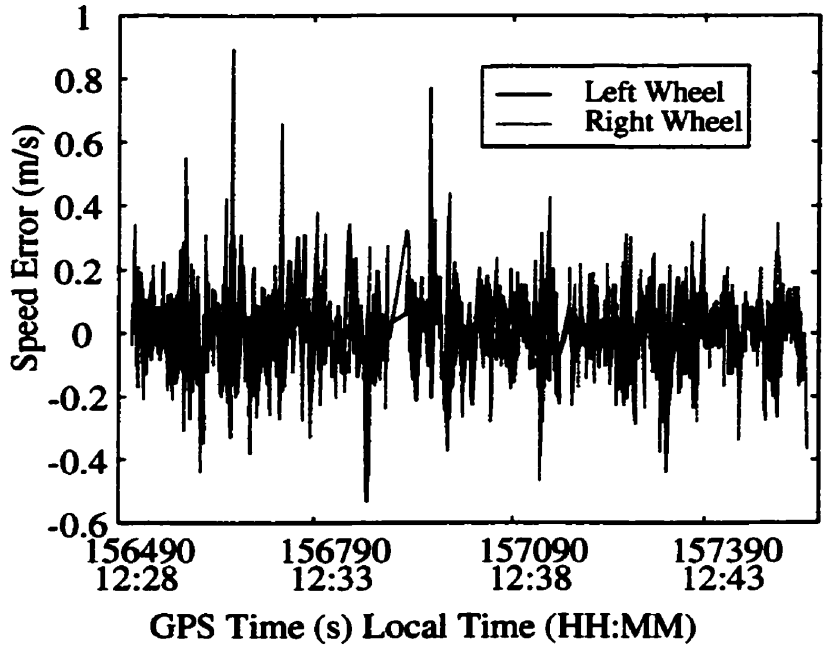

Figure 6.3 Strathcona Test 1 - Wheel Speed Errors

Table 6.1 Strathcona Test 1 - Wheel Speed Error Statistics (m/s)

Wheel	Max	Min	Mean	Std. Dev.
Left	0.77	-0.53	0.00	0.13
Right	0.90	-0.47	0.00	0.13

The wheel speeds are integrated to give a relative heading. As for the gyro, this should be a constant bias, but will drift due to errors in the wheel scale factor. Slippage of a single wheel will appear as a sharp change in the heading bias. The ABS heading bias for Test 1 and Test 2 are shown in Figure 6.4 and Figure 6.5. Again, there does not appear to be a large error in the wheel scale factors, but the wheel scale factors did seem to vary more than expected. With the wheel scale factor filter disabled, different fixed wheel scale factors were found to fit the heading data well for periods up to several minutes, but no wheel scale factor performs well throughout the test.

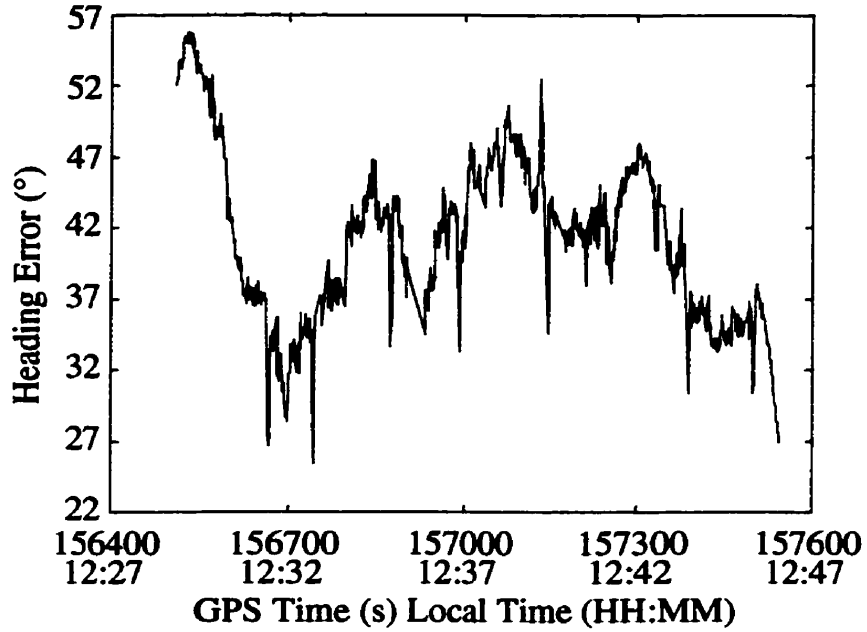

Figure 6.4 Strathcona Test 1 - Integrated Differential ABS Heading Bias

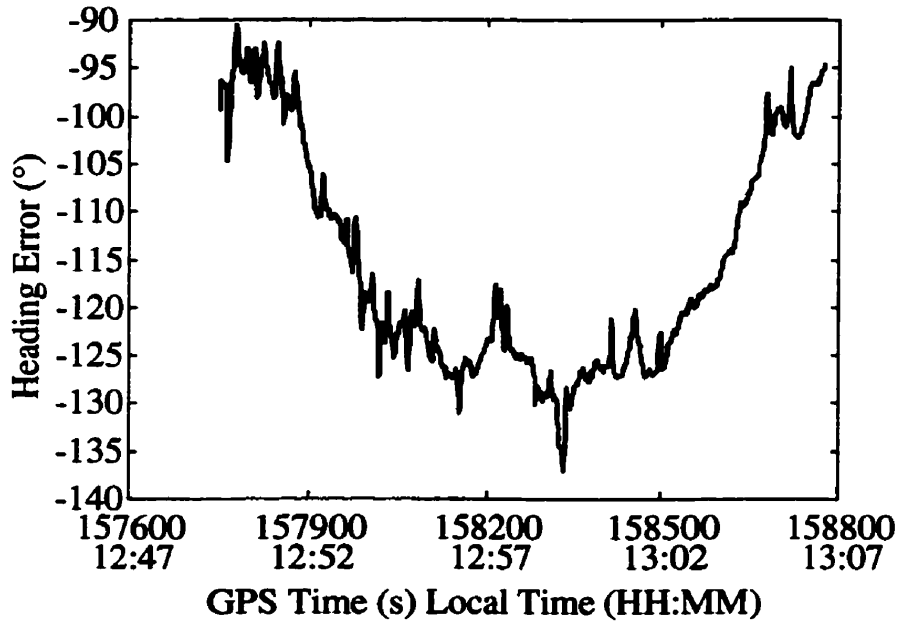

Figure 6.5 Strathcona Test 2 - Integrated Differential ABS Heading Bias

Figure 6.6 and Figure 6.8 show the trajectory results of processing with differential ABS for heading and speed, and GPS with 100 s GPS outages. The blue dots are the navigation filter positions, and the larger red dots are the GPS available to the filter. The corresponding results using Gyro heading information with ABS odometry for speed are shown in Figure 6.7 and Figure 6.9. The results for 200 s GPS outages are found in the Appendix, but error statistics are presented for all of the tests in Table 6.2. The reference trajectory was a floating ambiguity carrier phase solution created by the University of Calgary's FLYKIN software. The reference trajectory should be accurate to better than 1 m throughout the tests since the agreement between forward and reverse passes was better than 0.5 m.

Table 6.2 Strathcona Tests Filter Position Error Statistics

	Outage (s)	Differential ABS		Gyro and Odometry	
		Max (m)	RMS (m)	Max (m)	RMS (m)
Test #1	100	94	19	14	9
	200	116	26	61	13
Test #2	100	154	21	7	6
	200	170	29	20	7

As shown previously, the gyro heading information is superior to that from ABS differential odometry, so position domain results follow suit. The use of both gyro and differential odometry for heading information was also investigated. The gyro normally outperforms differential ABS due to the difficulty in determining the wheel scale factors, so little is gained by using both, and the errors were generally larger than using gyro alone. The combination might prove useful at times when the gyro's performance is degraded, such as when the car is inclined during turns. Differential ABS is presented as a zero sensor cost source for both speed and heading DR info, and the gyro with ABS odometry is presented as a low cost option with a measurable improvement in performance.

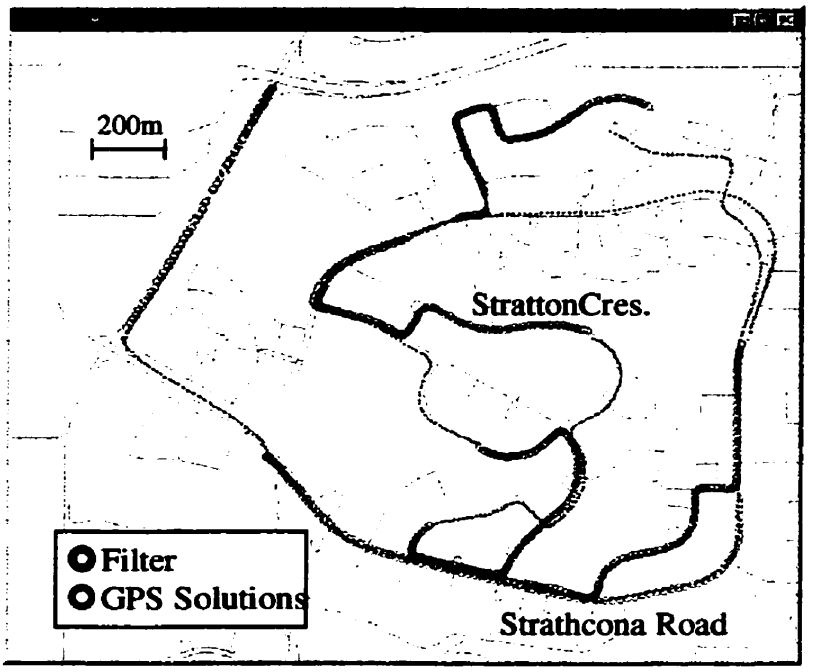

Figure 6.6 Strathcona Test 1 Trajectory – Differential ABS with 100 s GPS Masking

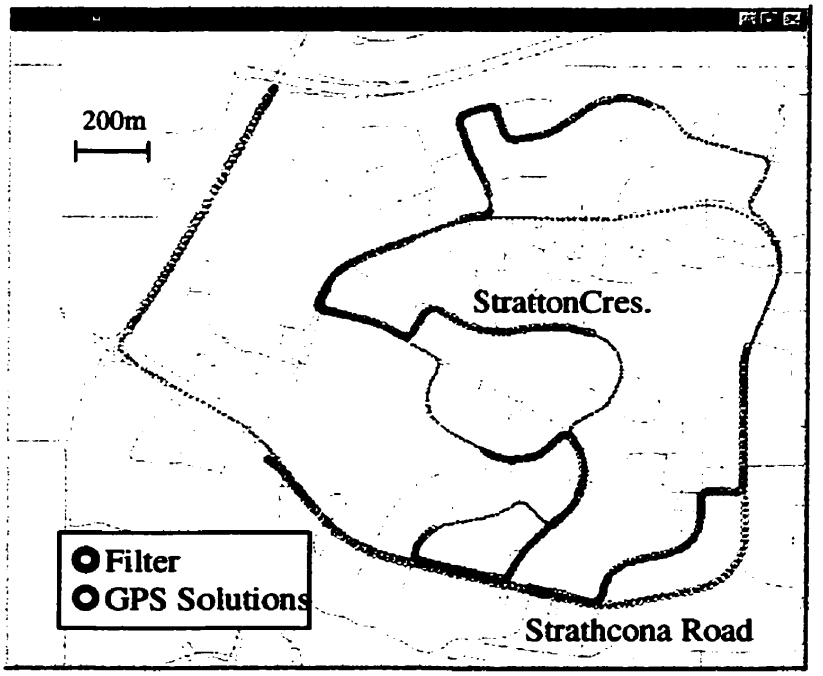

Figure 6.7 Strathcona Test 1 Trajectory – Gyro & Odometry with 100 s GPS Masking

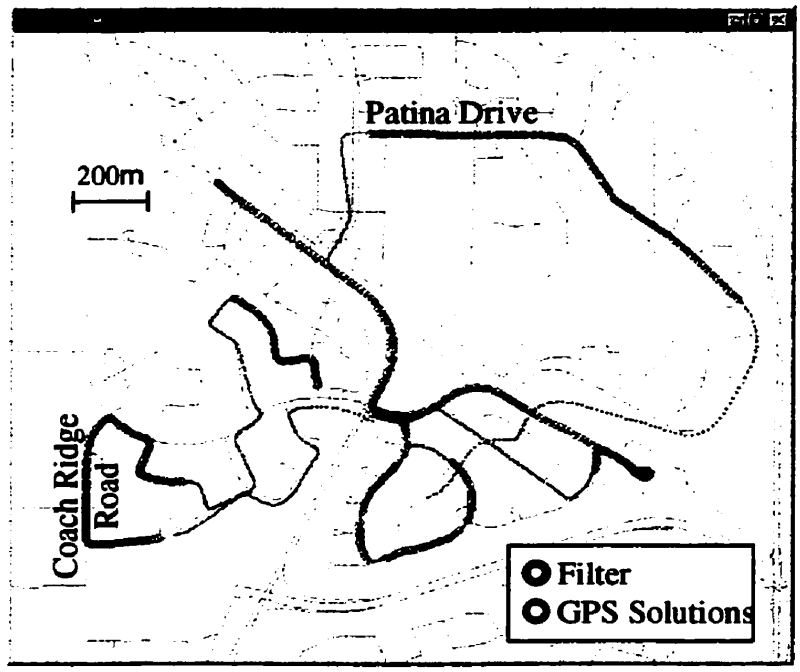

Figure 6.8 Strathcona Test 2 Trajectory – Differential ABS with 100 s GPS Masking

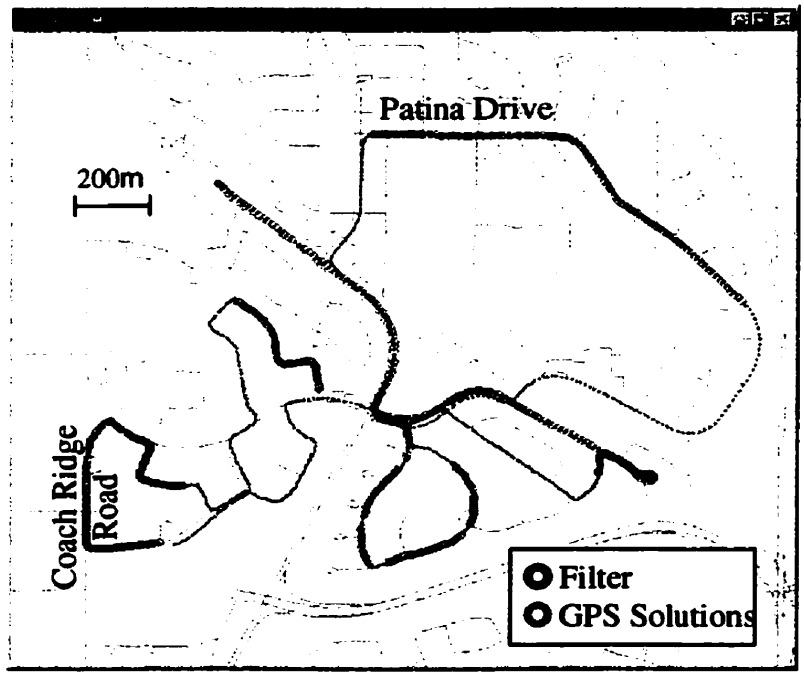

Figure 6.9 Strathcona Test 2 Trajectory - Gyro & Odometry with 100 s GPS Masking

The position and heading errors were also compared to their corresponding statistics estimated by the filter. Plots are presented for the Test 1 results in Figure 6.10 to Figure 6.13, while Test 2 results are presented in the Appendix. The figures show the errors in blue, and $\pm 2\sigma$ estimates in red. Error growth and a corresponding increase in the filter standard deviation are apparent during the GPS outages, growing more quickly for the differential ABS heading than the gyro. The error is generally contained within these bounds, which should contain about 95% of the errors, however occasional departures are noted. The statistics are generally pessimistic for the gyro results.

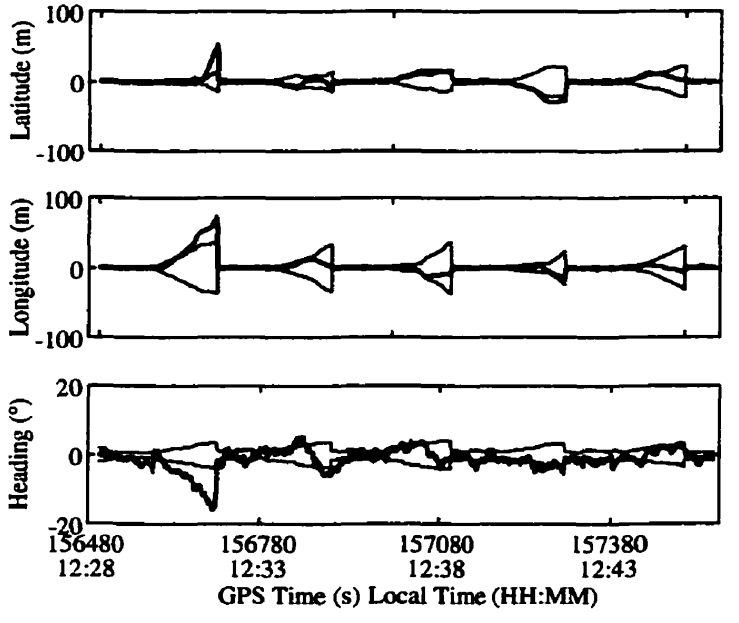

Figure 6.10 Strathcona Test 1 Errors and Estimated 2σ Bounds – Differential ABS with 100 s GPS Masking

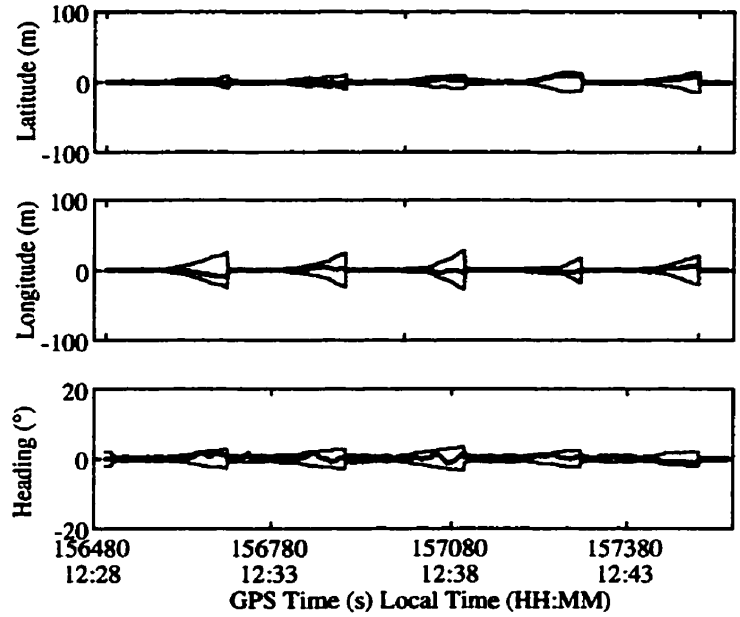

Figure 6.11 Strathcona Test 1 Errors and Estimated 2σ Bounds - Gyro & Odometry with 100 s GPS Masking

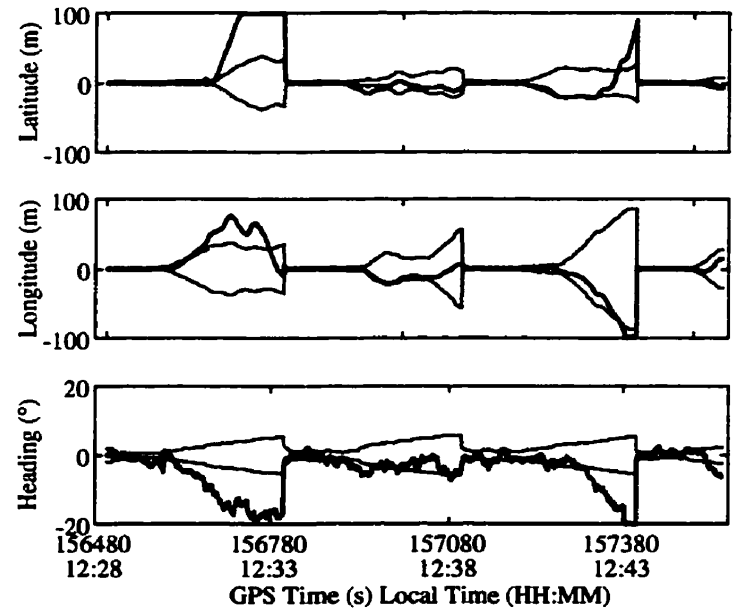

Figure 6.12 Strathcona Test 1 Errors and Estimated 2σ Bounds – Differential ABS with 200 s GPS Masking

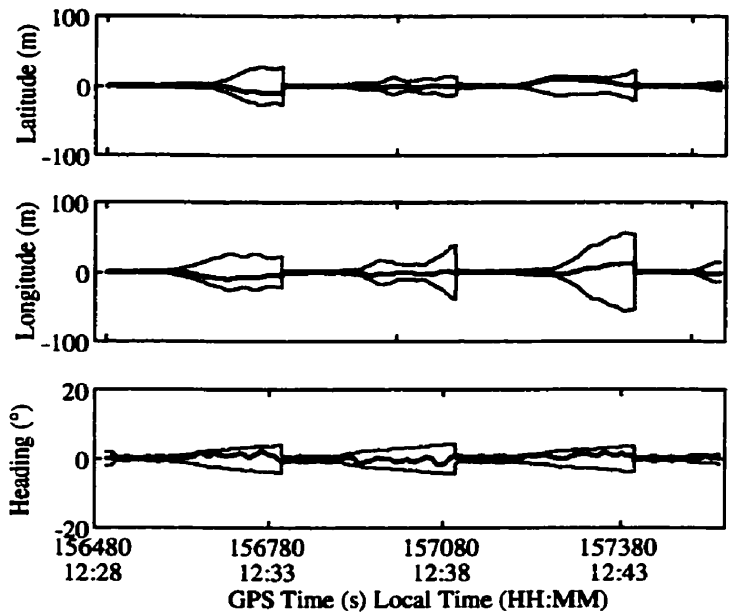

Figure 6.13 Strathcona Test 1 Errors and Estimated 2σ Bounds - Gyro & Odometry with 200 s GPS Masking

6.2 BOWNESS TESTING

The second set of tests was performed in Calgary's residential area of Bowness, on the afternoon of March 13th, 2000. At the time, it was about 3° C and the roads were partially covered with snow and ice. Two data sets were collected, each of approximately 15 minutes duration and separated by several minutes. Speeds during the tests reached approximately 50 km/h, with frequent stops. Bowness has a vertical variation of less than 10 m, so a fixed value height augmentation was used. The initial value for the gyro offset was taken from a previous test, but the wheel scale factors had to be recalibrated with some freeway driving. The wheel scale factors had changed by more than 1% because one of the tires had a slow leak repaired, and all of the tires pressures were adjusted. During the first test, a wait of over three minutes occurred at a busy intersection.

GPS positions were available throughout most of the test periods, although occasionally only two or three satellites were tracked. Signal masking similar to that found in an urban

canyon was therefore simulated. This methodology allows the performance of the filter under typical urban canyon masking to be evaluated, but a good reference trajectory was still available by not using the simulated mask. The difference between calibration values and filter states with and without masking were compared to see how much degradation occurs as a result of the lost GPS positions. The mask was generated based on the assumption of an infinite corridor of buildings on both sides of the road, with the road being 25 m wide and the buildings also 25 m high. This created a mask angle that varied from about 63° perpendicular to the road down to the horizon along the road. The mask is pictured in Figure 6.14 and can be computed using

$$\varepsilon_M = \left| \tan^{-1}\left(\frac{2h_c \cos(\alpha_v - \alpha_s)}{w_c} \right) \right| \tag{6.1}$$

where

h_c, w_c are the canyon height and width

α_v, α_s are the vehicle and satellite azimuth

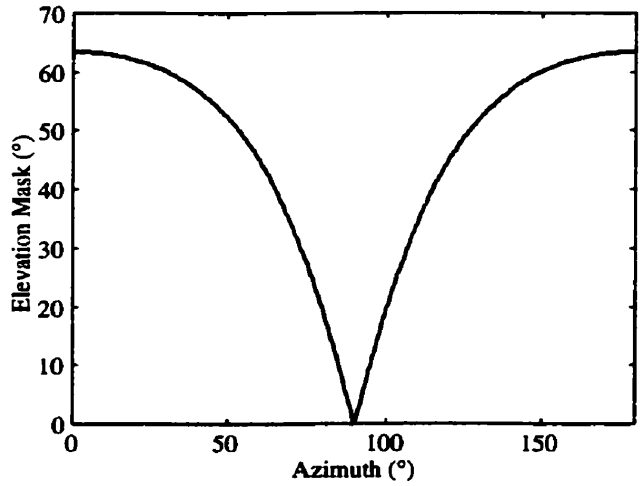

Figure 6.14 Simulated Urban Canyon Masking Profile

Three different schemes were used to augment the C^{3}NAVG^2 GPS position computation, namely fixed value height aiding, clock coasting, and the combination of the two. As

would be expected, the augmentations greatly increased the number of computed positions, sometimes because enough observations became available to compute a solution and others by improving the GDOP enough to meet the threshold. The number of positions successfully computed by C^3NAVG^2 for each augmentation is found in Table 6.3, with the theoretical number based on a rate of one position per second.

Table 6.3 Number of Computed Positions with Aiding

C^3NAVG^2 GPS Augmentation	Test 1	Test 2
Theoretical	986	996
No aiding	24	9
Height	387	260
Clock	383	305
Clock and height	397	331

Figure 6.15 and Figure 6.17 show the results of the first test using differential ABS, and Figure 6.16 and Figure 6.18 show the results using the gyro and ABS odometry. The results for the remaining augmentations and the second test are found in the Appendix. Again, it is clear that the gyro heading information outperforms ABS differential odometry. The wheel scale factors are difficult to calibrate to the accuracy needed to match the performance of the gyro. The observation variance for the gyro is much lower than that of the ABS due to wheel slippage occurrences and the discrete nature of ABS heading rate observations as described in Chapter 5. Therefore, noisy GPS headings are less heavily weighted when the gyro is used, and blunder GPS headings are more easily detected. An example of this can be seen by comparing differential ABS and gyro with no GPS augmentation. The differential ABS heading filter accepted a 10° blunder GPS heading just before GPS availability was lost at the beginning of the test (centre).

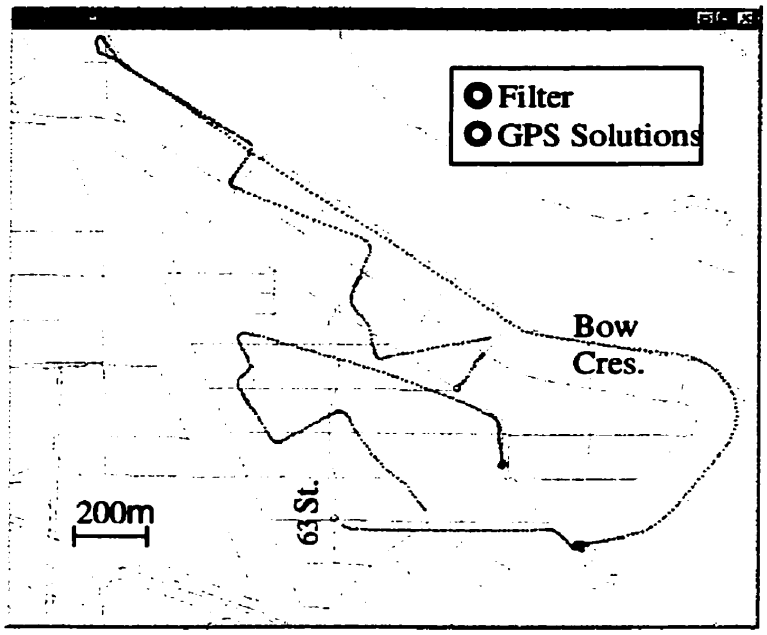

Figure 6.15 Bowness Test 1 ABS No Augmentation Trajectory

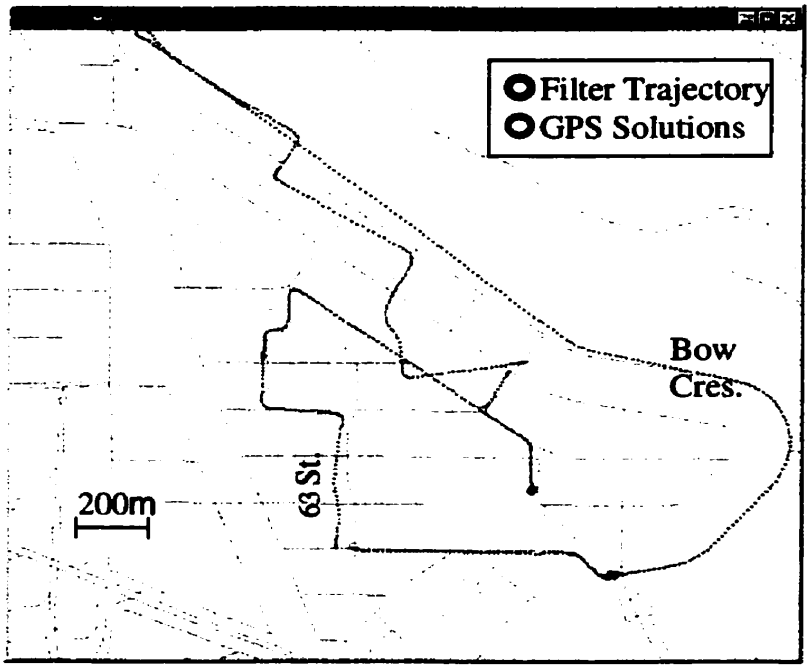

Figure 6.16 Bowness Test 1 Gyro No Augmentation Trajectory

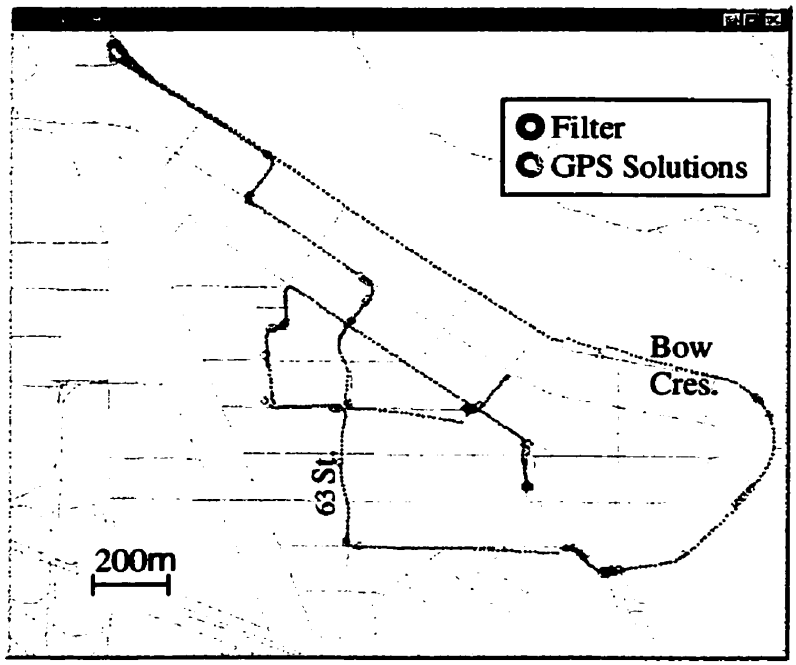

Figure 6.17 Bowness Test 1 ABS Height Augmentation Trajectory

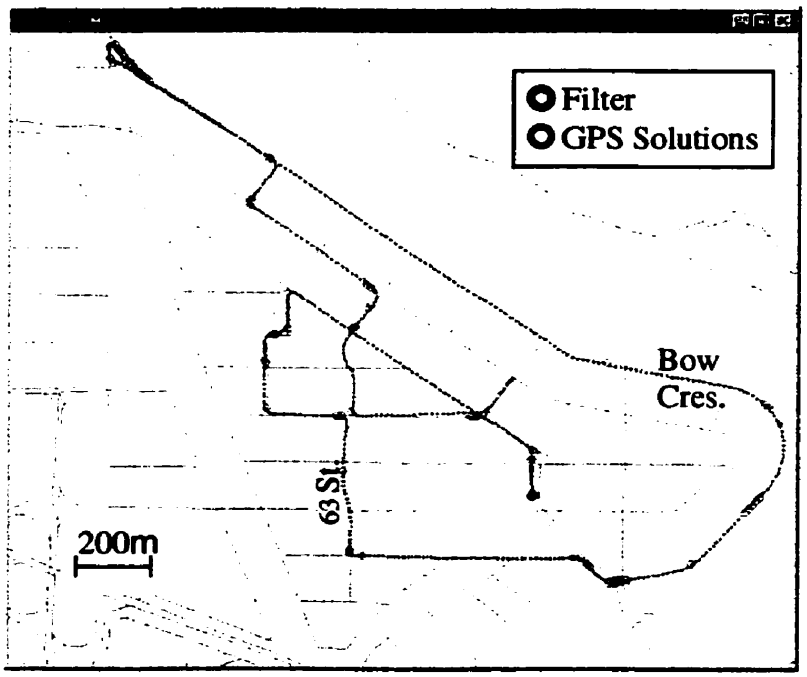

Figure 6.18 Bowness Test 1 Gyro Height Augmentation Trajectory

Cross-track errors were chosen to summarize the filter trajectory errors. Cross-track errors are computed by breaking the radial position error from the trajectory into a component in the direction of the heading (along-track) and perpendicular to this (across-track). The reference position was again a floating ambiguity carrier phase solution using FLYKIN. Table 6.4 contains the results for both Bowness tests. The case of no augmentation is not shown since there were many turns during the long GPS outage periods, leading to a confusion of cross-track and along-track errors.

Table 6.4 Bowness Test Cross-Track Filter Position Errors

C^3NAVG^2 GPS Augmentation	Dead Reckoning	Test 1		Test 2	
		Max (m)	RMS (m)	Max (m)	RMS (m)
Clock	Gyro and Odometry	49	7	54	8
	Differential ABS	65	9	32	7
Height	Gyro and Odometry	43	6	30	6
	Differential ABS	43	7	35	7
Clock and Height	Gyro and Odometry	34	5	44	8
	Differential ABS	33	5	35	6

It was noted that the addition of augmentations does not always improve performance. Individual examples of this are visible in the trajectory plots, and can be seen in the increased maximum errors in the statistics using both augmentations. The increase in error is due mainly to the inclusion of additional poor GPS solutions. Many of the GPS solutions that are added using augmentations just barely pass the maximum GDOP threshold, and as such are subject to larger errors. High GDOP solutions magnify the measurement and constraint errors into the position and velocity solutions. The GPS derived headings are notably affected. To confirm that GPS headings using both augmentations were showing increased errors, integrated heading observations were created as previously described. Much larger spikes are found in the integrated Gyro and ABS heading biases (Figure 6.19 and Figure 6.20) using augmentations than those from the Strathcona tests which did not use augmentations. The spikes are also clearly correlated between both the ABS and gyro plots, confirming that they are the result of GPS errors and not dead reckoning.

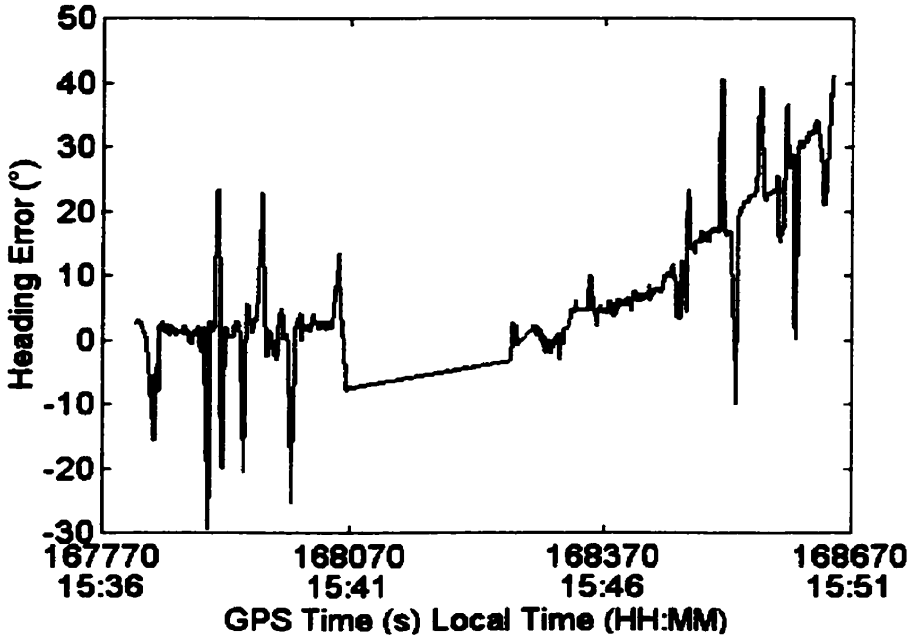

Figure 6.19 Bowness Test 1 Integrated Gyro Heading Bias

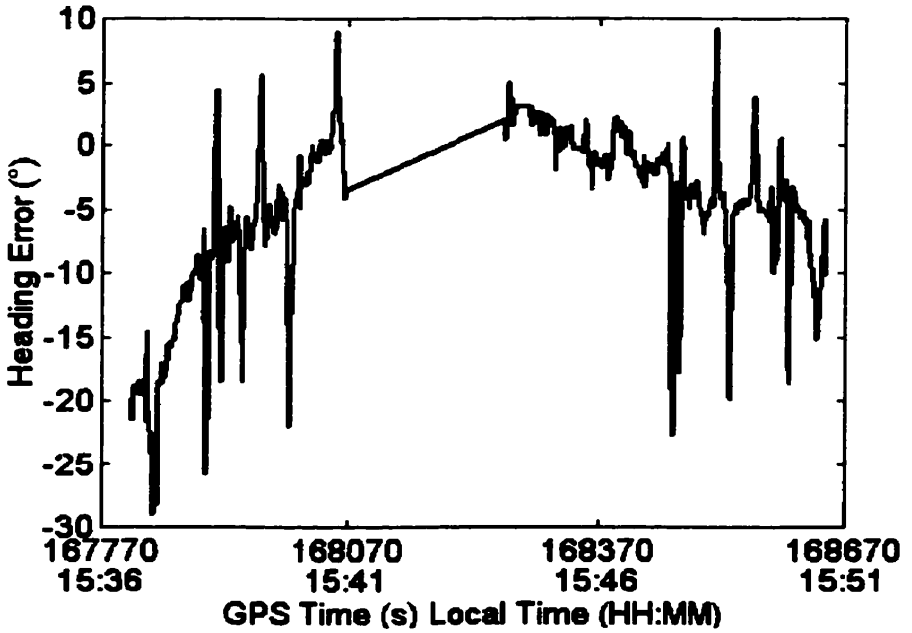

Figure 6.20 Bowness Test 1 Integrated Gyro Heading Bias

The position degradation of some of the augmented GPS results shows that the quality of GPS information is often more important than the quantity. Further testing will concentrate on height aiding and no aiding of the GPS solutions only. Height augmentation is more cost effective than clock coasting, requiring no additional sensors if the ground can be assumed flat in a local area (i.e. the last known height can be used) or a DTM is available. Clock coasting and height aiding generally perform almost equally since they are so highly correlated, but a precise oscillator adds substantial cost to the navigation system. The combination of the two augmentations does not provide much more benefit than either individually, because of the high correlation between them.

Error and estimated filter standard deviation plots were therefore generated for the cases of height augmentation and no augmentation of GPS, which will be concentrated on in further testing, and are found in the Appendix. The 2σ bounds again contain most of the errors, but there are extreme departures with no augmentation due to uncontrolled heading errors.

6.3 DOWNTOWN CORE TESTING (APRIL 4, 2000)

Testing was done in Calgary's downtown core to test the system under high multipath and masking conditions. On the test carried out on April 4, 2000, the connection to one of the wheel sensors was loose, so only the left wheel speed was reliably measured. Speed was therefore generated from observations from a single wheel, and only the gyro could be used for generating heading information. During this test, however, both narrow and wide correlator spaced GPS receivers were operated from a split antenna cable in the test vehicle. This would allow the multipath and tracking capabilities of the two receivers to be compared. To test the performance of ABS differential odometry, a second downtown test had to be performed on April 21, 2000.

Calgary's downtown core is located on a floodplain of the Bow River, so it is quite flat. There are several metres of slope over the test area, and a few overpasses along the test route that add roughly two more metres of vertical variation, so a fixed height

augmentation could reasonably be used. The temperature was about 10° C for both tests, and the roads were dry but patches of gravel remained from previous application. Three laps were done of a 6.4 km closed route, each taking about 20 minutes to complete. The test began in an area with open sky to allow for some calibration on each test run. Very little unobstructed sky was available during the test since numerous buildings have heights of 20 stories or more. Calgary's downtown city blocks are generally found to be 175 m in Easting and 100 m in Northing. Only a 5° isotropic mask angle was used in processing the GPS data.

The same three GPS augmentation schemes, in addition to unaugmented GPS, were used as in the previous testing. These solutions were combined with differential ABS odometry or gyro with ABS odometry dead reckoning in the navigation filter. The filtered trajectories from the first run, with height no GPS augmentation are presented in Figure 6.21 and Figure 6.23 for wide correlator spacing, and in Figure 6.22 and Figure 6.24 for narrow. A green line was added to mark the truth trajectory that was digitized on the map. The second and third test run trajectories can be found in the Appendix.

Since there was not enough GPS data to create an accurate carrier phase reference trajectory, errors had to be computed as the distance to the nearest line segment of the digital map. This method is less desirable than using a reference trajectory for several reasons. Only road centrelines are indicated on the map, but the vehicle does not follow the road centrelines closely and errors can be seen on the rounding of the sharp corners. In addition, the shortest computed distance is not always the distance to the road on which the vehicle was driving, i.e. a biased trajectory will show zero error every time it crosses another road on the map. Along-track errors are also not recognized. The large multipath error in the first run height aided results (bottom centre of Figure 6.23 and Figure 6.24) causes the position to be biased along the direction of the road. The bias is not apparent in the distance errors until the vehicle makes a right hand turn about three blocks later.

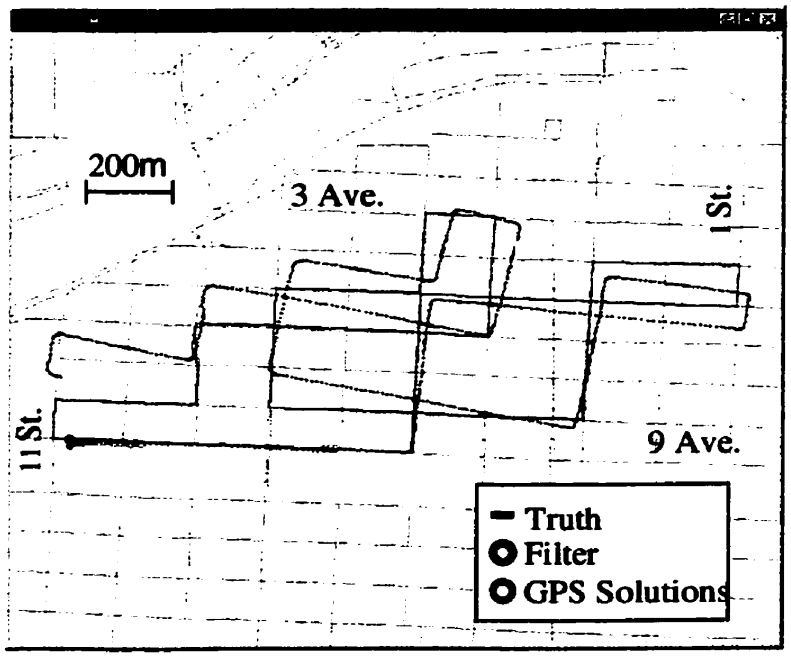

Figure 6.21 April 4 Downtown Test Run 1 Trajectory – Gyro & Odometry and Wide Correlator GPS with No Augmentation

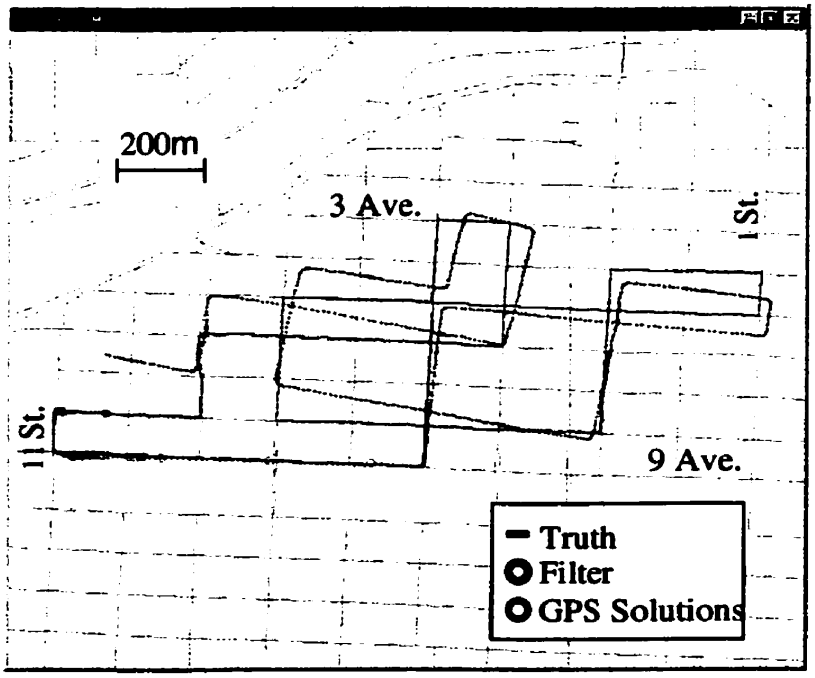

Figure 6.22 April 4 Downtown Test Run 1 Trajectory – Gyro & Odometry and Narrow Correlator GPS with No Augmentation

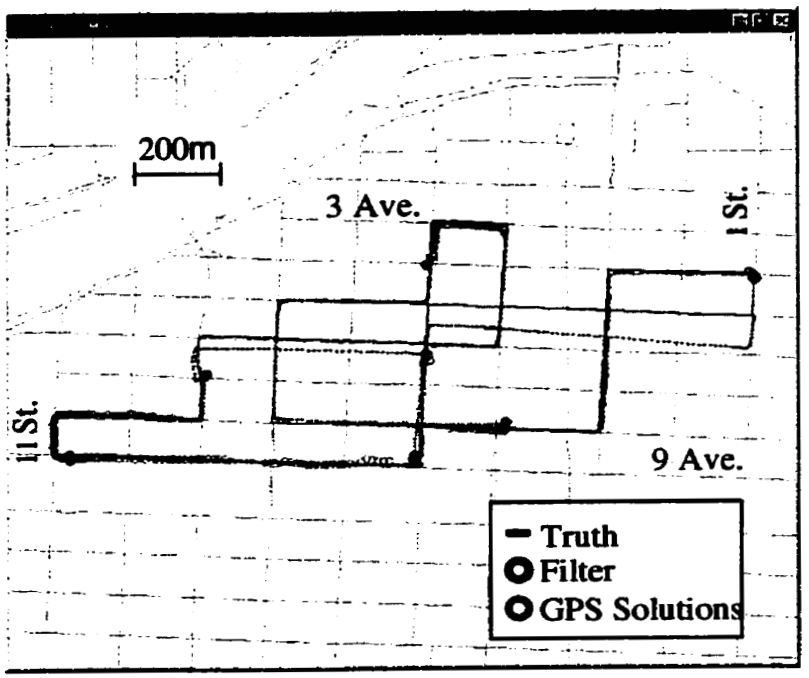

Figure 6.23 April 4 Downtown Test Run 1 Trajectory – Gyro and Odometry and Wide Correlator GPS with Height Augmentation

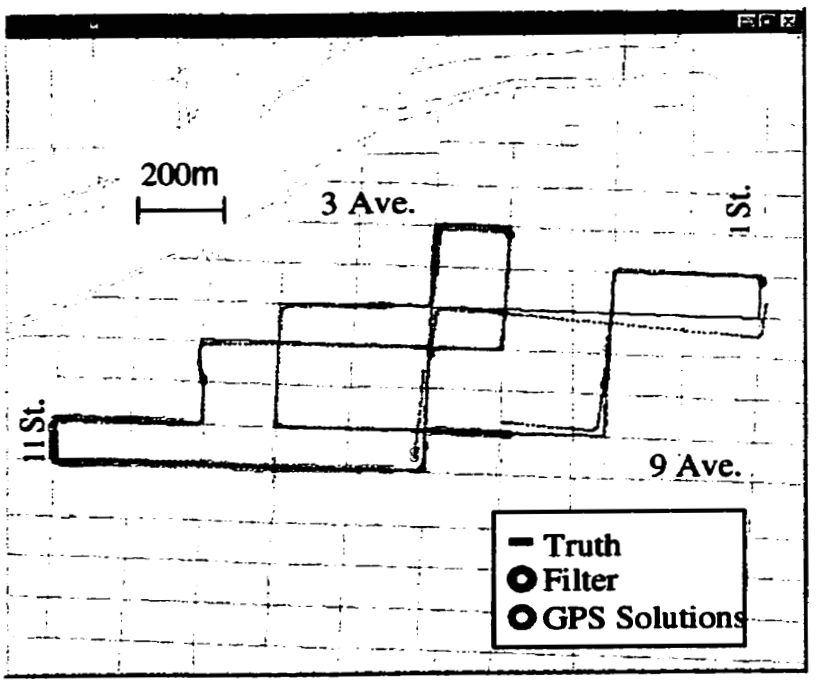

Figure 6.24 April 4 Downtown Test Run 1 Trajectory – Gyro & Odometry and Narrow Correlator GPS with Height Augmentation

Table 6.5 shows the number and quality of the C^3NAVG^2 GPS solutions for each of the four augmentation schemes for the three test runs combined. The narrow correlator clearly reduces the multipath error, but significant multipath error was experienced by even the narrow correlator at a few locations. It must be noted that the wide correlator receiver logging software stopped recording data half way into the third run, so the number of computed positions was affected. Table 6.6 summarizes the maximum and RMS errors for the filtered data, i.e. C3NAVG2 GPS positions with gyro and odometry dead reckoning.

Table 6.5 April 4 Combined Downtown Tests - Augmented C^3NAVG^2 GPS Statistics

C^3NAVG^2 GPS Augmentation	Correlator Spacing	# of Solutions	Max Error (m)	RMS Error (m)
None	Wide	804	18	5
	Narrow	996	14	3
Clock	Wide	1348	46	10
	Narrow	1518	34	6
Height	Wide	1332	24	7
	Narrow	1507	24	5
Both	Wide	1584	47	9
	Narrow	1761	37	6

Table 6.6 April 4 Combined Downtown Tests – Integrated Filter Position Errors (GPS with Gyro and Odometry Dead Reckoning)

C^3NAVG^2 GPS Augmentation	Correlator Spacing	Max (m)	RMS (m)
None	Wide	153	23
	Narrow	145	19
Clock	Wide	68	20
	Narrow	44	9
Height	Wide	81	13
	Narrow	44	9
Both	Wide	72	15
	Narrow	55	7

Figure 6.25 shows the trajectory derived using clock and height augmentation of the wide correlator spacing GPS receiver data. It was chosen to illustrate the problems associated with multipath that frequently corrupt filter positions and headings. Multipath causes

position errors of up to 80 metres, typically at intersections where visibility is temporarily increased. Individual range errors are very difficult to identify, since there is usually little or no redundancy. Consecutive epochs are also plagued by the slowly changing multipath geometry when stationary, giving tight spacing (precision) to consecutive solutions. Although multipath corrupted GPS position residuals would likely be flagged as outliers, a reasonable assumption might be made by the filtering algorithm that the trajectory was in error due to prior multipath or dead reckoning sensors drifts. Using the map, the distinction can usually be made visually as to whether the trajectory, GPS positions, or both are in error, but it is still difficult to impart these decision making criteria to a map aiding algorithm.

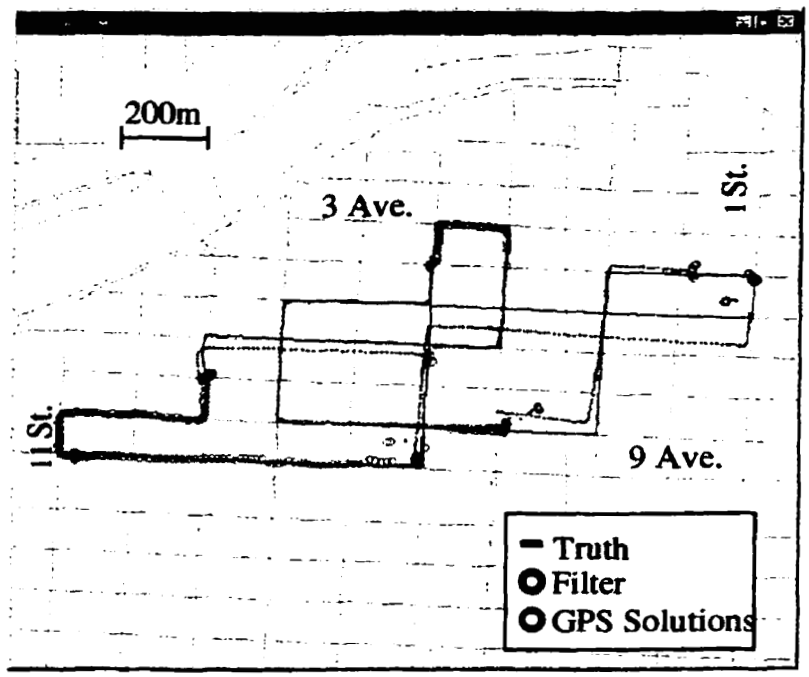

Figure 6.25 April 4 Downtown Test Run 1 – Gyro & Odometry and Wide Correlator GPS with Clock & Height Augmentation

6.4 DOWNTOWN CORE TESTING (APRIL 21, 2000)

The test runs on April 21 were used to compare the differential ABS and gyro with odometry filter performance under urban canyon conditions. A second Murata gyro unit

was also tested to confirm the performance of the primary tested unit. The loop used for the previous downtown testing was again run three times, and the data was processed using the same augmentations previously described. Only a wide spacing correlator was used for this testing. Trajectory results are presented in Figure 6.26 to Figure 6.29.

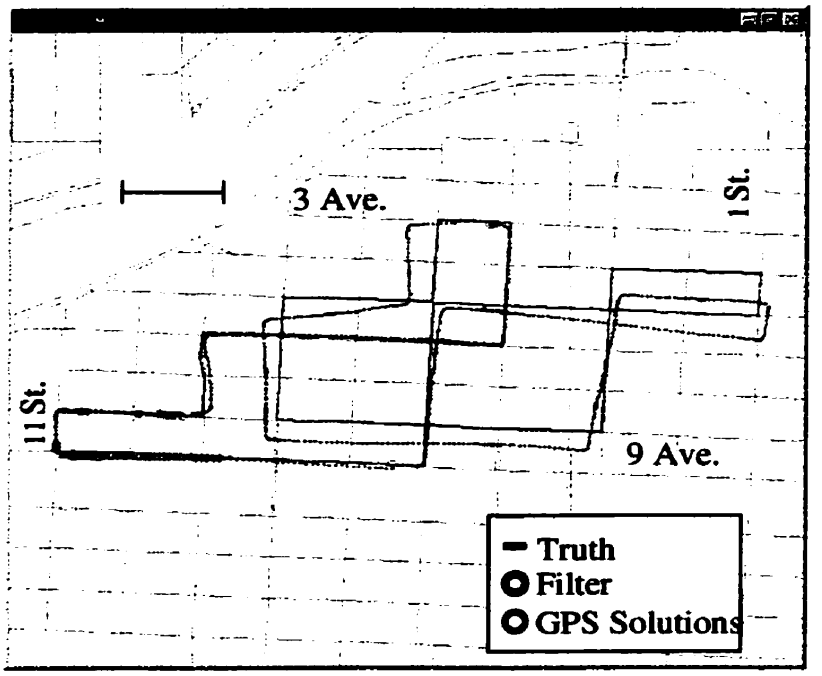

Figure 6.26 April 21 Downtown Test Run 1 – Differential ABS with No GPS Augmentation

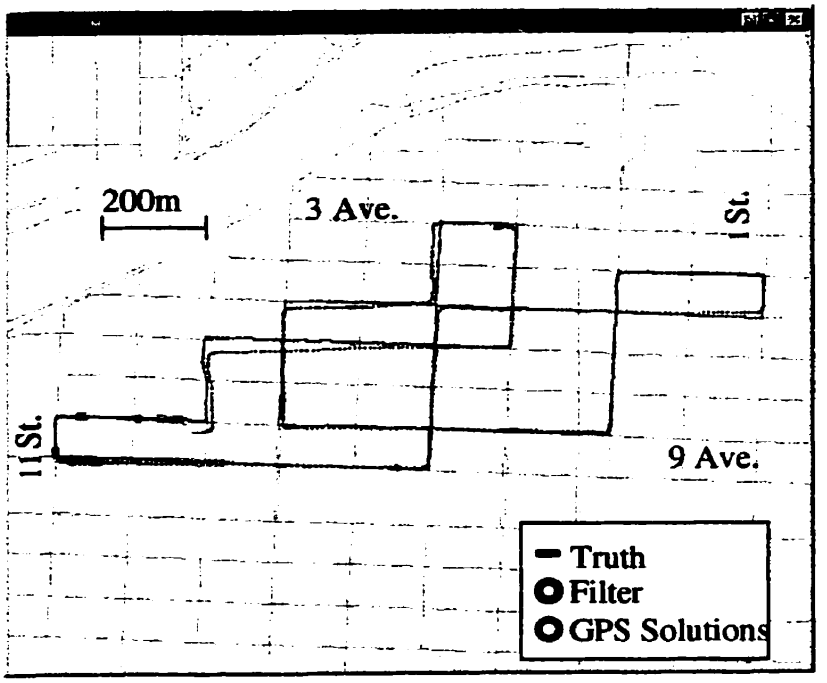

Figure 6.27 April 21 Downtown Test Run 1- Gyro 1 & Odometry with No GPS Augmentation

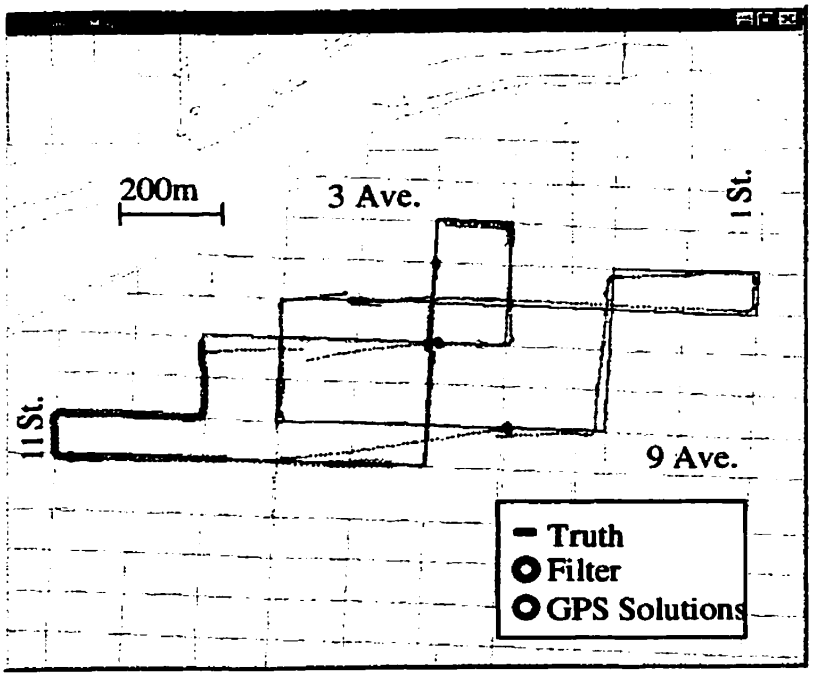

Figure 6.28 April 21 Downtown Test Run 1 – Differential ABS with Height GPS Augmentation

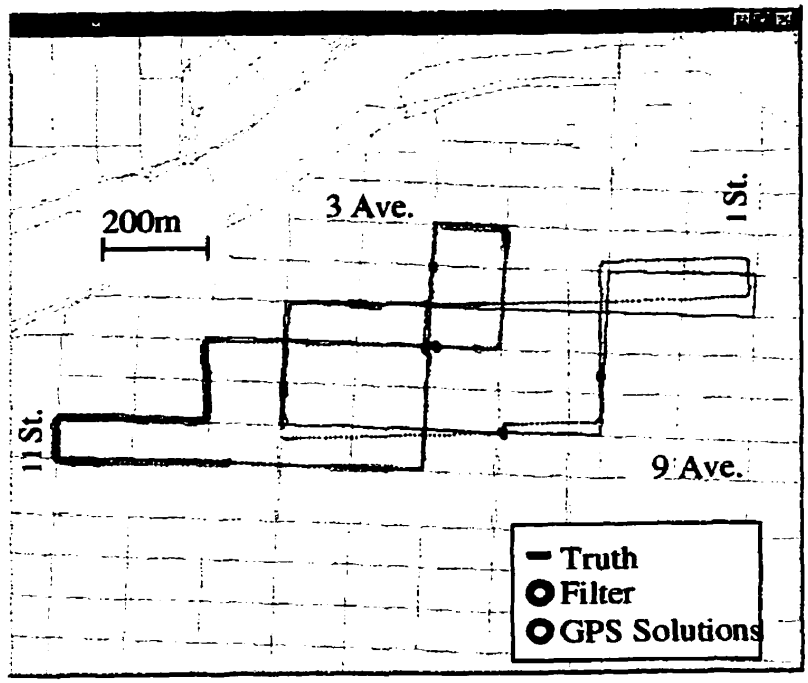

Figure 6.29 April 21 Downtown Test Run 1- Gyro 1 & Odometry with Height GPS Augmentation

Again, the gyro outperforms ABS since the zero-rotation bias can be calibrated every time the vehicle is forced to stop, while drifts due to wheel scale factors are apparent in the ABS results. The errors were computed, as in the previous downtown testing, as the distance to the nearest road segment, and are summarized in Table 6.7 for the combination of all three test runs. It is noted that the best statistics were achieved using the second gyro and no GPS augmentation, confirming that the quality of GPS information is generally more important than the quantity. Table 6.8 shows that the maximum GPS position error using no augmentation is less than half of those with augmentations, likely accounting for a great deal of the filter position errors.

Table 6.7 April 21 Combined Downtown Tests – Integrated Filter Position Errors

C^3NAVG^2 GPS Augmentation	Dead Reckoning	Max (m)	RMS (m)
None	Differential ABS	191	34
	Gyro 1 & Odometry	68	14
	Gyro 2 & Odometry	55	9
Clock	Differential ABS	114	19
	Gyro 1 & Odometry	76	15
	Gyro 2 & Odometry	79	15
Height	Differential ABS	115	18
	Gyro 1 & Odometry	67	14
	Gyro 2 & Odometry	69	14
Both	Differential ABS	115	17
	Gyro 1 & Odometry	67	13
	Gyro 2 & Odometry	69	12

Table 6.8 April 21 Combined Downtown Tests - Augmented C^3NAVG^2 GPS Statistics

C^3NAVG^2 GPS Augmentation	# of Solutions	Max Error (m)	RMS Error (m)
None	1018	21	6
Clock	1672	56	8
Height	1689	47	7
Both	1730	47	7

7 CONCLUSIONS AND RECOMMENDATIONS

7.1 CONCLUSIONS

Dead reckoning can provide continuous navigation coverage with little added cost to the production of a vehicle. The ABS systems which are now standard on many vehicles can provide this information solely, or odometer information can be used with a heading sensor such as a gyro to provide the distance and heading information needed. Test results show that position errors can be kept at the 20 m level most of the time, but deviations of 100 m or more are possible in urban canyons where GPS information is not available or not reliable. Map matching information could likely control this error to meet a 20 m requirement.

Distance or speed measurements can be obtained easily to a level that meets the needs typical of vehicle navigation, but calibration of the heading sensors using GPS is the primary challenge. This is especially true since low cost sensors normally have more instability due to environmental factors such as temperature and vibration. GPS availability is poor in urban canyons and under dense foliage, so GPS data receives a large weighting in the filter when it becomes available. Unfortunately, GPS positions are often corrupted by multipath or constraint errors that are amplified by poor geometry. A GDOP threshold was used for this research, but it was not sufficient to keep GPS blunders out of the navigation solution. Since the quality of GPS information was found to be generally more important than the quantity or frequency, a threshold on a reliability parameter such as the MHE would likely lead to improved results.

A low cost gyro outperforms differential odometry for several reasons. Firstly, the drift can be calibrated very quickly if the vehicle makes a stop. Wheel slippage plagues wheel rotation measurements, and in many environments there are factors such as gravel and ice on roads that make slippage inevitable. Wheel scale factors can be better calibrated at high speeds, but slippage and radius of the tires are variable with speed, temperature, and

other factors, making calibration less effective. The gyro does suffer from errors associated with pitch and roll on turns, but vehicle tilt is generally small when driving at normal speeds in urban centres.

Augmentations such as height information or the use of a precise oscillator can greatly improve the availability of GPS solutions. Geostationary satellite augmentation is of little help in urban canyons, due to signal masking. The GLONASS constellation is presently quite small, and adds little advantage for the increase in cost of the receiver. Other proposed systems, such as GALILEO, may provide a substantial improvement to GNSS availability and reliability, but quantifying this improvement is still speculation. Height and clock coasting augmentations perform almost identically, due to the large correlation between the two. The combination of the two provides little advantage over either individually, for the same reason. Height observations can be acquired from the last computed height, a lookup table or DTM, or barometric height filtering. Clock augmentation with a precise oscillator is a more expensive option due to the cost of the oscillator, but can add fast reacquisition of signals as a benefit.

7.2 RECOMMENDATIONS FOR FURTHER RESEARCH

The most substantial improvement to the results could likely be achieved using map information as a navigation aid. Headings could be constrained by road direction, and biases could be removed by shifting the position to the appropriate road. More intelligent algorithms to detect and discard multipath corrupted GPS positions could also be designed. Further analysis of the heading errors from poor GPS solutions would be beneficial, since GPS headings appear to be largest source of error in many of the tests. This would involve using a digital map, or preferably INS heading, to quantify the nature of GPS heading errors and improve adaptive features and blunder detection algorithms to minimize their effect on filter performance.

The alternative of fully integrating the GPS positioning module with dead reckoning in a single filter could be tested, since it might be possible to detect individual multipath

biased ranges in a centralized filter. Temperature repeatability of the gyro voltage offset and scale factors could be tested. If temperature stability could be shown, a calibration function or table could be developed over time that would minimize the need for constant calibration found in previous testing. A study of whether or not an accurate calibration table could be developed for a compass could also be investigated. Drifting of the relative heading instruments could be controlled, with clear advantages in places without digital map coverage.

REFERENCES

Analog Devices (2000). Accelerometer Design and Applications. http://www.analog.com/industry/iMEMS/library/Sensor971.pdf.

Analog Devices (2000). ADXL210/202 Specifications. http://www.analog.com/pdf/ADXL202_10_b.pdf.

Bennett, S., R. Dyott, D. Allen, J. Brunner, R. Kidwell, R. Miller (1998). Fiber Optic Gyros as Replacements for Mechanical Gyros. American Institute of Aeronautics and Astronautics.

Braasch, M. (1996). Multipath Effects. Global Positioning System Theory and Applications, Progress in Astronautics and Aeronautics, Vol. 163. American Institute for Astronautics and Aeronautics, Washington, D.C.

Braasch, M. (1998). GPS Multipath Characterization and Mitigation. Guest Lecture, July 8, 1998, Department of Geomatics Engineering, University of Calgary.

Brown, R., P. Hwang (1992). Introduction to Random Signals and Applied Kalman Filtering. John Wiley and Sons Inc., Toronto.

Bullock, J.B. (1995). A Prototype Portable Navigation System Utilizing Map Aided GPS. M.Sc. Thesis, UCGE Report 20081, Department of Geomatics Engineering, University of Calgary.

Cannon, M.E. (1997). ENGO 561 Lecture Notes. Department of Geomatics Engineering, University of Calgary.

Carlson, N. (1993). Federated Kalman Filter Simulation Results. Proceedings of the 49th Annual Meeting of the Institute of Navigation. Pp. 421-435.

Etak (2000). Etak Real-Time Traffic Information Delivery System. http://www.etak.com/traffic/present1.html.

G&S Engineering (1998). Magnetic Sensors. http://www.bioweb.net/mfg/gsengineering/hall.shtml.

Ganssle, J. (1989). Anatomy of a Fluxgate. Ocean Navigator. No. 28. Pp 75-79.

Gao, Y., J. Krakiwsky, M. Abousalem, J. McLellan (1993). Comparison and Analysis of Centralized, Decentralized, and Federated Filters. Navigation: Journal of the Institute of Navigation, Vol. 40, No. 1. Pp. 69-86.

Geier, G.J. (1998). GPS Integration with Low Cost Sensor Technology for Automotive Applications. Lecture Notes, Department of Geomatics Engineering, University of Calgary.

Gelb, A. (1974). Applied Optimal Estimation. The M.I.T. Press, Cambridge.

GMH Engineering (2000). Application Note 1001: Using Non-Contact Speed Sensing to Measure Vehicle Ground Speed. http://www.gmheng.com/pdf/an1001.pdf.

GMH Engineering (2000). Delta Speed Sensor. http://www.gmheng.com/pdf/delta.pdf.

Hall, T., B. Burke, M. Pratt, P. Misra (1997). Comparison of GPS and GPS+GLONASS Positioning Performance. ION GPS-97 Proceedings. Pp. 1543-1550.

Harris, C. (1989). Prototype for a Land Based Automatic Vehicle Location and Navigation system. M.Sc Thesis, UCGE Report 20033, Department of Geomatics Engineering, University of Calgary.

Harvey, R. (1998). Development of a Precision Pointing System Using an Integrated Multi-Sensor Approach. M.Sc. Thesis, UCGE Report 20117, Department of Geomatics Engineering, University of Calgary.

Hatch, R. (1982). The Synergism of GPS Code and Carrier Measurements. Proceedings of the 3rd International Geodetic Symposium on Satellite Doppler Positioning. Pp. 1213-1231.

Hayashi, N. (1996). Augmentation of GPS with a Barometer and a Heading Rate Gyro for Vehicular Navigation. M.Sc. Thesis, UCGE Report 20098, Department of Geomatics Engineering, University of Calgary.

Heidemann, R., A. Jeje, F. Mohtadi (1987). An Introduction to the Properties of Fluids and Solids. University of Calgary Press.

Infineon (2000), Microwave Motion Sensors.
http://www.infineon.com/products/sensors/pdf/kmy10.pdf.

Infineon (2000). SFH9201 Product Specifications.
http://www.infineon.com/products/opto/pdf/sfh9201.pdf.

Janicke, J.M. (1998). The Magnetic Measurements Handbook. Magnetic research, Inc. Butler, NJ.

Kaplan, E.D. (1996). Understanding GPS, Artech House, Norwood, MA.

Kealy, A., M. Tsakiri, M. Stewart (1999). Land Navigation in the Urban Canyon – A Kalman Filter Solution using Integrated GPS, GLONASS and Dead Reckoning. ION GPS-99 Proceedings.

Konno, M., S. Sugawara (1987). Equivalent Circuits of a Piezoelectric Vibratory Gyro. Ultrasonics Symposium Proceedings, IEEE Ultrasonics Ferroelectric & Frequency Control Society. New York, NY. Pp 539-542.

Krakiwsky, E.J. (1987). Papers for the CISM Adjustment and Analysis Seminars. Canadian Institute of Geomatics, Ottawa.

Kyle, T. (1991). Atmospheric Transmission, Emission and Scattering. Pergamon Press.

Lachapelle, G., S. Ryan, M. Petovello, J. Stephen (1997). Augmentation of GPS for Vehicle Navigation Under Signal Masking. ION GPS-97 Proceedings. Pp. 1511-1519.

Lan, H. (1996). Development of a Real-Time Kinematic GPS System. M.Sc. Thesis, UCGE Report 20107, Department of Geomatics Engineering, University of Calgary.

Lancaster, D. (1996). Tech Musings. http://www.tinaja.com/glib/muse103.pdf. Vol 103.

Lappin, Todd (1999). The New Road Rage. Wired, July 1999. Condé Nast Publications, Inc., San Francisco.

Lee, Y. (1993). RAIM Availability for GPS Augmented with Barometric Altimeter and Clock Coasting. Navigation: Journal of the Institute of Navigation, Vol. 40, No. 2. Pp. 179-198.

Leick, A. (1995). GPS Satellite Surveying. John Wiley and Sons, Inc., Toronto.

Lutgens, F.K., E.J. Tarbuck (1992). The Atmosphere - An Introduction to Meteorology. Prentice-Hall, Inc., Toronto.

Marchent, B.G, M. Foster (1985). Compass. U.S. Patent 4,539,760. Sept. 10, 1985.

Mark, J., D. Tazartes, B. Fidric, A. Cordova (1991). A Rate Integrating Fiber Optic Gyro. Navigation: Journal of the Institute of Navigation, Vol. 38, No. 4. Pp. 341-352.

Matthews, A. (1990). Utilization of Fiber Optic Gyros in Inertial Measurement Units. Navigation: Journal of the Institute of Navigation. Vol. 37, No. 1. Pp. 17-38.

Maybeck, P. (1994). Stochastic Models, Estimation, and Control, Volume I and II. Navtech Book and Software Store.

McLellan, J. (1992). Design and Analysis of A Low cost GPS Aided Navigation System. M.Eng. Thesis, UCGE Report 20097, Department of Geomatics Engineering, University of Calgary.

Misra, P., M. Pratt, B. Burke, R. Ferranti (1995). Adaptive Modeling of Receiver Clock for Meter-Level DGPS Vertical Positioning. ION GPS-95 Proceedings. Pp 1127-1135.

Motta, L. (2000). Coriolis Force. www.treasure-troves.com/physics/CoriolisForce.html.

Murata (1998). Murata's Piezoelectric Products. Catalog P19E-6.

Murata (1999). Piezoelectric Vibrating Gyroscope ENC Series. Catalog S42E-2.

Nakamura, T. (1990). Vibration Gyroscope Employs Piezoelectric Vibrator. JEE, September, 1990.

National Instruments (1995). DAQPad-MIO-16XE-50 User Manual. PN 320935A-01.

National Instruments (1992). Field Wiring and Noise Considerations for Analog Signals. Application Note 025.

NOAA NGDC (National Geophysical Data Centre) (2000). GeoMag Calculator. http://www.ngdc.noaa.gov/seg/potfld/java/GeoMag.shtml.

Petovello, M., G. Lachapelle (2000). Estimation of Clock Stability Using GPS. GPS Solutions, Vol. 3 No 4., John Wiley and Sons, Toronto.

Phillips A. (1993). Low Cost dead Reckoning Sensors. ION NTM-93 Proceedings. Pp. 145-149.

Phillips, A., T. O'Brien, C. Taylor (1989). Two-Axis Angular Rate Gyroscope. U.S. Patent 4,811,613. Mar. 14, 1989.

Precision Navigation Inc. (1997). TCM2 Electronic Compass Module User's Manual.

Press, W.H., S.A. Teukolsky, W.T. Vetterling, B.P. Flannery (1992). Numerical Recipes in C. 2nd Ed., Cambridge University Press, Cambridge.

Ripka, P. (1992). Review of Fluxgate Sensors. Sensors and Actuators, Vol. 33, No. 3. Pp. 129-141.

Ray, Jayanta K. (2000). Mitigation of GPS Code and Carrier Phase Multipath Effects using a Multi-Antenna System, Ph.D. Thesis, UCGE Report No. 20136. Department of Geomatics Engineering, University of Calgary.

RepairWare (1998). What happened to the air in my tires? http://www.ultramall.com/Autodig/autotipf/airprssr.html.

Ribbens, W. (1994). Understanding Automotive Electronics. SAMS Publishing.

Rüeger, J. (1982). Inertial Sensors Part I: Gyroscopes. Technical report, UCGE Report 30002, Department of Geomatics Engineering, University of Calgary.

Russian Space Command (1995). Global Orbiting Navigation Satellite System Interface Control Document.

Ryan, S., J. Stephen, G. Lachapelle (1999). Testing and Analysis of Reliability Measures for GNSS Receivers in a Marine Environment. Canadian Aeronautics and Space Journal, Vol. 45, No. 3. Pp. 270-280.

Savage, P. (1978). Strapdown Sensors. NATO AGAARD Lecture Series No. 95.

Schleppe, J., J. McLellan, J. Battie, J. McCarron (1995). Twenty-Centimetre Kinematic DGPS Positioning with Digital Barometry. ION GPS-95 Proceedings. Pp. 1831-1840.

Shearwood, C. (1997). Micromachined Rotating Gyroscope. http://www.shef.ac.uk/uni/projects/mesu/projects/gyro.html.

Skaloud, J. (1995). Strapdown INS Orientation Accuracy with GPS Aiding. M.Sc. Thesis, UCGE Report 20079, Department of Geomatics Engineering, University of Calgary.

Spilker, J. Jr. (1996). GPS Signal Structure and Theoretical Performance. Global Positioning System Theory and Applications, Progress in Astronautics and Aeronautics, Vol. 163. American Institute for Astronautics and Aeronautics, Washington, D.C.

St. Lawrence, W. (1993). Understanding Tilt Vectors and Processing Tiltmeter Data. Applied Geomechanics, Santa Cruz.

Sushko, M. (1993). Dead Reckoning Navigation Assistance for GPS Based AVL Systems. ION NTM-93 Proceedings.

Systron Donner (1998). BEI GyrochipII Product Specifications. http:\\www.systron.com\gyroiids.pdf.

Thomas, Jr., J.B. (1991). Digital Phase-Lock Loop. U.S. Patent 5073907.

U.S. Digital (2000). E5M Rugged Incremental Optical Encoder Product Specifications.

U.S. JPO (1995). Global Positioning System Standard Positioning Service Signal Specification. Navtech Book and Software Store.

van Nee, R. (1992). GPS Multipath and Satellite Interference. Proceedings of the 48th Annual Meeting of the Institute of Navigation. Pp. 167-177.

Viatran (1997). Specification Sheet – Model 246/346 Pressure Transmitter.

Viatran (1998). Technical Note 2: Pressure Sensing Technologies.

Williams, C., C. Shearwood, P. Mellor, R. Yates (1997). Modeling and Testing of a Frictionless Levitated Micromotor. Sensors and Actuators, Vol. A61. Pp 469-473.

Young, H.D. (1992). University Physics. Addison-Wesley. Don Mills, ON.

Zavoli, W.B., K.A. Milnes, G.R. Peterson (1988). Method and Apparatus for Measuring Relative Heading Changes in a Vehicular Onboard Navigation System. U.S. Patent 4788645.

Zhang, Z. (1997). Impact of Rubidium Clock Aiding on GPS Augmented Vehicular Navigation. M.Sc. Thesis, UCGE Report 20112, Department of Geomatics Engineering, University of Calgary.

Zickel, R., N. Nehemia (1994). GPS Aided Dead Reckoning Navigation. ION NTM-94 Proceedings. Pp. 577-586.

APPENDIX

STRATHCONA TESTS (JANUARY 10, 2000)

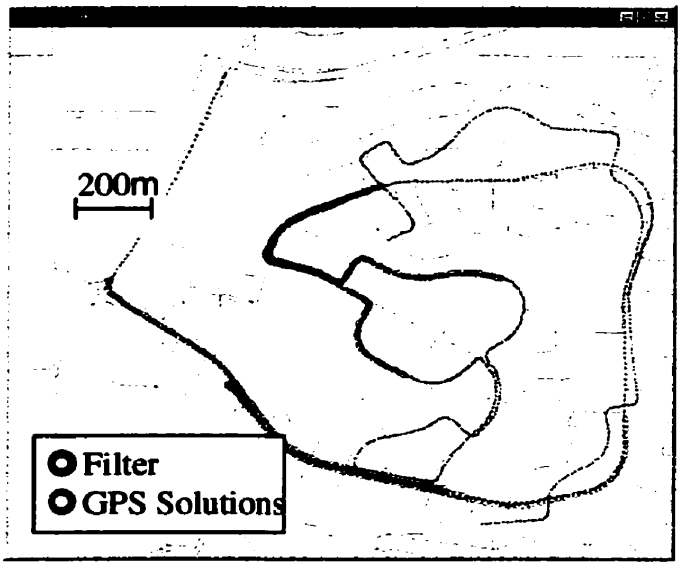

Figure A.1 Strathcona Test 1 Trajectory – Differential ABS with 200 s GPS Masking

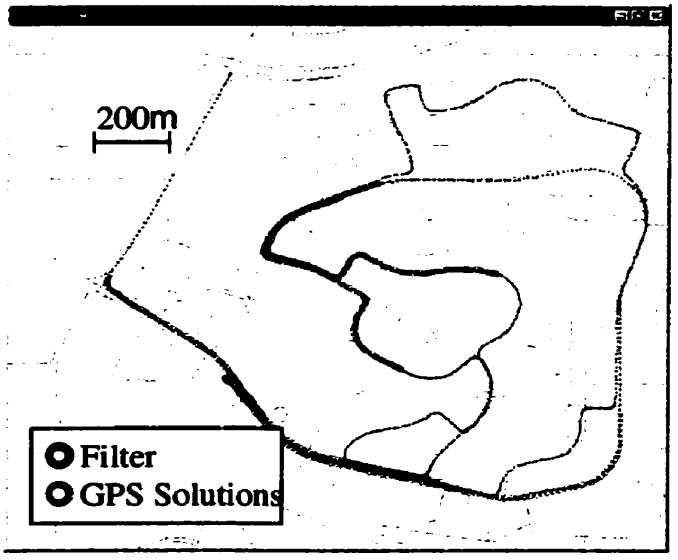

Figure A.2 Strathcona Test 1 Trajectory - Gyro & Odometry with 200 s GPS Masking

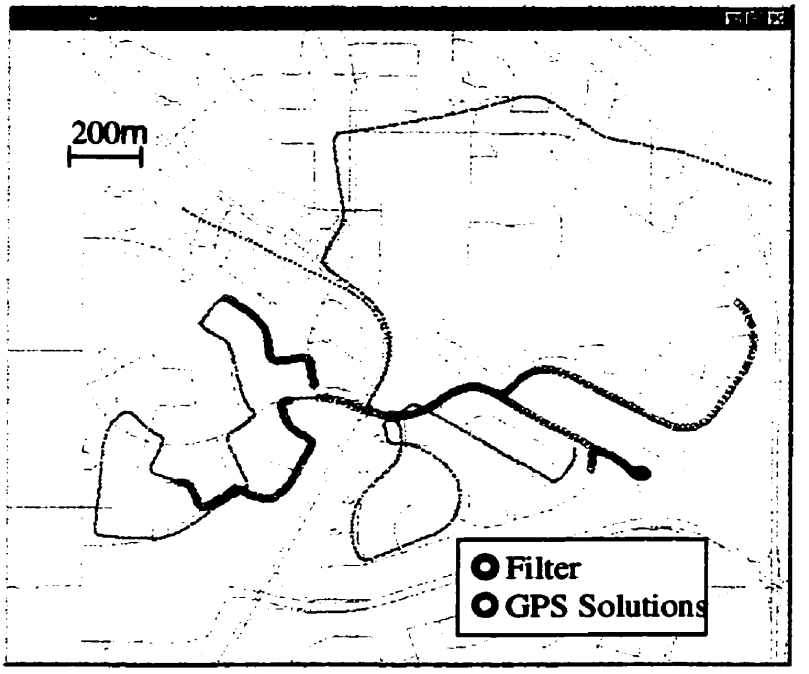

Figure A.3 Strathcona Test 2 Trajectory – Differential ABS with 200 s GPS Masking

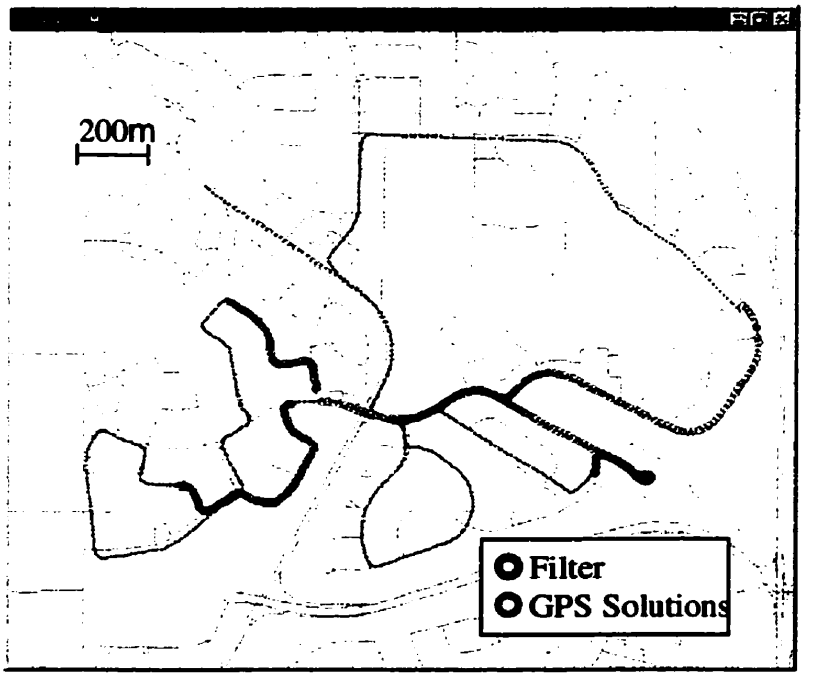

Figure A.4 Strathcona Test 2 Trajectory - Gyro & Odometry with 200 s GPS Masking

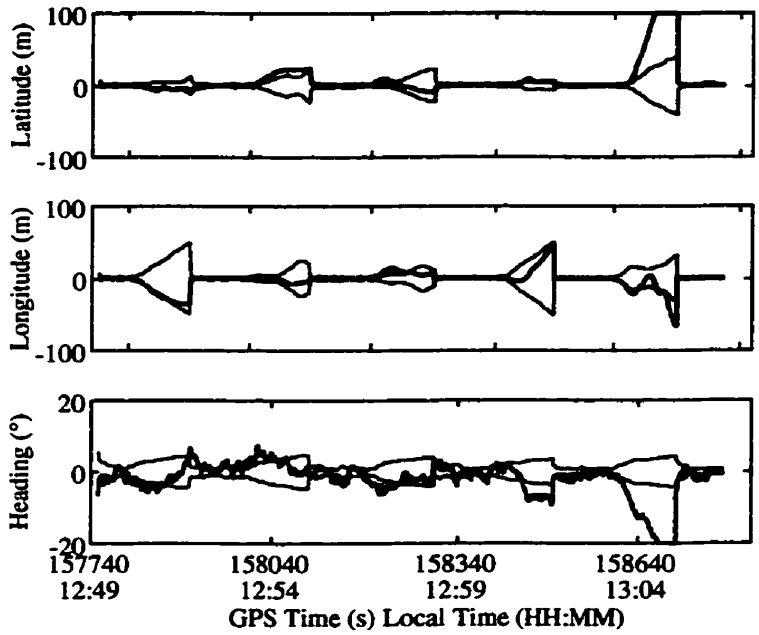

Figure A.5 Strathcona Test 2 Errors and Estimated 2σ Bounds – Differential ABS with 100 s GPS Masking

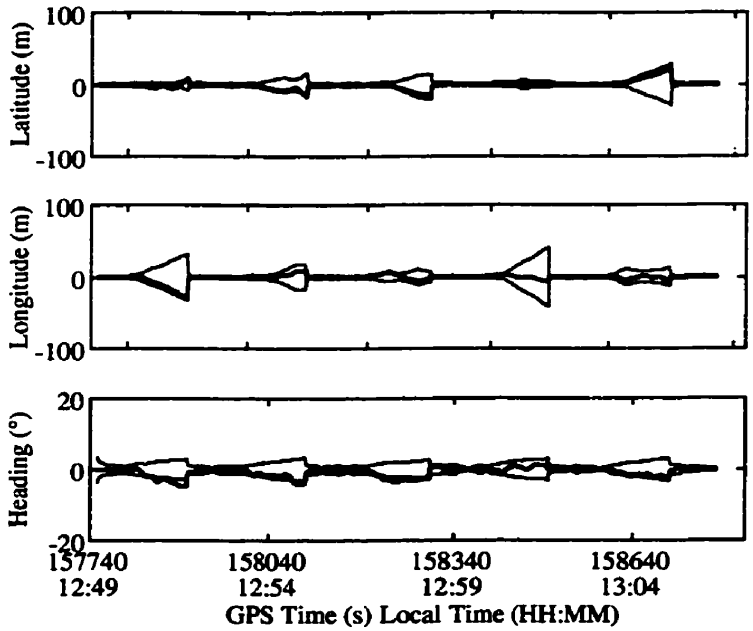

Figure A.6 Strathcona Test 2 Errors and Estimated 2σ Bounds - Gyro & Odometry with 100 s GPS Masking

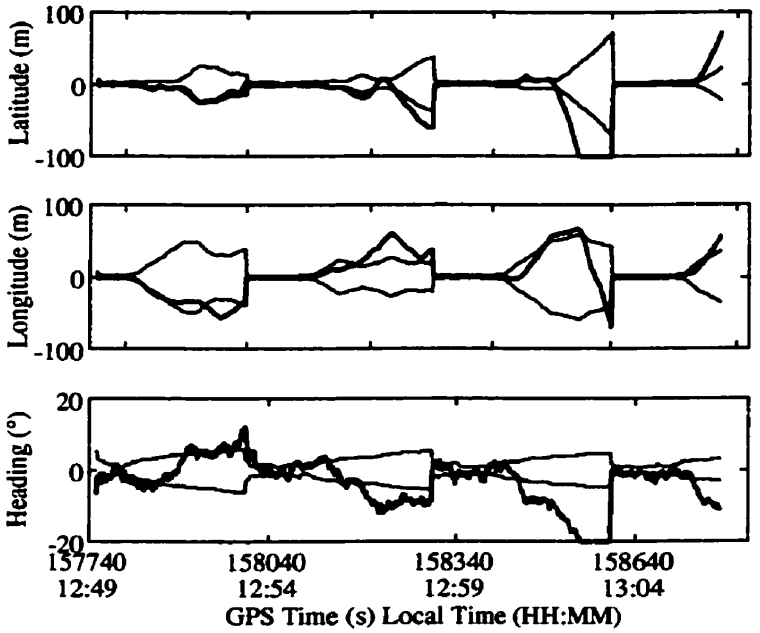

Figure A.7 Strathcona Test 2 Errors and 2σ Bounds – Differential ABS with 200 s GPS Masking

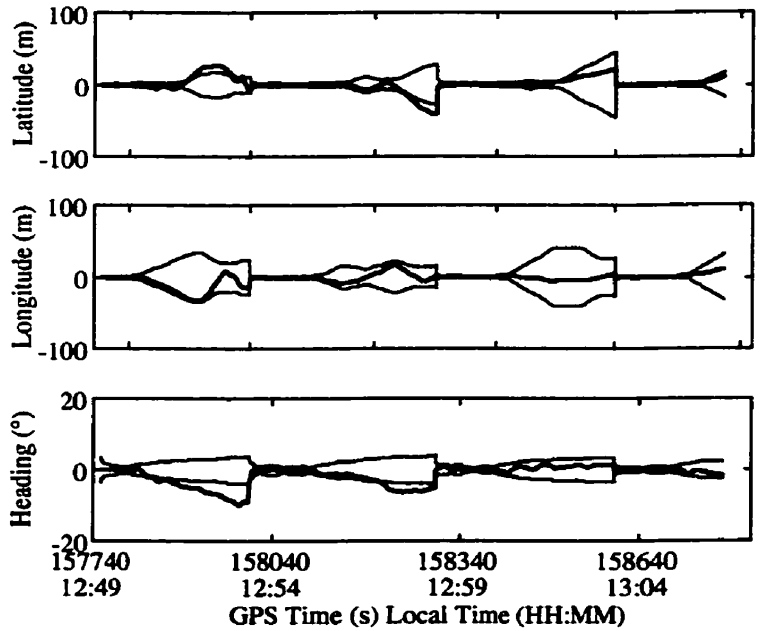

Figure A.8 Strathcona Test 2 Errors and 2σ Bounds - Gyro & Odometry with 200 s GPS Masking

BOWNESS TESTS (MARCH 13, 2000)

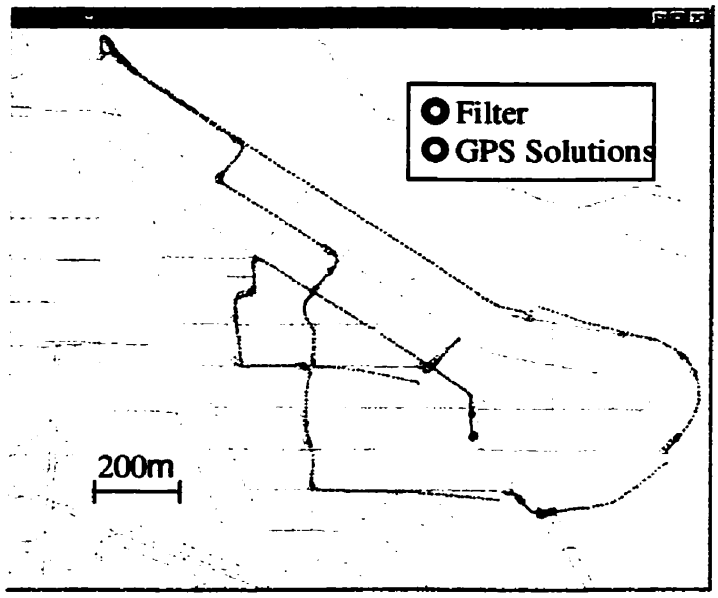

Figure A.9 Bowness Test 1 Trajectory – Differential ABS with Clock GPS Augmentation

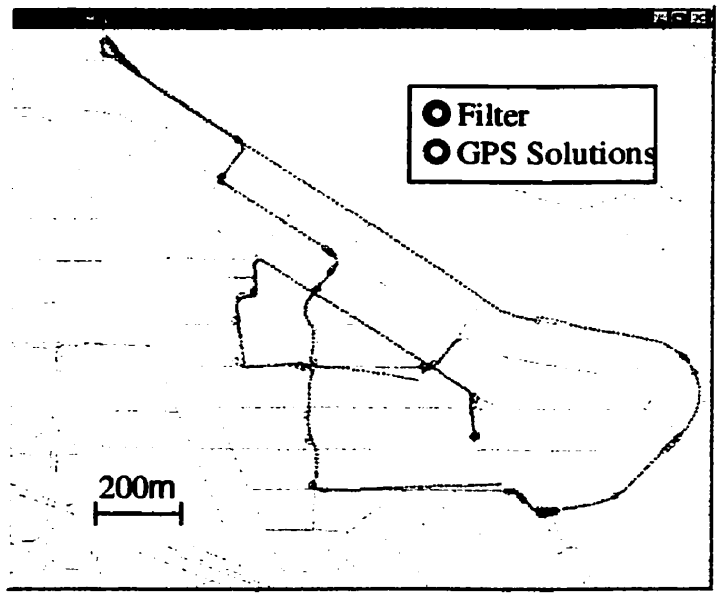

Figure A.10 Bowness Test 1 Trajectory – Differential ABS with Clock & Height GPS Augmentations

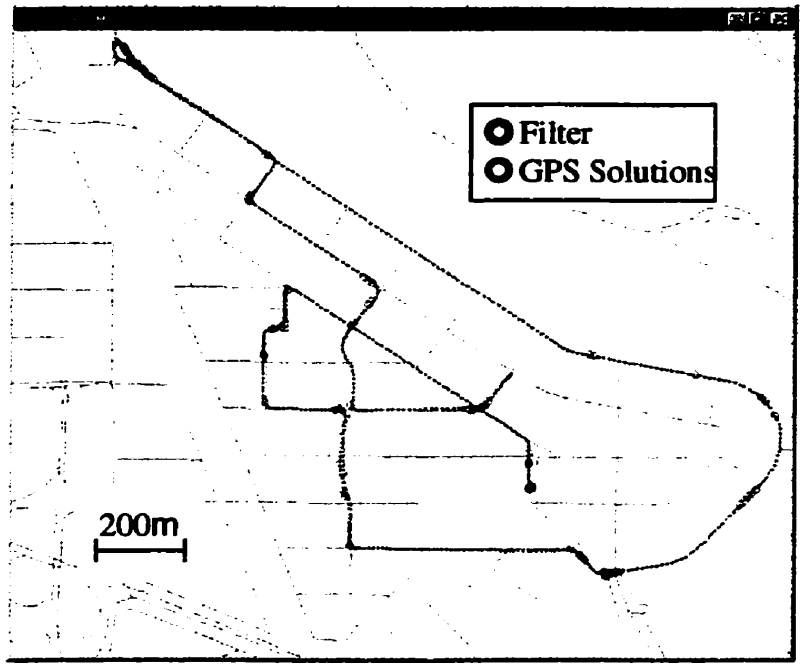

Figure A.11 Bowness Test 1 Trajectory - Gyro & Odometry with Clock GPS Augmentation

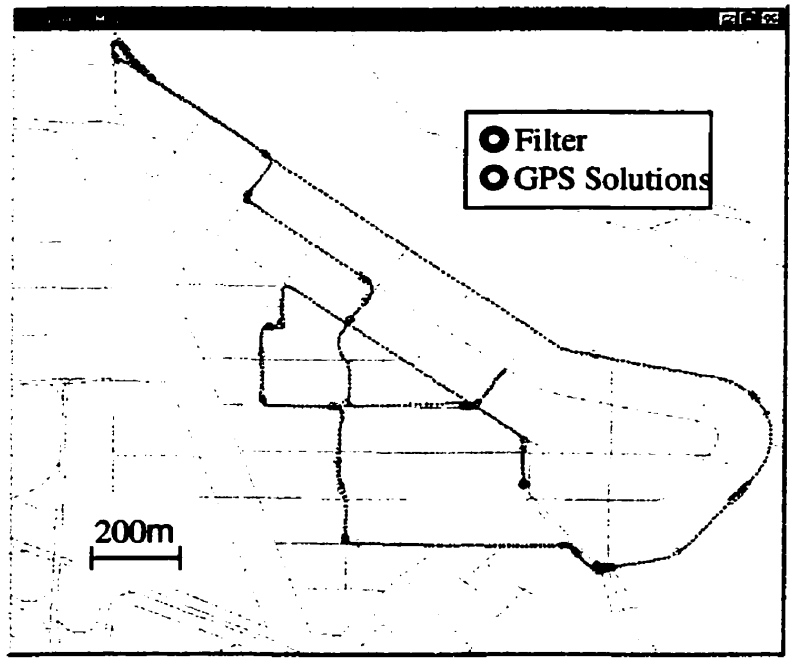

Figure A.12 Bowness Test 1 Trajectory - Gyro & Odometry with Clock & Height GPS Augmentations

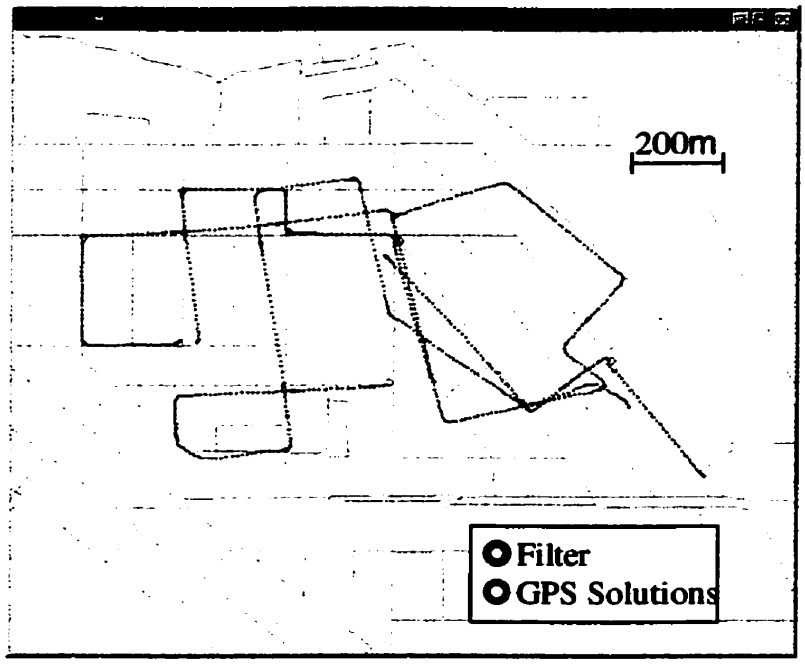

Figure A.13 Bowness Test 2 Trajectory - Differential ABS with No GPS Augmentation

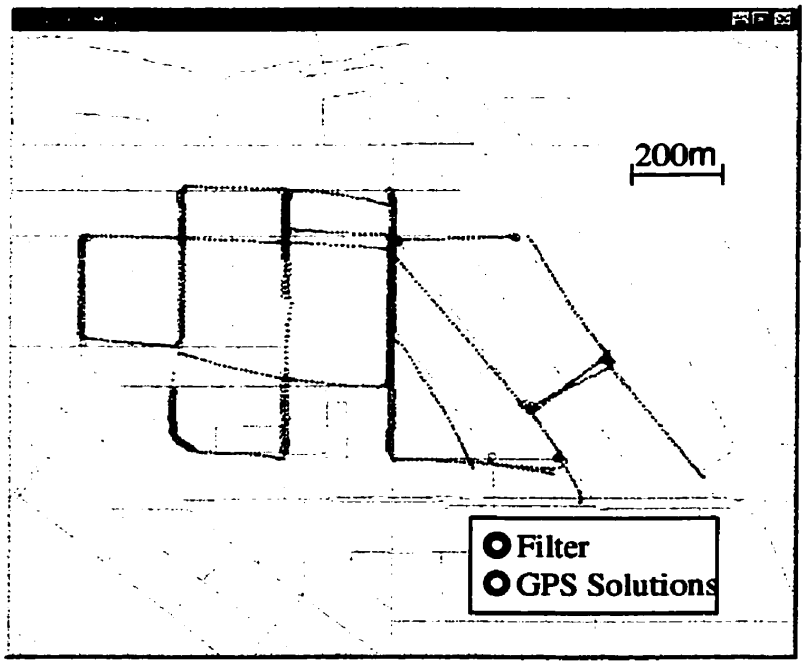

Figure A.14 Bowness Test 2 Trajectory – Differential ABS with Clock GPS Augmentation

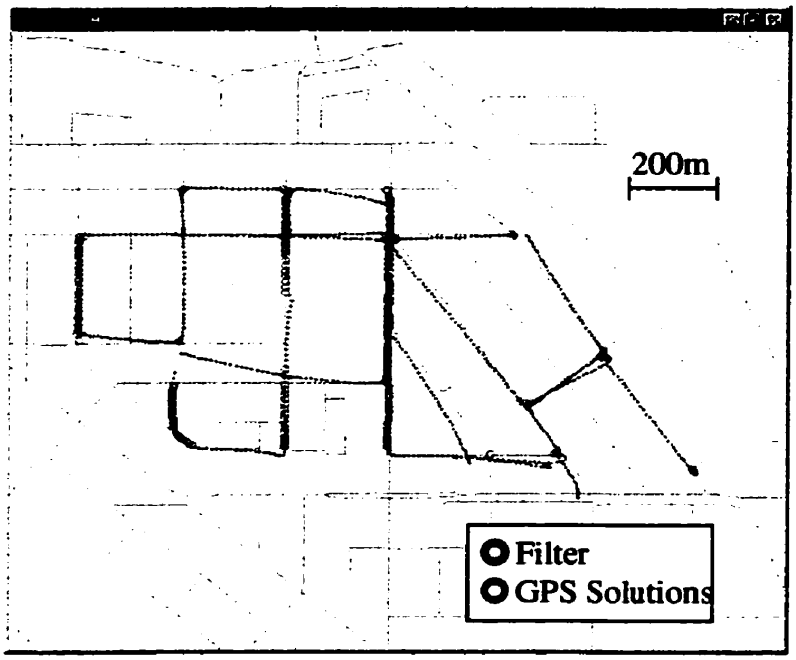

Figure A.15 Bowness Test 2 Trajectory – Differential ABS with Height GPS Augmentation

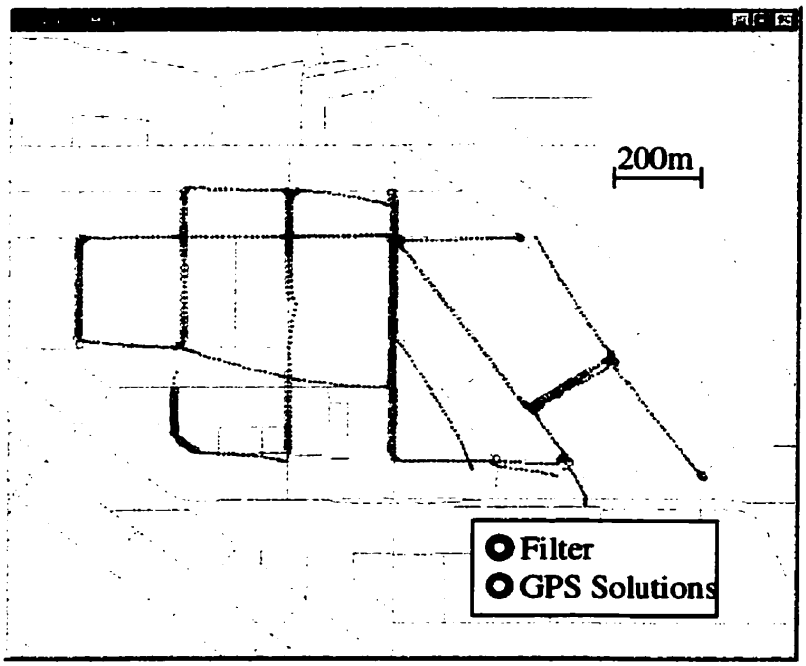

Figure A.16 Bowness Test 2 Trajectory – Differential ABS with Clock & Height GPS Augmentations

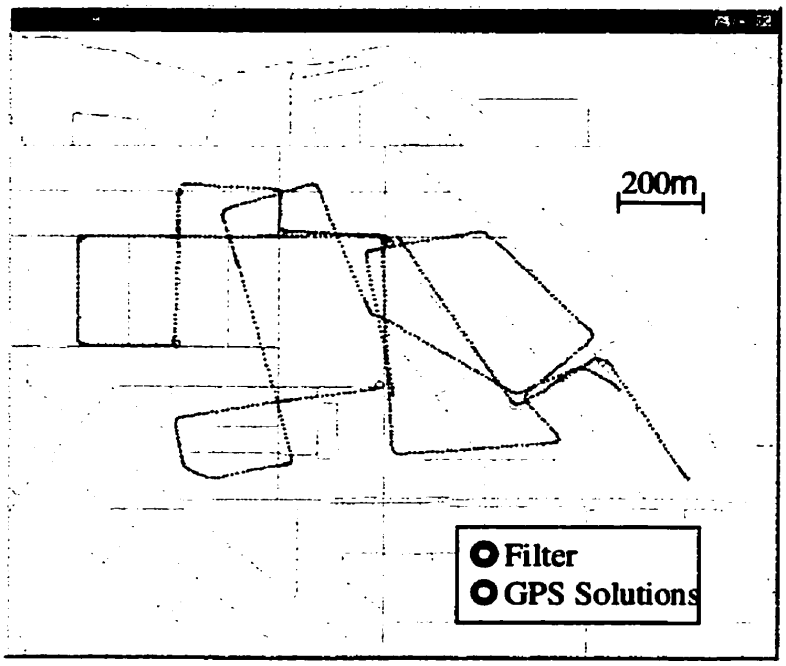

Figure A.17 Bowness Test 2 Trajectory - Gyro & Odometry with No GPS Augmentation

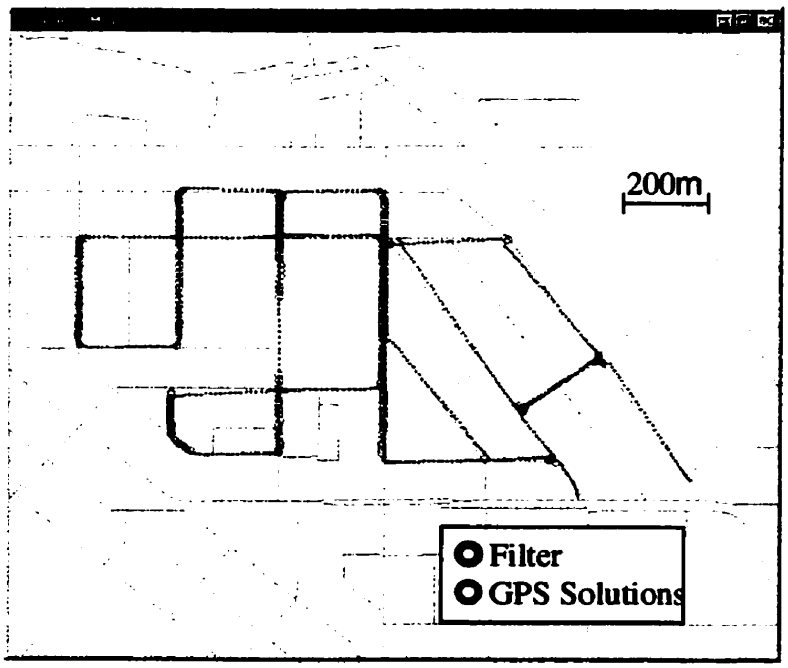

Figure A.18 Bowness Test 2 Trajectory - Gyro & Odometry with Clock GPS Augmentation

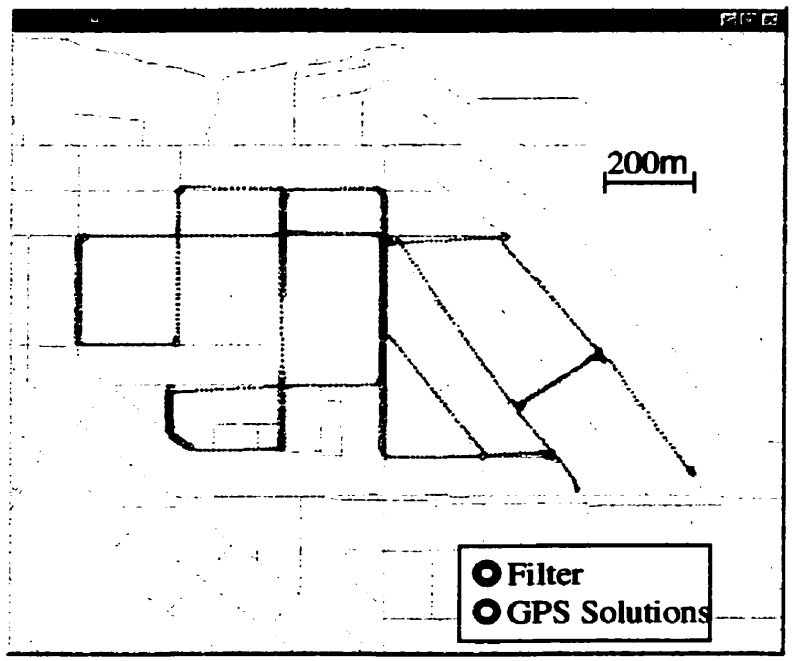

Figure A.19 Bowness Test 2 Trajectory - Gyro & Odometry with Height GPS Augmentation

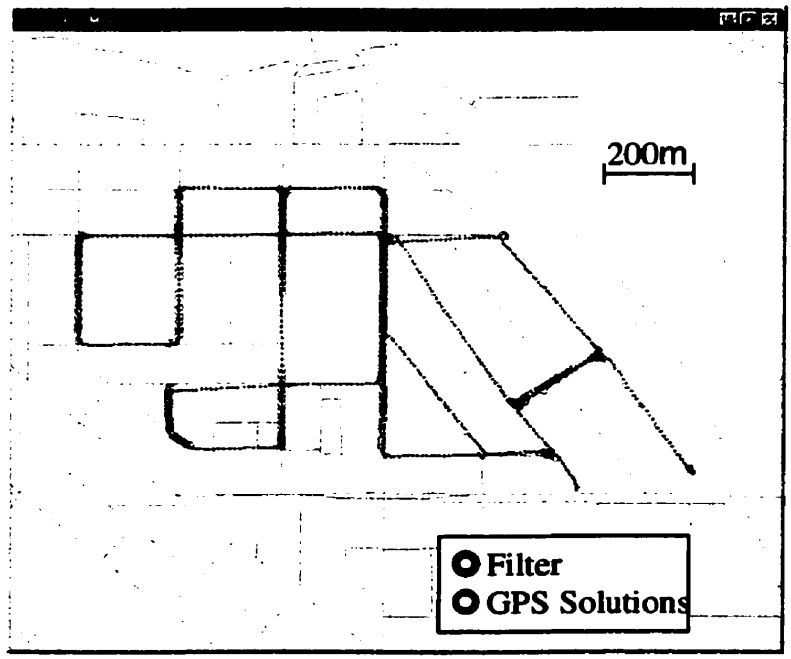

Figure A.20 Bowness Test 2 Trajectory - Gyro & Odometry with Clock & Height GPS Augmentations

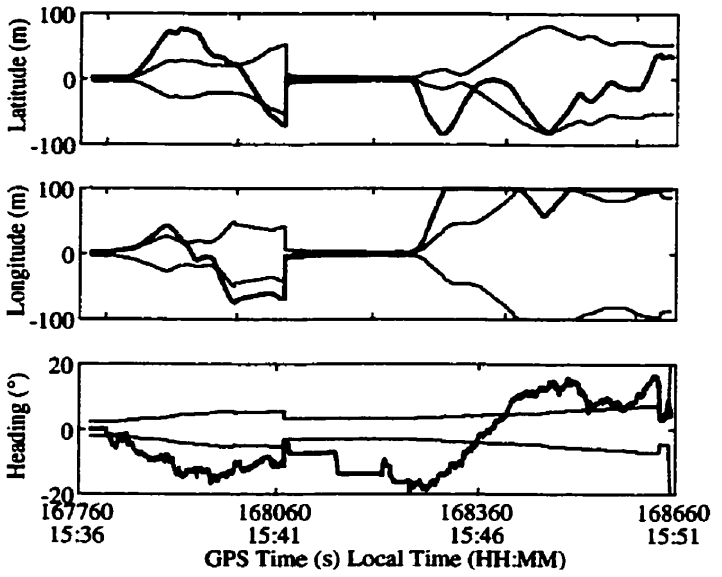

Figure A.21 Bowness Test 1 Errors and Estimated 2σ Bounds – Differential ABS with No GPS Augmentation

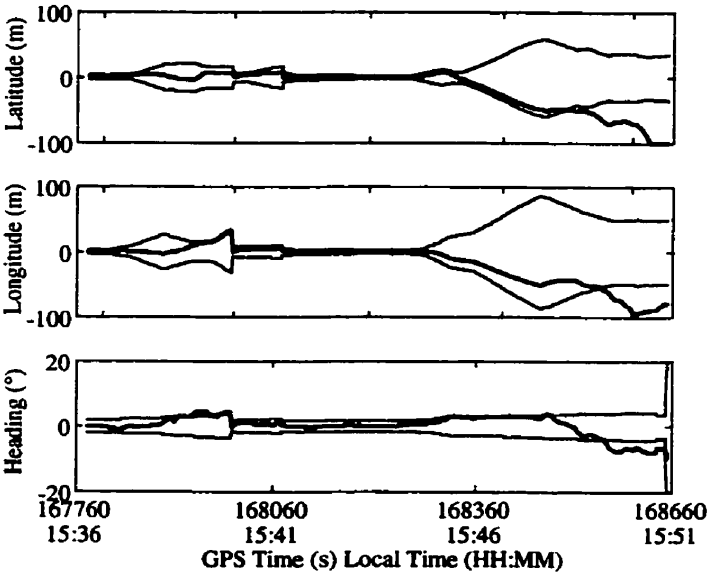

Figure A.22 Bowness Test 1 Errors and Estimated 2σ Bounds - Gyro & Odometry with No GPS Augmentation

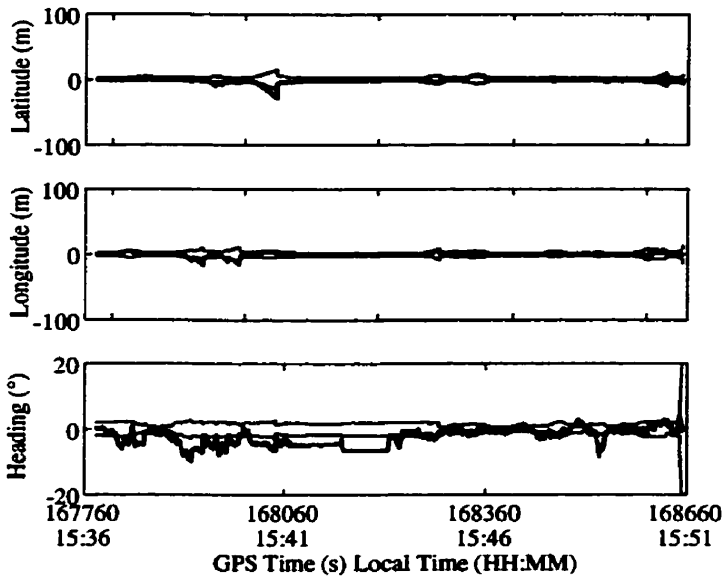

Figure A.23 Bowness Test 1 Errors and Estimated 2σ Bounds – Differential ABS with Height GPS Augmentation

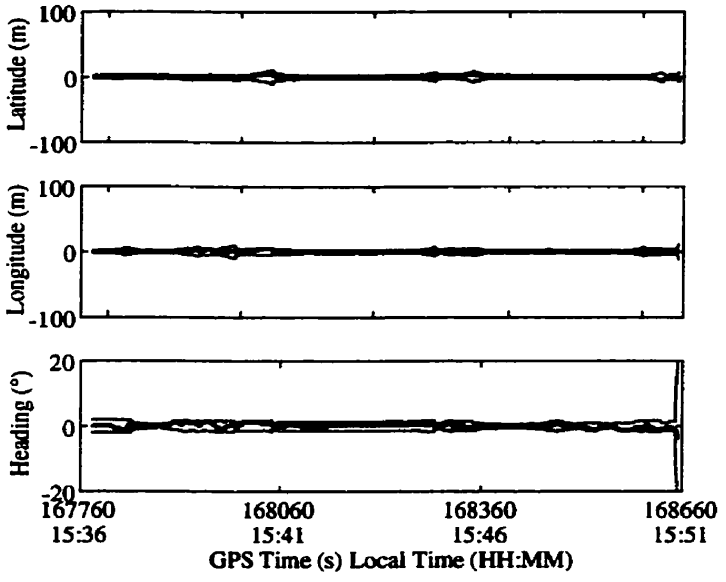

Figure A.24 Bowness Test 1 Errors and Estimated 2σ Bounds - Gyro & Odometry with Height GPS Augmentation

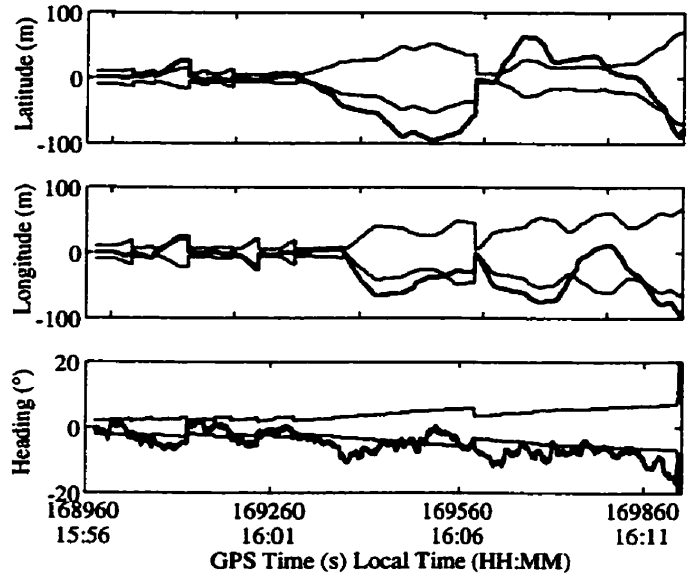

Figure A.25 Bowness Test 2 Errors and Estimated 2σ Bounds – Differential ABS with No GPS Augmentation

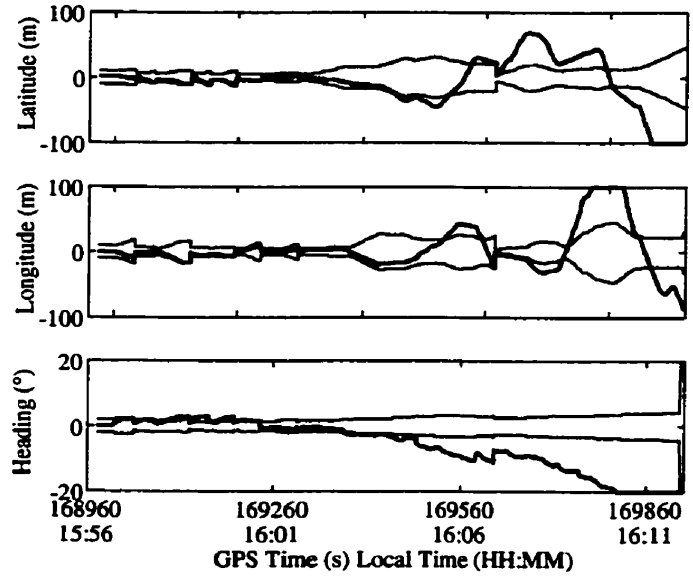

Figure A.26 Bowness Test 2 Errors and Estimated 2σ Bounds - Gyro & Odometry with No GPS Augmentation

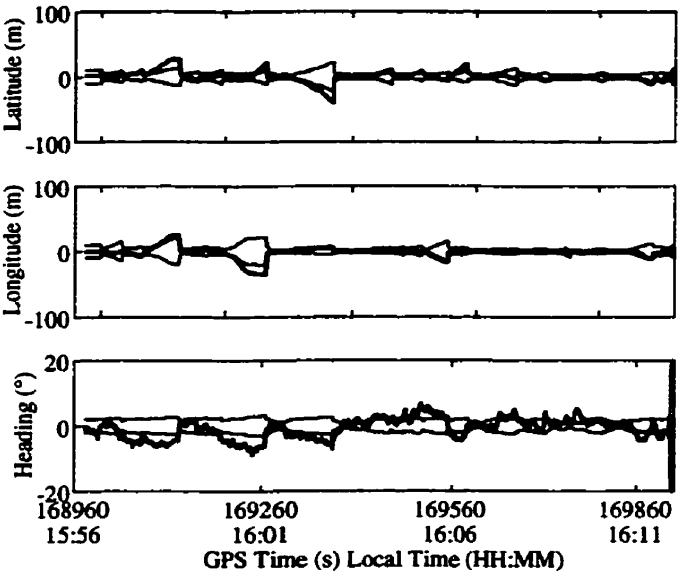

Figure A.27 Bowness Test 2 Errors and Estimated 2σ Bounds - Differential ABS with Height GPS Augmentation

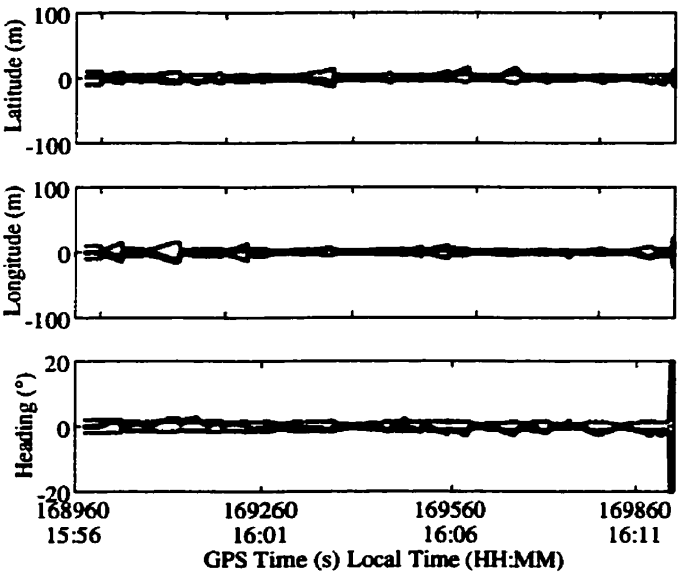

Figure A.28 Bowness Test 2 Errors and Estimated 2σ Bounds - Gyro and Odometry with Height GPS Augmentation

DOWNTOWN TESTS (APRIL 4, 2000)

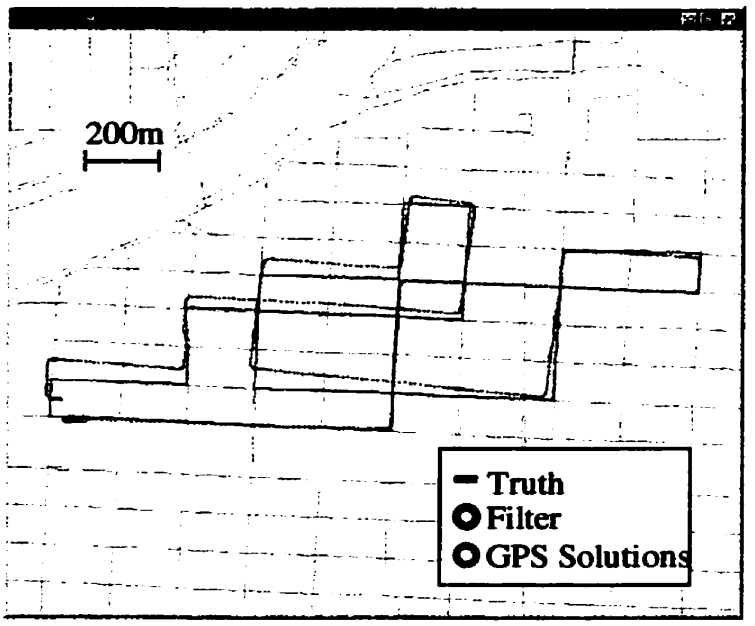

Figure A.29 April 4 Downtown Test Run 2 Trajectory – Gyro & Odometry and Narrow Correlator GPS with No Augmentation

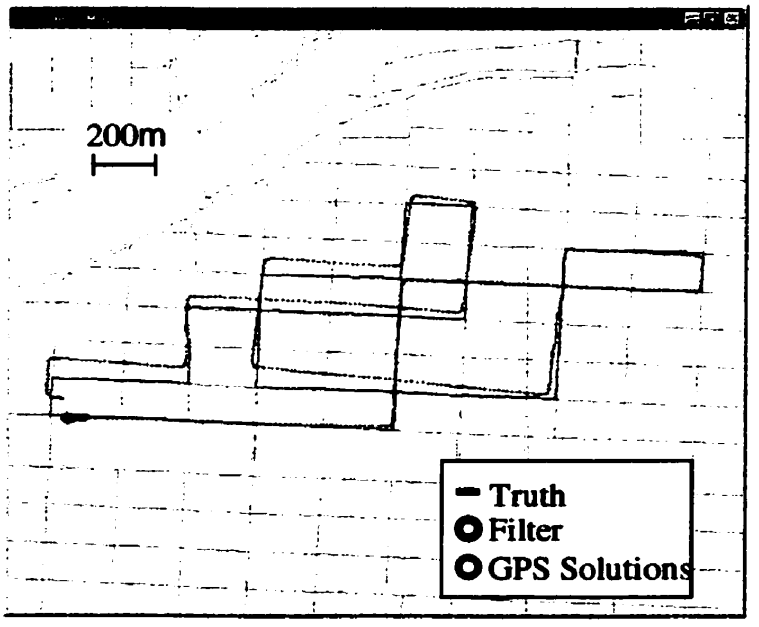

Figure A.30 April 4 Downtown Test Run 2 Trajectory – Gyro & Odometry and Wide Correlator GPS with No Augmentation

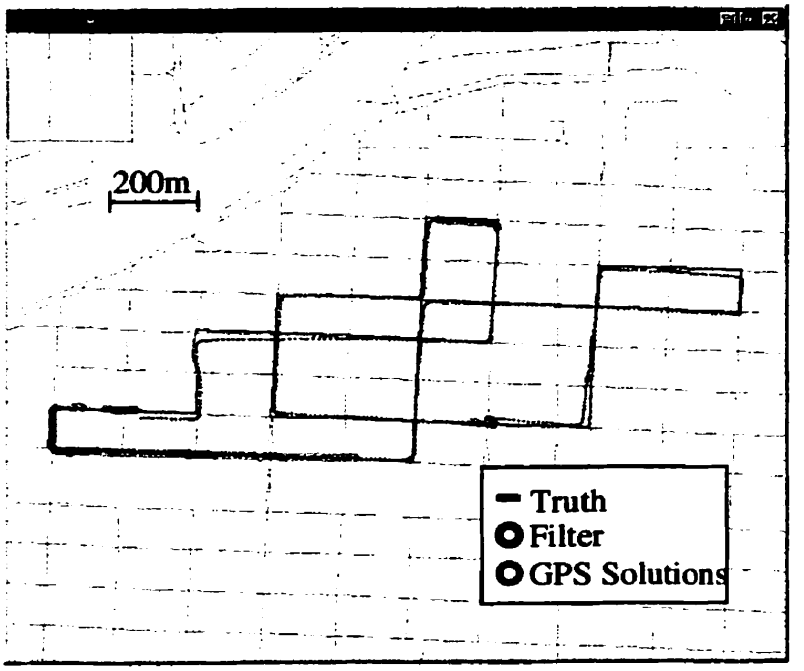

Figure A.31 April 4 Downtown Test Run 2 Trajectory – Gyro & Odometry and Narrow Correlator GPS with Height Augmentation

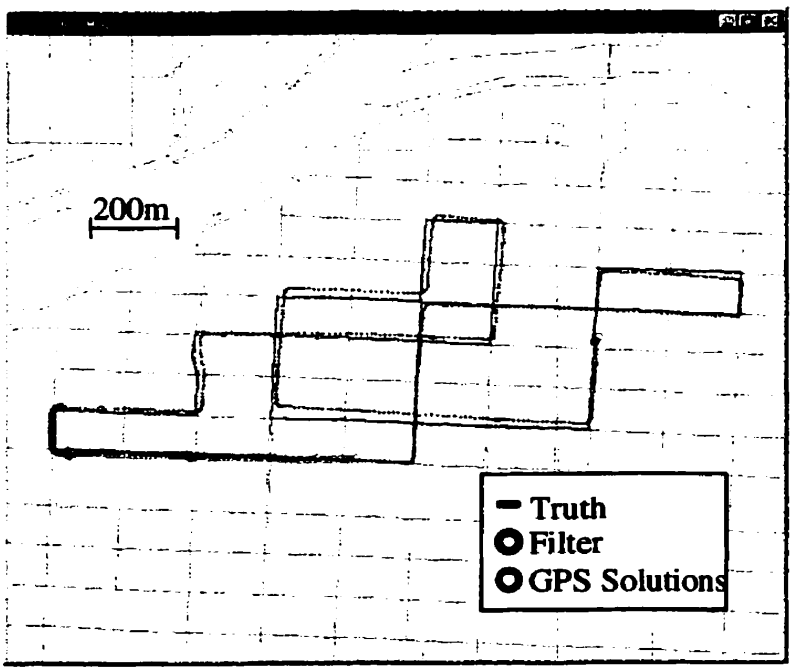

Figure A.32 April 4 Downtown Test Run 2 Trajectory – Gyro & Odometry and Wide Correlator GPS with Height Augmentation

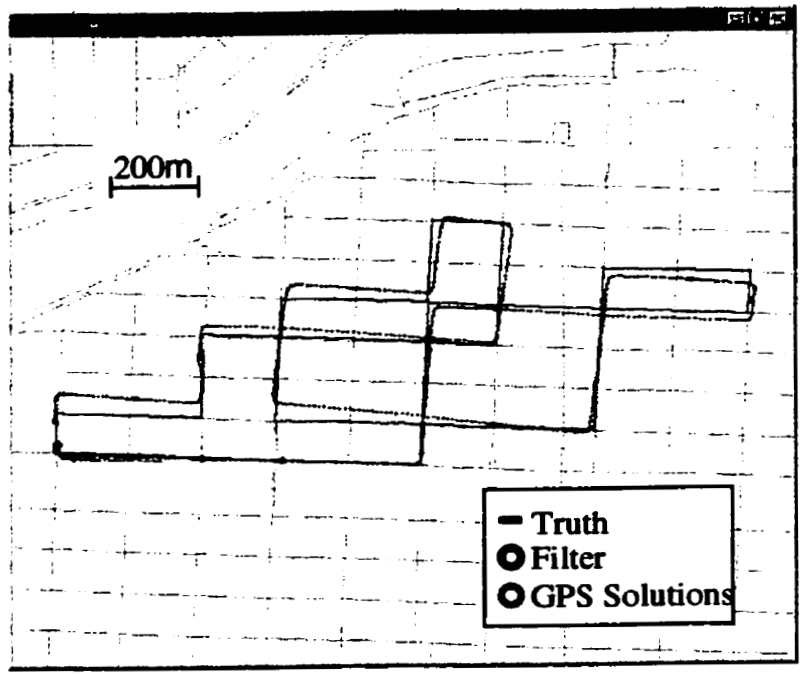

Figure A.33 April 4 Downtown Test Run 3 Trajectory – Gyro & Odometer and Narrow Correlator GPS with No Augmentation

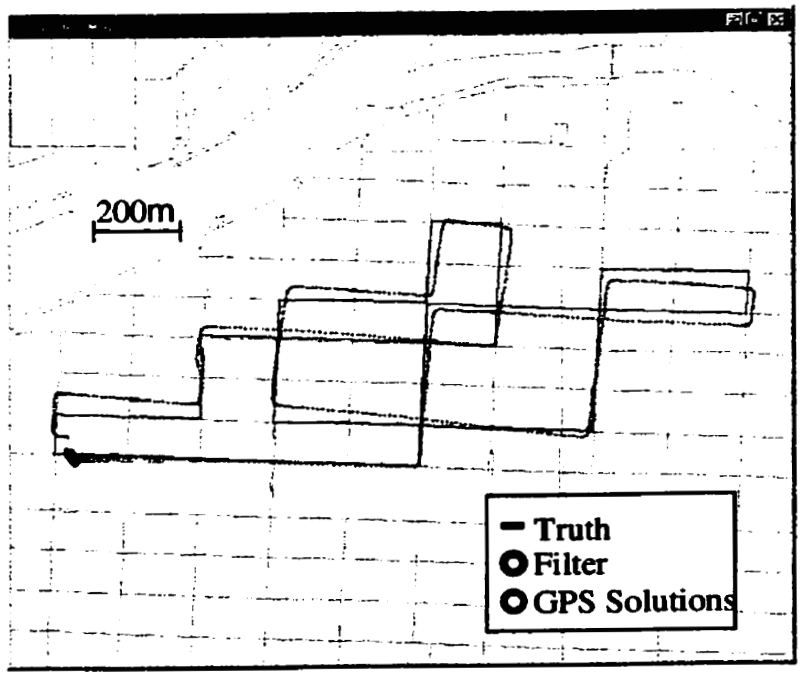

Figure A.34 April 4 Downtown Test Run 3 Trajectory – Gyro & Odometry and Wide Correlator GPS with No Augmentation

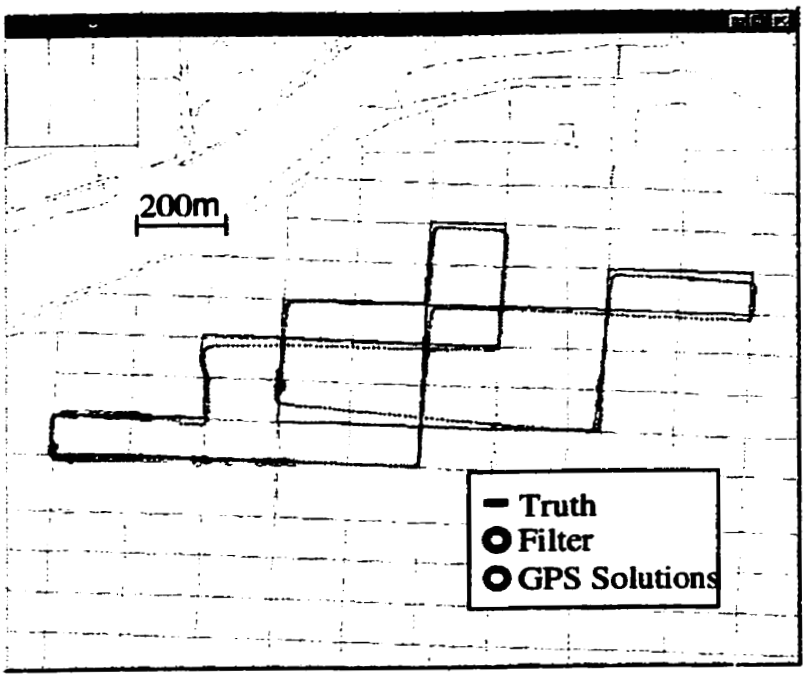

Figure A.35 April 4 Downtown Test Run 3 Trajectory – Gyro & Odometry and Narrow Correlator GPS with Height Augmentation

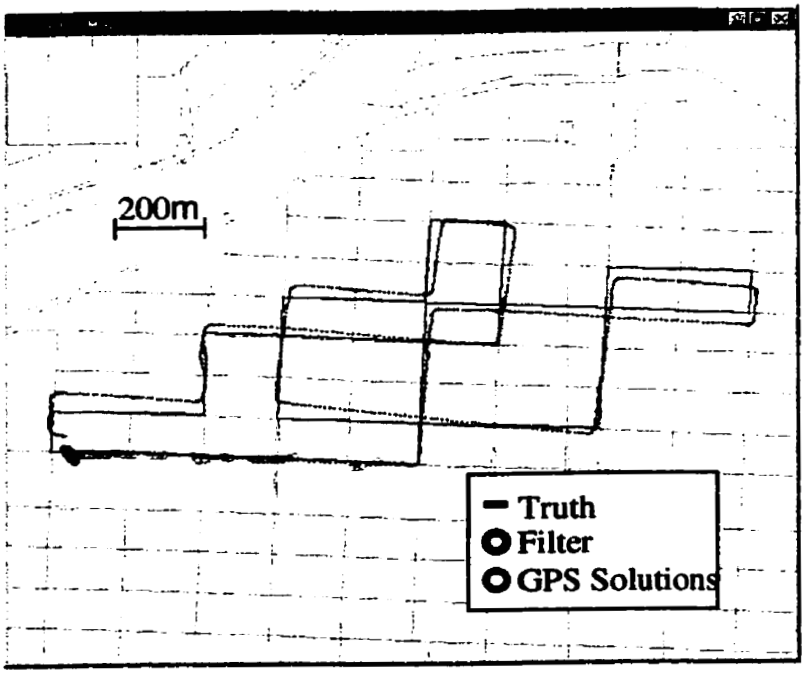

Figure A.36 April 4 Downtown Test Run 3 Trajectory – Gyro & Odometry and Wide Correlator GPS with Height Augmentation

DOWNTOWN TESTS (APRIL 21, 2000)

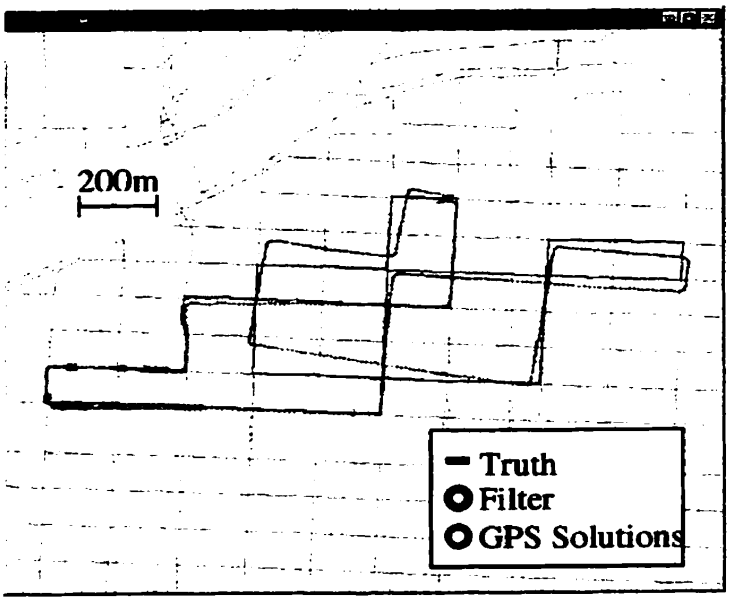

Figure A.37 April 21 Downtown Test Run 1 Trajectory - Gyro 2 & Odometry with No GPS Augmentation

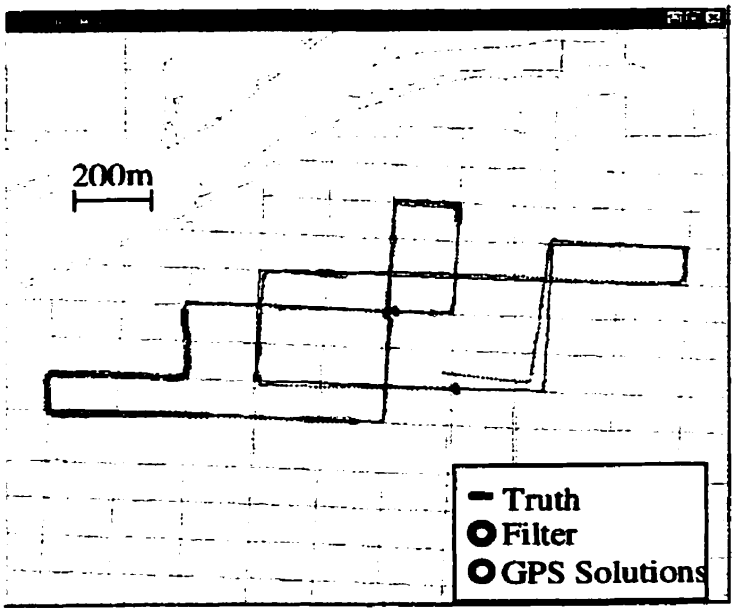

Figure A.38 April 21 Downtown Test Run 1 Trajectory - Gyro 2 & Odometry with Height GPS Augmentation

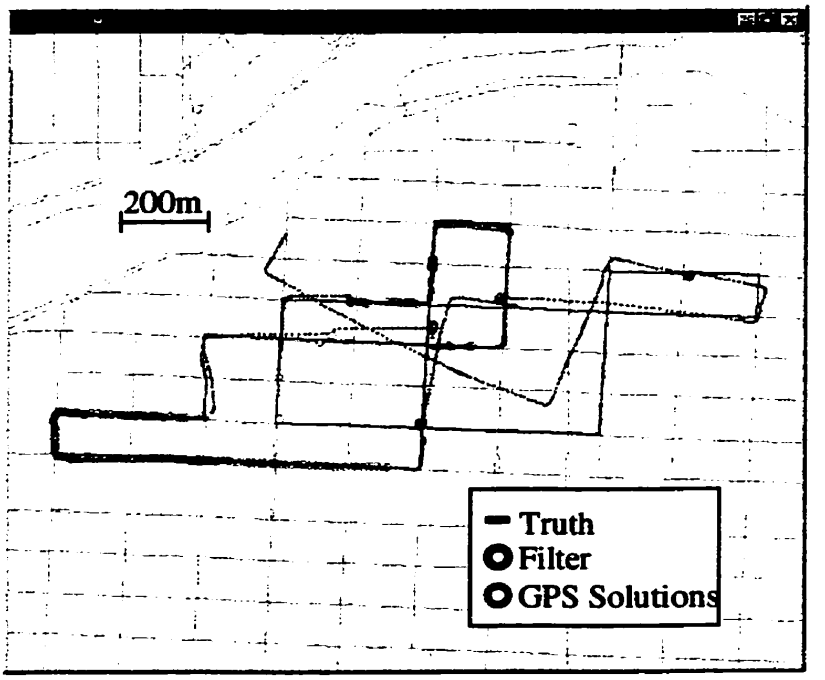

Figure A.39 April 21 Downtown Test Run 2 Trajectory – Differential ABS with No GPS Augmentation

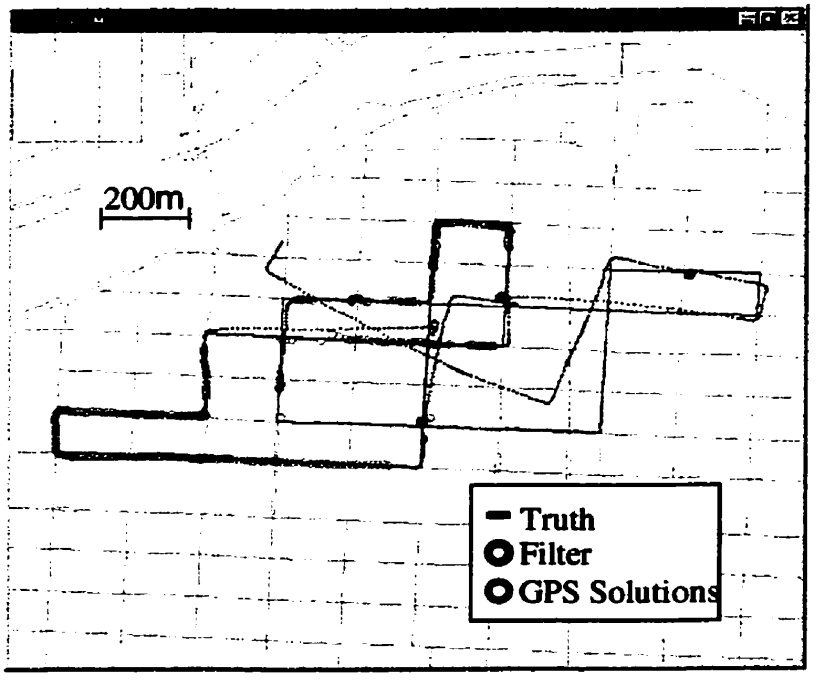

Figure A.40 April 21 Downtown Test Run 2 Trajectory – Differential ABS with Height GPS Augmentation

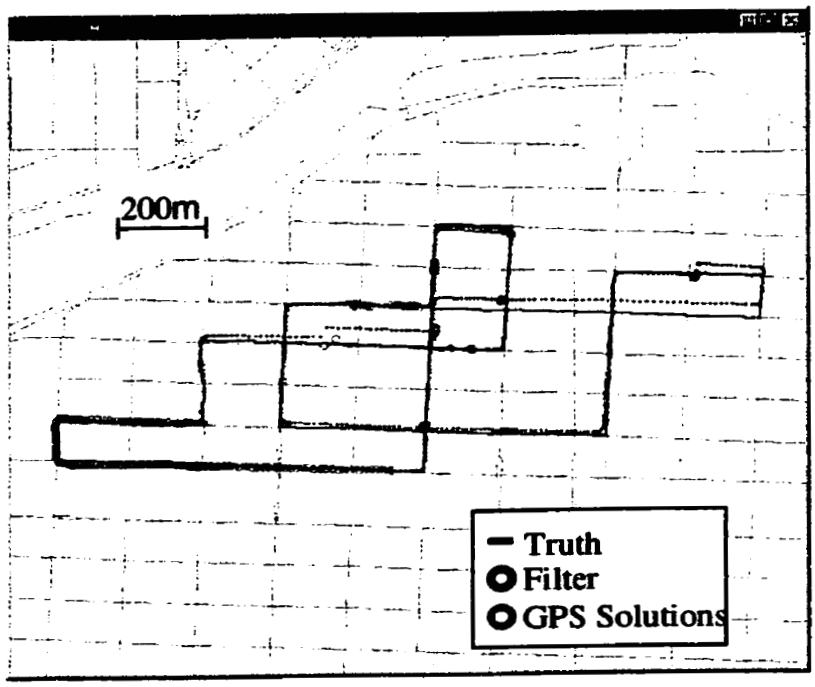

Figure A.41 April 21 Downtown Test Run 2 Trajectory – Gyro 1 & Odometry with No GPS Augmentation

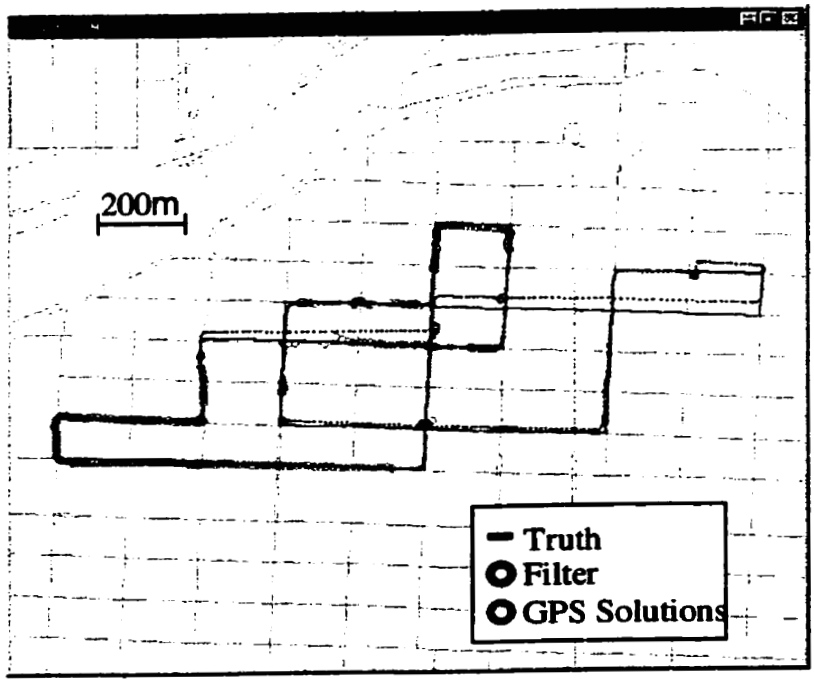

Figure A.42 April 21 Downtown Test Run 2 Trajectory - Gyro 1 & Odometry with Height GPS Augmentation

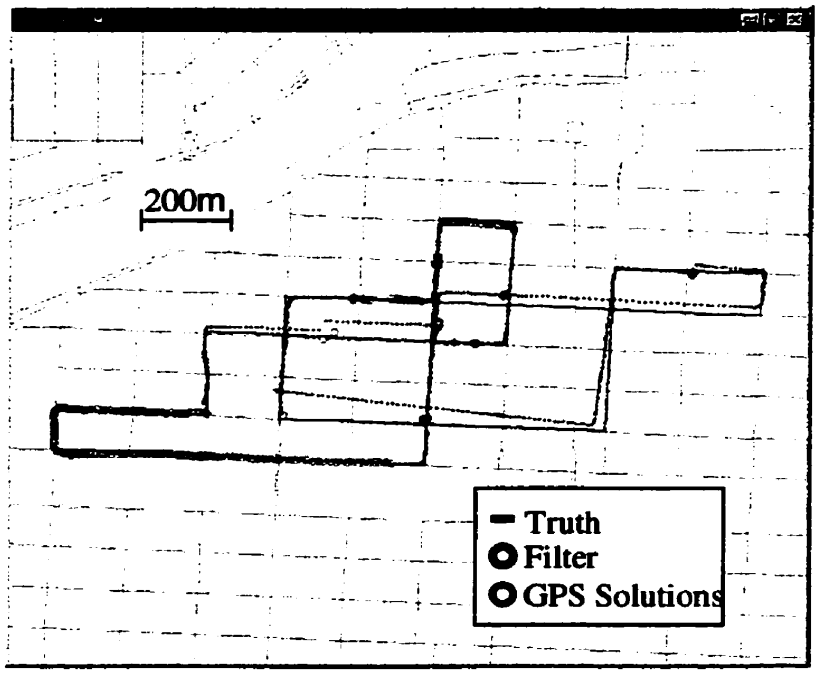

Figure A.43 April 21 Downtown Test Run 2 Trajectory - Gyro 2 & Odometry with No GPS Augmentation

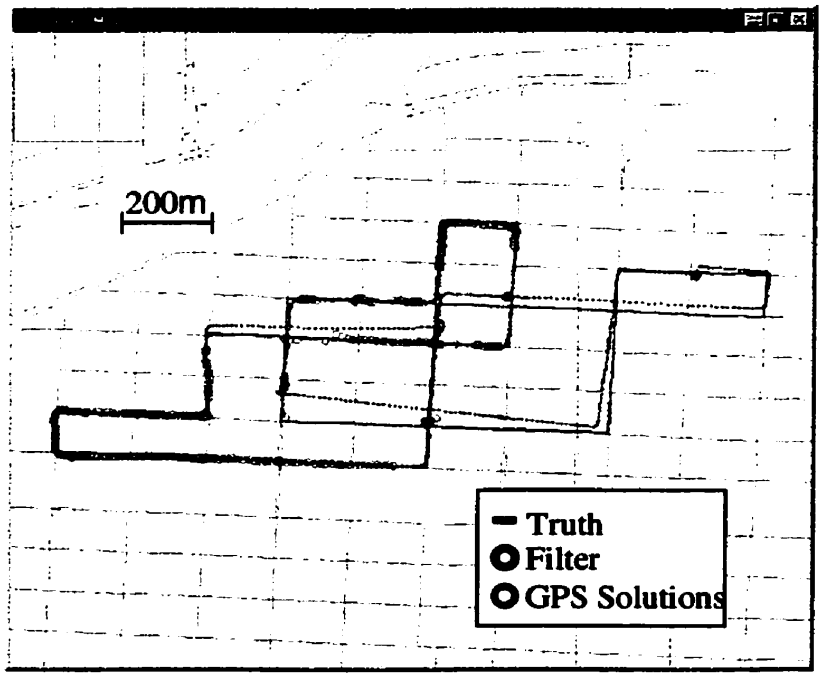

Figure A.44 April 21 Downtown Test Run 2 Trajectory - Gyro 2 & Odometry with Height GPS Augmentation

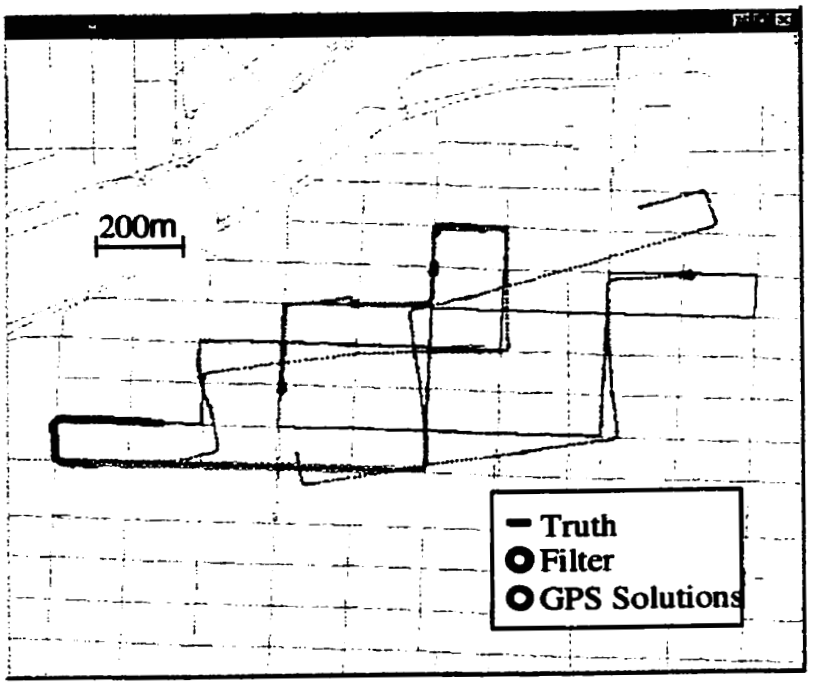

Figure A.45 April 21 Downtown Test Run 3 Trajectory - Differential ABS with No GPS Augmentation

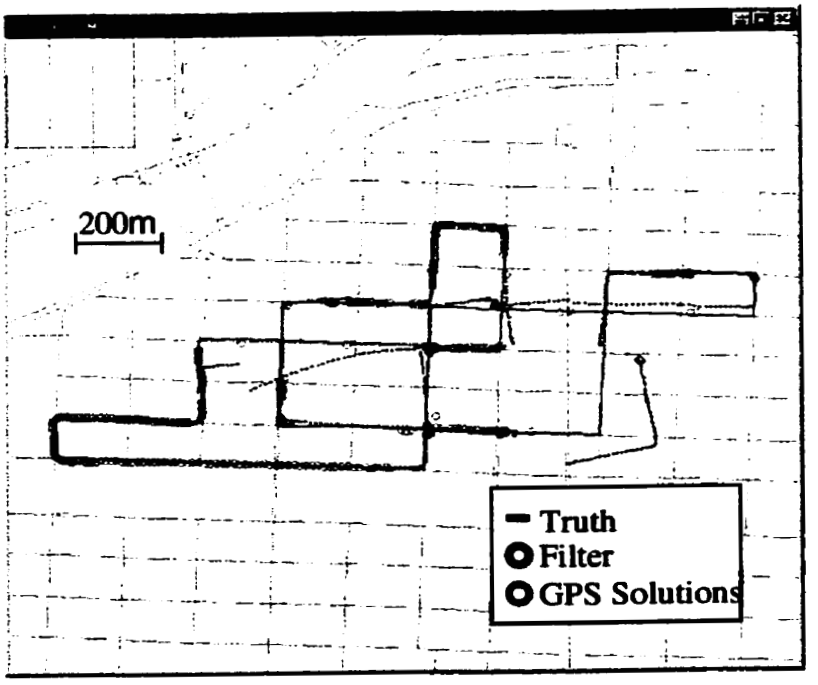

Figure A.46 April 21 Downtown Test Run 3 Trajectory - Differential ABS with Height GPS Augmentation

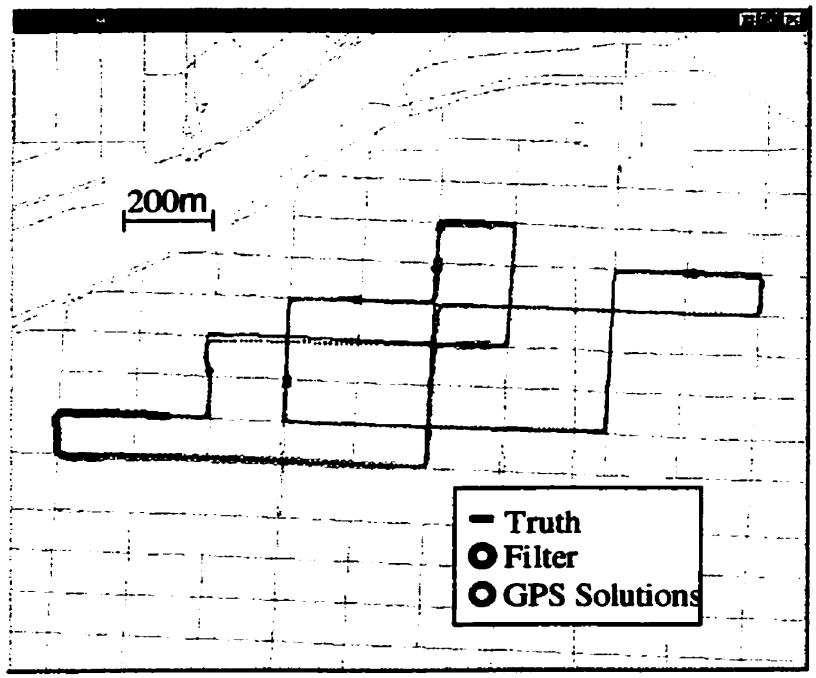

Figure A.47 April 21 Downtown Test Run 3 Trajectory - Gyro 1 & Odometry with No GPS Augmentation

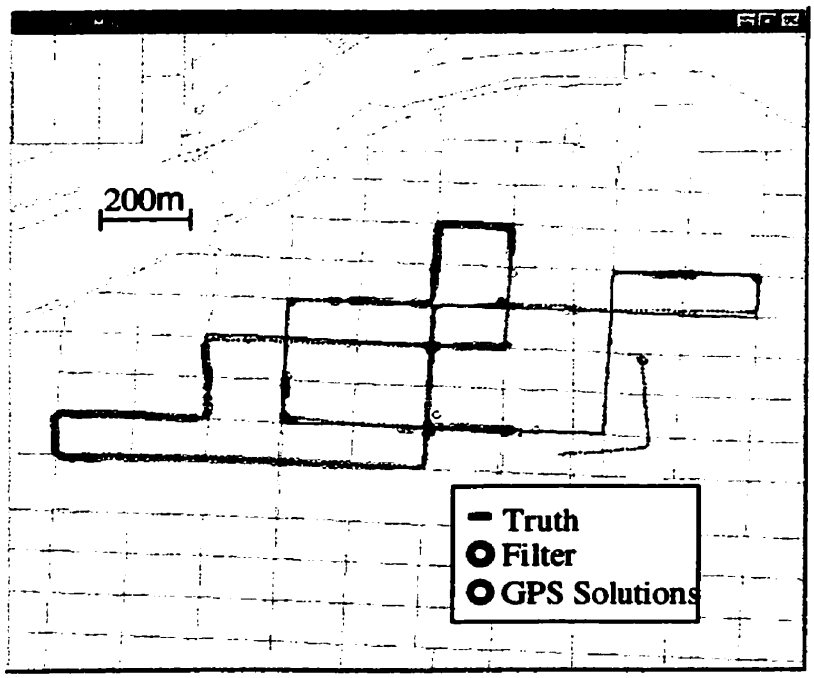

Figure A.48 April 21 Downtown Test Run 3 Trajectory – Gyro 1 & Odometry with Height GPS Augmentation

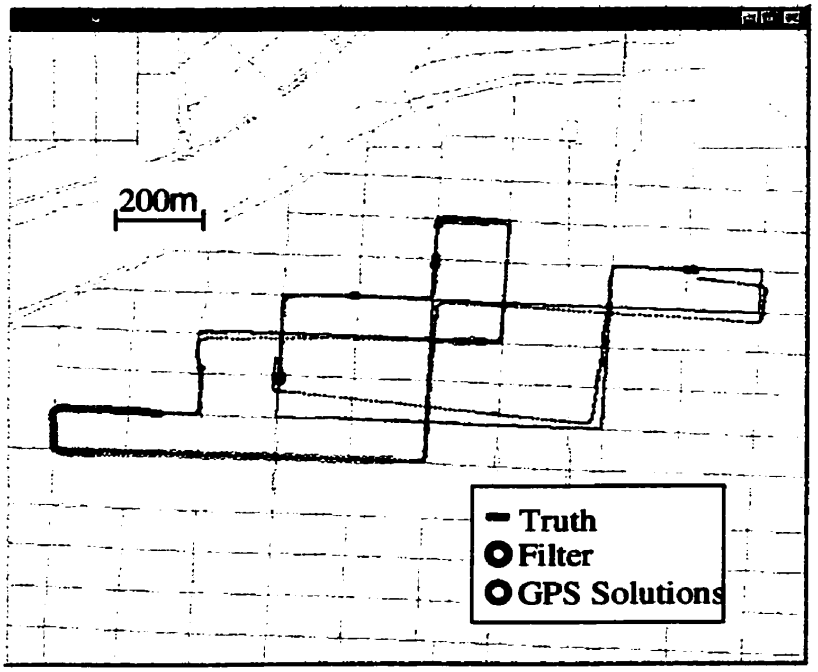

Figure A.49 April 21 Downtown Test Run 3 Trajectory - Gyro 2 & Odometry with No GPS Augmentation

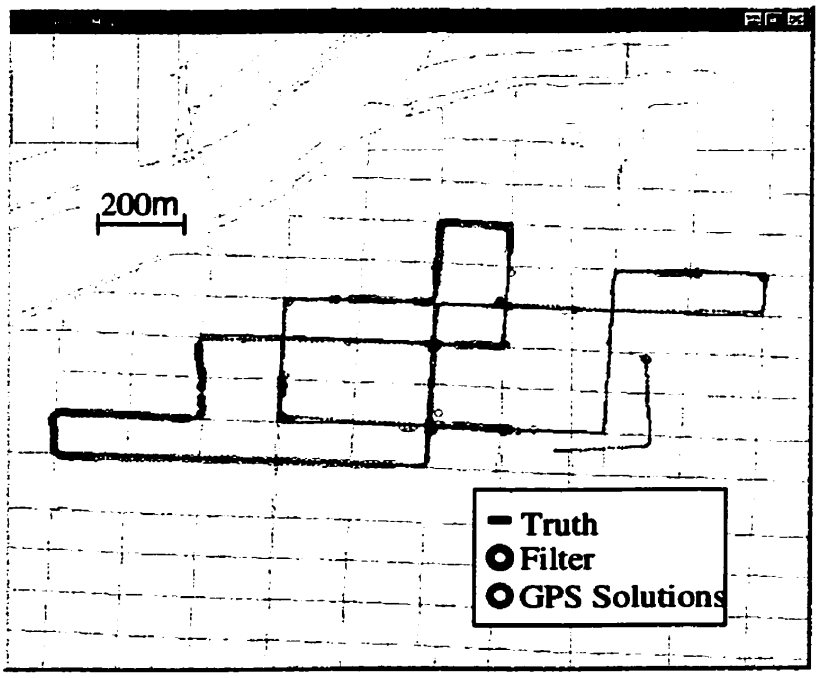

Figure A.50 April 21 Downtown Test Run 3 Trajectory - Gyro 2 & Odometry with Height GPS Augmentation