Development and Testing of a GPS-Augmented Multi-Sensor Vehicle Navigation System

J. Stephen and G. Lachapelle

(*Department of Geomatics Engineering*, *The University of Calgary*)

An integrated multi-sensor vehicle navigation system is presented that uses a low-cost rate gyro and differential odometry to supplement GPS under signal masking conditions such as tree foliage and urban canyons. Signal masking is often accompanied by extreme multi-path in urban centres with tall buildings, and is also found in wooded areas, enclosed car parks, tunnels, etc. The purpose of the system tested is to provide an accuracy of better than 20 metres almost 100 % of the time throughout these interruptions, which are assumed to last up to a few minutes. The equipment used is discussed in detail, as is the method used for filtering measurements. Results are presented from tests carried out in an urban core with relatively long periods of signal loss – up to several minutes over a 6-km test circuit. Tests in urban canyons demonstrate that it is difficult to reach the above specifications with aiding from differential odometry alone due to the high precision of the wheel-scale factor required. However, with the use of a rate gyro and odometry, RMS errors are below 20 metres while availability is nearly 100 %. Some of the large deviations could probably be better controlled if GPS multi-path errors were detected before they are allowed to corrupt the filtered solution.

KEY WORDS

1. GPS. 2. Augmentation. 3. Vehicle Navigation.

1. INTRODUCTION. The prototype navigation system described attempts to improve the inadequacies of GPS vehicle navigation through augmentation with DR (dead reckoning) sensors. GPS signals, transmitted in the upper part of the UHF spectrum, are essentially line-of-sight signals that suffer attenuation from signal masking and reflection (multi-path) in areas such as urban canyons, wooded areas, tunnels, and enclosed car parks. The problem is further compounded by a free space transmission loss of 184 dB from the satellites located some 20 000 km above the Earth. As a consequence, current GPS receiver technology is still limited to line-of-sight, although emerging technologies are expected to result in significant improvements in signal tracking under masking conditions.

Many applications, such as those listed in Table 1, require continuous and accurate information that cannot be achieved using GPS alone. Accuracy requirements vary by application, but most require knowledge of the exact street and block. Others that require much higher accuracy, such as lane identification, are not considered.

Various sensors are available for augmenting GPS, some of which are listed in Table 2. They vary greatly in accuracy and price, so they are chosen as a function of

Table 1. Vehicle position data applications.

Application	Desired accuracy (m)
Commercial fleet management	100
Ambulance/Fire/Police calls	10
Location of services, e.g. restaurants	10
Route finding	25
Retrieval of lost/stolen vehicles	10
Detection of road departures	0·1

Table 2. Sample aiding sensors (Bullock, 1995).

	Accuracy	Comments
Heading Sensors		
Vibrating rate gyro	0·1°/s	Temperature Sensitive
Mechanical rate gyro	0·03°/s	Short life
Fiber-optic gyro	0·001°/s	Expensive
Differential Odometry	0·01°/m	Difficult to calibrate
Fluxgate compass	2–4°	Sensitive to iron
Distance Sensors		
Wheel odometer	0·01 %	Some have cut-off speed
Transmission odometer	0·05 %	No differential capability
Doppler radar	1 %	Road irregularities
Other Aiding		
Barometer-height	3 m	Local variations
Inclinometer-attitude	2°	Acceleration error

the application (Bullock, 1995). For vehicle location and navigation, sensors must be chosen that cause minimal additional cost to the production or modification of a vehicle, while delivering a high position availability with an accuracy of 10–20 metres.

2. INTEGRATED SYSTEM COMPONENTS. The dead reckoning sensors selected to augment differential GPS were: a low-cost piezoelectric vibrating gyro, namely a Murata Gyrostar (Murata, 1999), and differential odometry provided by the anti-lock break system (ABS) of the test vehicle. Differential GPS was used as the tests took place prior to termination of GPS Selective Availability on May 1, 2000. Additional equipment included two NovAtel receivers for differential GPS operation, a 16-bit National Instruments A/D converter and two laptop PCs for the data recording.

Attempts to calibrate the local magnetic field of the test vehicle in order to use a digital compass were unsuccessful, as was experimentation with a bubble-level inclinometer. Data was processed off-line to allow for a comparison of different filtering techniques and tuning of the filter parameters. An extended Kalman filter was used to combine computed GPS positions with DR data, while independent filters were employed to calibrate the aiding sensors. This implementation is very robust and requires no intervention or tuning.

A front-wheel drive 1998 Honda Civic with ABS system was used as the test vehicle. The external sensors were fixed to an aluminum plate, which was tied to a rack on the roof of the vehicle. The laptop computers, A/D converter, and GPS receiver were located on the back seat of the vehicle. The ABS system was located under the dashboard on the passenger side. 12V power supplies were used to power the equipment.

2.1. *Differential Odometry*. The test vehicle has Hall effect-type odometers that measure the rotation of each of the four wheels. The cut-off speed on this type of odometer is very low, namely about 0·5 m/s. A sinusoidal signal is produced by the wheel pick-ups with nearly constant amplitude and a frequency proportional to wheel speed. The voltage levels of the rear ABS sensors were sampled at 2000 Hz. A digital counter was implemented in software that gave the number of zero-crossings of the sinusoid at 10 Hz intervals. The sinusoid has an amplitude of about 3 V, so a hysteresis of 0·5 V was used to avoid problems with noise causing multiple crossings. Each wheel's rotation is available with a discrete quantization of about 6°, or 2·5 cm, since approximately 60 zero-crossings occur for each rotation of the wheel. Other ABS system inputs, including a brake strength indicator, reverse indicator, and park brake indicator, were available at the ABS control box. These were useful in interpreting the data; for example, when to calibrate and whether the measured speed is in the forward or reverse direction.

The ratio between the count of zero-crossings and the corresponding distance travelled by a wheel must be estimated as a scale factor. The wheels have slightly different sizes and slippages due to tread wear, weight distribution in the car, and numerous other factors. Slippage occurs when the contact point between road and tyre must support a transverse load and is increased during acceleration; for instance, a heading or speed change. Slippage is dependent on both type and condition of the tyres and road surface. With modern tyres and normal driving habits, slippage seldom exceeds 1 % of the distance travelled, with the exception of skidding occurrences. Due to acceleration and frictional deceleration of the engine, wheel slippage is much higher for the powered wheels (connected to the powertrain) than for non-powered wheels. To minimize the error caused by slippage, the non-powered wheels are normally used to aid in navigation. The differential allows the wheels to rotate at different speeds in order to minimize the amount of wheel slippage on corners, but this same mechanism also allows the wheels to slip differently due to the surface varying under the tyres. The road surface variations caused by water, ice, snow, and gravel cause the friction to vary greatly between tyres and change 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 rotating.

Due to the complex nature and variability of slippage related errors, potential slippage occurrences should be identified whenever possible. Even with data monitoring, slippage (particularly skidding) requires that the user must always be cautious with this data. Large slippage occurrences normally amount to a decimetre or less, so they do not constitute a substantial error source for simple distance odometry. For differential odometry, such an error can be shown to have a large impact.

In conjunction with some varying amount of slippage, the circumference of the wheels also varies between tyres and as a function of environmental conditions. Modern radial tyres have fairly constant size and shape, unlike the bias-ply tyres that

preceded them. Studies have shown that several factors affect the tyre radius, including temperature, speed, pressure, and weight loading. The estimated individual contribution of these factors is given in Table 3 (Stephen, 2000).

Table 3. Wheel circumference variation sources.

Error factor	Possible error in radius
Pressure	1 mm/psi
Temperature	2 mm Expansion + 1 psi/10 °C
Wear	5 mm
Speed	1 mm + 50 °C or more
Weight	1 mm/100 kg

The shape of a tyre is slightly affected by speed, but the temperature (and thereby pressure) varies greatly due to the increased frictional heat that must be dissipated by the tyres. The air temperature inside the tyres can change by more than 50 °C due to speed, as well as 50 °C in ambient temperature. A change to the loading of the vehicle or adjustment of the air pressure will also affect the effective radius of the tyre. It is necessary to calibrate the wheel size continually since the above variations amount to more than 1% error in distance alone, or in 10 m of error for each km of driving.

A digital LPF (Low Pass Filter) of the ratio between GPS-derived speed and the frequency of zero-crossings was used to calibrate the tyre sizes. The GPS speed was derived from the horizontal components of a DGPS Doppler velocity solution. A cut-off speed was used to prevent calibration when the vehicle was stationary or the relative accuracy of GPS was low. The filter can be written as:

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

where: w_s is the wheel-scale factor. The gain factor α was set to $1{\cdot}0 \times 10^4$. It was also found that the scale factor should not be calibrated while turning, since it would often change by more than 0·05% on a single turn because of the action of the differential. If the gain were set lower, the average of left and right turns would likely cancel out this effect, but a higher gain factor was needed to follow the changes in scale factor that occurred during the tests.

The action of the differential allows the inside wheel to travel a shorter distance than the outside wheel during a turn to avoid slippage and skidding. The distance travelled by each wheel therefore varies from the overall distance travelled by the vehicle while turning. To compensate for this, the average of the left and right distances was used as an estimate of the distance travelled over an interval, i.e.

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

The difference between the distances travelled by each of the two tyres also gives a measure of the heading change. The method suggested by Harris (1989), illustrated in Figure 1, can be formulated as

$$\Delta az = \frac{\Delta d_L - \Delta d_R}{Track\ Width}. \tag{3}$$

https://doi.org/10.1017/S0373463301001357 Published online by Cambridge University Press

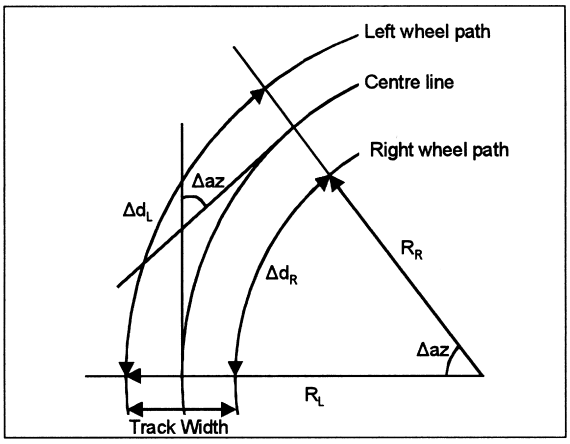

Figure 1. Geometry of differential odometry.

The track width is the effective turning width of the car (approximately the distance between tyre 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 travelled and heading change (Zavoli *et al*., 1988) that accounts for the change in effective track width of the vehicle as a function of heading rate. Banking of the road surface decreases the difference in distance travelled by the two wheels. Using a 90° bank, there would be no difference, while the maximum difference would occur when the surface is level. This is not usually a problem, since GPS signal masking is not usually associated with areas where high speed banked turns occur.

Scale factor errors and slippage become a serious concern in differential odometry. If a vehicle with a track width of 1·5 m is coming to a stop on patchy ice, and one wheel slips while the other is locked for just 10 cm, a 4° heading error occurs. The result of a 0·5% error in the scale factor for a single tyre is 0·2°/m, accumulating 2° for each 100 m segment travelled. To meet the specifications desired for the system, scale factors would need to be calibrated to an accuracy of several hundredths of a percent, which is a difficult task given that GPS-derived speed measurements are made with an accuracy of 1 to 5%, and previously discussed variations in wheel size. Some amount of drift and slippage error is inevitable using differential odometry for heading determination.

Since the distance travelled by each wheel is measured over an interval, but a speed is desired, the variance of the speed can be estimated from the variance of the distance measurement as:

$$\sigma_{speed} = \frac{\sigma_{dist}}{\Delta t}. \tag{4}$$

Figure 2 depicts the difference between the GPS-derived heading and the integrated ABS heading for an open-sky test (the GPS heading should be accurate to about 1°). Throughout the test, the integrated bias changes by over 30°. Without absolute heading calibration from GPS, the estimated vehicle heading would be in error by more than 30° at the end of the 20-minute test.

https://doi.org/10.1017/S0373463301001357 Published online by Cambridge University Press

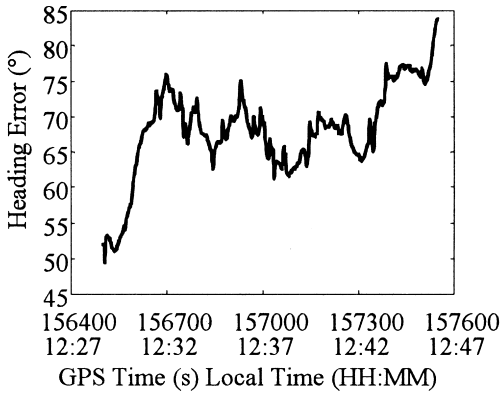

Figure 2. ABS heading error.

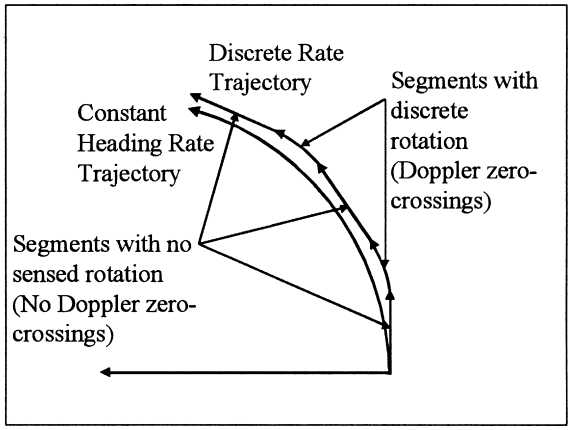

Figure 3. Continuous vs discrete heading rate.

Using the accumulated heading results in a serious shortcoming when it comes to statistical treatment. Accumulated heading should not be dealt with as a relative heading since slippage and loss of lock create errors that cannot be accounted for. After the first few tests, it was decided that a heading rate model should be used, since the rate model described the relative nature of the measurements. The derivative of the heading change equation (3) for differential odometry was employed.

A phase lock loop (PLL) was initially implemented to track the ABS sinusoidal signal that gave much more accurate estimates of the rotation rate than the discrete count of zero-crossings. Unfortunately, loss of signal lock occurred frequently because the bandwidth had to be made large enough to handle the high frequency dynamics, thereby making the filter unstable. If lock was lost on a single wheel during a corner, no heading change information was available for periods of up to several seconds while signal lock was reacquired, resulting in large errors. Instead, heading rate measurements were created by accumulating the zero crossings of the sinusoids over an interval of 0·1 s. Counting zero crossings creates a measurement that is much noisier than using frequency estimates from the PLLs, but the observations are continuously available.

Heading rates are therefore quantized in discrete steps of about 10°/s because the difference between distances travelled by the two wheels is measured in zero crossings of about 2·5 cm. The difference of the two signals can be interpreted as a Doppler signal; the number of zero-crossings of this Doppler signal during the interval gives 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 travelled 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 would be observed every second interval, while no rotation would be 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 3. However, the discrete nature of this process is somewhat countered by the constant heading rate model. A variance of 0·1°/s was chosen for the heading rate observation variance after extensive experimentation.

2.2. *Rate Gyro*. The Murata Gyrostar piezoelectric vibrating 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, which is sampled at 2000 Hz by the A/D converter. The samples were averaged at 10 Hz to give an average heading rate over the interval. The specifications of the automotive version of the sensor (ENV-05D series) are given in Table 4 (Murata, 1999).

Table 4. Murata Gyrostar specifications.

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 Rate = 0	2·2	2·5	2·8	VDC
Scale Factor	19·3	22·2	25·1	mVs/°C
Asymmetry	–	–	3	°/s
Temperature Scale	–	–	±10	% FS
Temperature Drift	–	–	9	°/s
Operating Temp.	−30	–	80	°C
Noise Level (7 kHz)	–	–	20	MV
Dimensions	18 × 30 × 41 mm			

Piezoelectric materials exhibit the PE (piezoelectric) effect, in which a vibrational motion of the crystals creates an electric potential within the material. The reverse PE effect is also used, whereby application of an electric field to a PE 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. Applications usually make use of resonant or harmonic frequencies of the piece of material, which are a function of its size, shape, and dielectric properties.

Two basic types of PE gyros 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 a reverse 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. 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 by the cross product

$$\bar{F}_c = 2m(\bar{v} \times \bar{\Omega}), \tag{5}$$

where:

m is the mass of the object,
v is the velocity vector of the object, and
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 transverse vibration of the piezoelectric crystal. This vibration can be sensed as a varying voltage that is 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), which has also been integrated into low cost IMUs (Inertial Measurement Units).

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 material, which claims a better temperature stability and electro-mechanical coupling than other PE materials (Murata, 1999). The bar is excited into motion from two sides, and the third side is used to drive a feedback loop that controls the bar oscillation. The drive ceramics are also used to detect the motion of the bar. Only the difference of the two output signals is measured, thereby removing variation caused by multiple measurement circuits. A diagram of the bar is shown in Figure 4 (Hayashi, 1996).

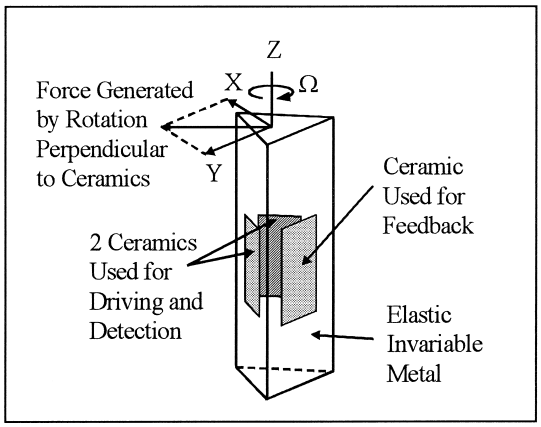

Figure 4. Gyrostar Free-Free-Bar and Ceramics.

https://doi.org/10.1017/S0373463301001357 Published online by Cambridge University Press

The gyro has instability in both the zero-rotation output voltage and the scale factor, mainly due to temperature variations. The gyro has also demonstrated limited endurance to shock and some susceptibility to vibration (Geier, 1998). Vibration in vehicles varies greatly in both magnitude and frequency, depending on engine speed, temperature, and numerous other factors. The specifications for the zero-rotation bias and scale factor vary by 20 %, so independent calibration of each unit is required. High temperature variations that occur in automotive applications necessitate regular calibration; however, it is very difficult 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 was 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. This scale factor appears to be very stable as a function of temperature for the two units that were tested, as the data sets were all processed using the same scale factor, although the temperature varied by as much as 20 °C between tests.

The zero offset was continually calibrated through a conditional LPF. The filter only operates when the vehicle appears to be driving relatively straight. 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 motorway 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, and a low gain is used in 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^{-3}, while a corresponding value of 10^{-6} was used during straight portions, 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.

The gyro also suffers from a misalignment with respect to the vertical plane containing the gravity vector. It is therefore not sensing the entire rotation in heading in the horizontal plane, but also sensing rotations in pitch and roll. This situation is shown in Figure 5 and the relationship between the measured rotation and the desired heading change can be given by (St. Lawrence, 1993)

$$\omega_G = \omega_\psi \cos\xi + \omega_\varphi \cos\beta \sin\xi + \omega_\theta \sin\beta \sin\xi. \tag{6}$$

Rearranging for the heading rate, one obtains:

$$\omega_\Psi = \frac{\omega_G}{\cos\xi} - \frac{\cos\xi + \omega_\varphi \cos\beta \sin\xi + \omega_\theta \sin\beta \sin\xi}{\cos\xi}, \tag{7}$$

https://doi.org/10.1017/S0373463301001357 Published online by Cambridge University Press

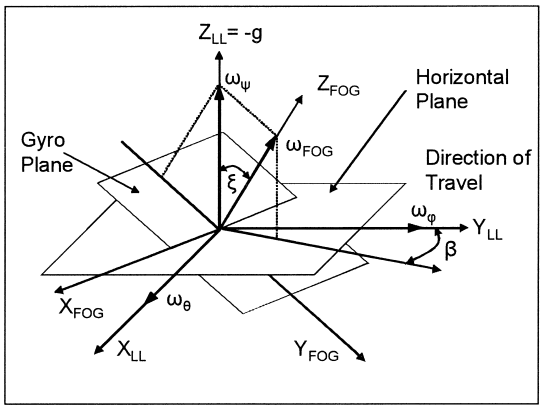

Figure 5. Geometry of a tilted gyro.

where:

ω_G is the measured rotation rate,
$\omega_{\dot{\psi}}$ is the desired azimuth rotation rate,
$\omega_{\dot{\theta}}$ is the pitch rotation rate,
$\omega_{\dot{\varphi}}$ is the roll rotation rate,
β is the horizontal component of the gyro and local level (LL) plane misalignment,
ξ is the misalignment between the gyro and LL vertical vectors.

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. Errors in the heading due to roll should result in a random walk behaviour (Harvey, 1998). These errors are normally considered an acceptable loss in vehicle navigation applications, since adding complementary gyros would be costly and result in only a marginal improvement. Ignoring these rotations drops the second term from the previous equation. The remaining error can be derived from the direction cosines of the gyro axes on to the local level axes with pitch angle (p) and roll angle (r) *as* (*ibid.*)

$$\sin^2 p + \sin^2 r + \cos^2 \xi = 1, \tag{8}$$

$$\beta = \tan^{-1}\left(\frac{p}{r}\right), \tag{9}$$

$$\cos \xi = \sqrt{1 - \sin p^2 - \sin r^2}, \tag{10}$$

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

$$\omega_\Psi = \frac{\omega_G}{\sqrt{1-\vartheta^2}}, \tag{11}$$

with the tilt ϑ defined as:

$$\vartheta^2 = p^2 + r^2. \tag{12}$$

This remaining correction relies on the observation of the vehicle's pitch and roll, which could be acquired using an inclinometer. Vehicle tilt has the effect of decreasing

https://doi.org/10.1017/S0373463301001357 Published online by Cambridge University Press

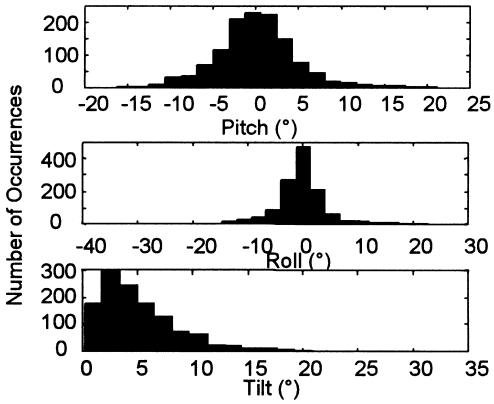

Figure 6. Vehicle motion under normal driving conditions.

the sensed rate of rotation during a turn by a small amount. For a measured rotation rate of 20°/s, which is seldom exceeded, with a vehicle tilt of 20°, the actual heading rate should be 21·3°/s, which results in a heading error of 1·3°/s, i.e. up to a 6° error for a 90° turn. At a tilt of 10°, this error drops to 1·3° for a 90° turn. It is difficult to equate heading errors to positional errors since they vary greatly depending on vehicle dynamics, but it can be easily seen that small heading errors can quickly lead to large position errors. If tilt compensation is not used, the scale factor for gyro output should be increased slightly to compensate for a RMS tilt value. An inclinometer would be necessary if positions were required inside covered or underground car parks. A sample distribution of vehicle motion experienced in various residential and motorway driving environments is shown in Figure 6. It can be seen that vehicle tilt seldom exceeds 10° under normal driving conditions.

2.3. *GPS.* NovAtel MiLLennium receivers were used for both the reference station and the vehicle. The remote receiver was used in a wide-correlator mode to simulate low-cost GPS sensor performance. The differential mode was used to reduce satellite and atmospheric propagation errors more effectively. The GPS solutions were obtained using the University of Calgary's C^{3}NAVG^2 single-difference carrier-smoothed code processing software. This software creates independent epoch solutions and does not filter the positions or velocities. Filtering at this point would complicate the use of the dynamics model in the external DR navigation filter. A 3-D position computation is initially performed. A 2-D solution, with the height component obtained from a previous position or from a digital map, is computed if there are not enough satellites or if the geometry is poor. A position will be output only if the DOP is better than five. The software also estimates the receiver velocity using Doppler measurements.

Computed latitude and longitude are used as observations for the DR integration filter, and GPS North and East velocity components are used to generate speed and heading observations for the filter. Only the horizontal component of velocity was used, since the vertical component tends to be small but noisy and can bias the results. Even on a 7% gradient, the error caused in speed by ignoring the vertical velocity component would amount to less than 0·25% of the ground speed. GPS speed was calculated as:

$$s = \sqrt{v_n^2 + v_e^2}, \tag{13}$$

https://doi.org/10.1017/S0373463301001357 Published online by Cambridge University Press

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

$$\Psi_g = \tan^{-1}\left(\frac{v_n}{v_e}\right). \tag{14}$$

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

$$\sigma_{sg}^2 = \sigma_{vn}^2 \frac{v_n^2}{s^2} + \sigma_{ve}^2 \frac{v_e^3}{s^2}, \tag{15}$$

and the GPS heading observation variance is:

$$\sigma_{\Psi g}^2 = \sigma_{vn}^2 \frac{1}{v_e^2\left(1+\frac{v_n^2}{v_e^2}\right)^2} + \sigma_{ve}^2 \frac{v_n^2}{s^4}. \tag{16}$$

3. A/D CONVERTER. The analogue data acquisition was carried out using the National Instruments DAQ-Pad MIO 16XE-50. Some specifications of the device are presented in Table 5.

Table 5. 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
Max Sampling Rate	20 kS/s
Relative Accuracy	$\pm$1 LSB max
Settling Time	50 μs max to $\pm$1 LSB
Analog Outputs	2, 12 bit
Counter/Timers	2 up/down, 24 bit
Digital I/O	8 TTL/CMOS input/output
Power Supply	9-42 VDC, 325 mA at 12 V

Differential signal sampling was used to minimize interference with normal ABS operation. All leads and wires were formed into twisted pairs where 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 ends of each differential input to minimize noise. Only the rate gyro was powered from the A/D converter 5 V supply. All other equipment was powered directly from the independent 12 V battery.

4. DATA COLLECTION SCHEME. The data acquisition program performed a scan of the four channels used at two kHz, with individual samples allowing sufficient 1 LSB settling time at 16 kHz. The acquisition was driven by the internal 20 MHz oscillator, which was also used to trigger a pulse generator at 1 s intervals. This pulse was routed to the mark input pin on the GPS 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, allowing for simple integration of the data. Two laptop computers were used to

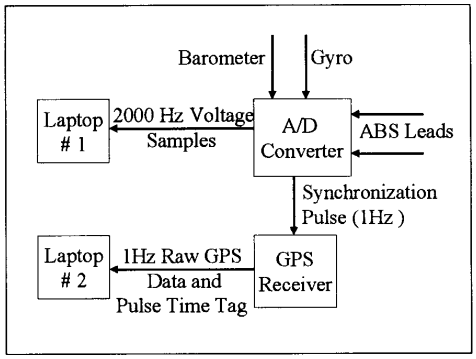

Figure 7. Data collection system.

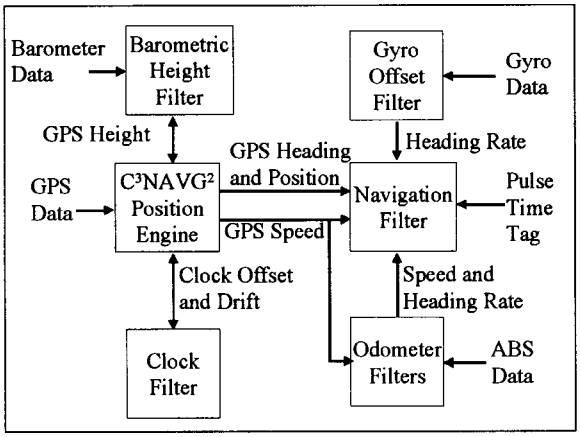

Figure 8. Data processing system.

record the data, one for GPS data, and the other, for the analogue data. A diagram of the data collection configuration is shown in Figure 7.

5. DATA PROCESSING SCHEME. The processing was done off-line to avoid the issue of broadcasting differential GPS corrections. No smoothing algorithms were used to ensure that the results would be identical to those obtained in real-time. The data processing of a 20-minute data set took approximately one minute on a 200 MHz P-II PC, so real-time implementation with a few modifications could have been easily accomplished. A flow chart depicting the data flow to the filter is presented in Figure 8. The C^3NAVG^2 software provides the GPS solutions that are integrated with dead reckoning in the navigation filter. The DR sensors are calibrated with filters external to the main navigation filter.

Attempts were made to include the calibration of the DR sensors into the navigation filter, but poor results were obtained. High correlations between the sensor bias states and the navigation parameters occurred, leading to misbehaviour of the filter after long periods of GPS outage. A primary example of the consequence of this is the occurrence of large position jumps during heading changes.

https://doi.org/10.1017/S0373463301001357 Published online by Cambridge University Press

6. NAVIGATION FILTER. The navigation filter employs an EKF (Extended Kalman Filter). The EKF uses the values of the filter states to predict the future states through a dynamics model (prediction), and the observations can be used, as they become available (update). The discrete time filter equations are given by (Gelb, 1974). Given a system model with no deterministic input:

$$\dot{x} = f(x,t) + u(t), \tag{17}$$

$$F = \frac{\partial f}{dx}, \tag{18}$$

$$\Phi_k = L^{-1}\{[sI - F]^{-1}\}_{t=\Delta t}, \tag{19}$$

and a linear measurement model:

$$z_k = H_k x_k + v_k, \tag{20}$$

supplied with initial conditions P_0 and X_0, the prediction equations are given by:

$$x_{k+1}^- = \Phi_k x_k^+, \tag{21}$$

$$P_{k+1}^- = \Phi_k P_k^+ \Phi_k^T + Q_k, \tag{22}$$

$$Q_k = \int_{t_k}^{t_{k+1}} \Phi_k(\xi) w w^T \Phi(\xi)^T \, d\xi, \tag{23}$$

and the update equations by:

$$x_k^+ = x_k^- + K_k[z_k - H_k x_k^-], \tag{24}$$

$$P_k^+ = P_k^- - K_k H_k P_k^-, \tag{25}$$

$$K_k = P_k^- H_k^T [H_k P_k^- H_k^T + R_k]^{-1}, \tag{26}$$

where:

x is the vector of unknown parameters,
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,
w is the vector of spectral densities,
Φ is the state transition matrix.

An EKF is necessary because the dynamics model is not linear. Since the degree of non-linearity in navigation applications can be high, the linearized EKF can easily become unstable. Great care must be taken in choosing the necessary statistics to prevent divergence of the filter while maintaining adequate filter performance.

Research has shown that the constant velocity model is adequate to handle most land vehicle dynamics (Lan, 1996); however, some implementations have used constant acceleration models or complex models designed for handling and traction control simulations (Julier and Durant-Whyte, 1995). The trade-off must be made between a higher order dynamics model, meaning more variables to estimate, and the need for the errors to be Gaussian. 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 non-linearity is high, and heading information becomes useless at low speeds. A constant heading

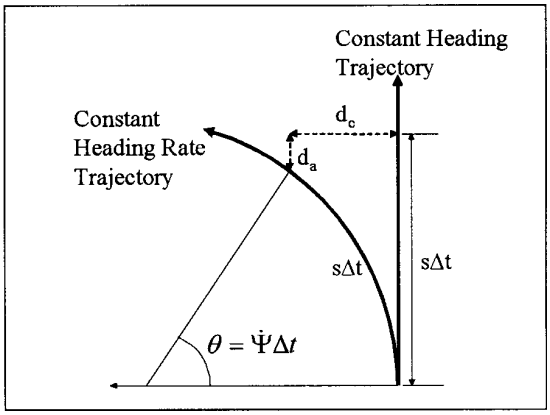

Figure 9. State prediction error.

rate and speed model was used to describe vehicle dynamics. The states for the model were:

$$x = \begin{bmatrix} \textit{North Position } (m) \\ \textit{East Position } (m) \\ \textit{Heading Rate } (rad/s) \\ \textit{Heading } (rad) \\ \textit{Speed } (m/s) \end{bmatrix} = \begin{bmatrix} \varphi \\ \lambda \\ \dot{\Psi} \\ \Psi \\ s \end{bmatrix}. \tag{27}$$

The latitude and longitude were transformed to units of length (metres) to make the propagation model simple. The headings were defined with origin East and counter-clockwise positive. The state transition matrix must be derived from the dynamics model, which is given by:

$$\varphi_{k+1} = \varphi_k + s_k \sin \Psi_k \Delta t, \tag{28}$$

$$\lambda_{k+1} = \lambda_k + s_k \cos \Psi_k \Delta t, \tag{29}$$

$$\Phi_{k,k+1} = \begin{bmatrix} 1 & 0 & 0 & a\Delta t & b\Delta t \\ 0 & 1 & 0 & c\Delta t & d\Delta t \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & \Delta t & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \tag{30}$$

$$a = s \cos \Psi_k, \tag{31}$$

$$b = \sin \Psi_k, \tag{32}$$

$$c = -s \sin \Psi_k, \tag{33}$$

$$d = \cos (\Psi_k). \tag{34}$$

This state transition matrix is used for propagating the covariance, but the dynamics model from Equations 28 and 29 is used directly to propagate the state vector, since the dynamics model is nonlinear. 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, depicted in Figure 9, was found to be insignificant. In Figure 9, the radius of the circle can be determined from:

$$r = \frac{s}{\dot{\Psi}}. \tag{35}$$

https://doi.org/10.1017/S0373463301001357 Published online by Cambridge University Press

The along track and across track errors over an interval Δt are:

$$d_c = r(1-\cos\theta), \tag{36}$$

$$d_a = s\Delta t - r\sin\theta. \tag{37}$$

Since dead reckoning data is processed at 10 Hz, the propagation time intervals are very small. Taking a sharp corner at 5 m/s and 20°/s of heading rate approximates the worst error normally caused, which is only 8·7 mm across track and 0·4 mm along track for each 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.

The Q matrix can be computed from (Gelb, 1974):

$$Q_k = E\{w_k w_k^T\} \tag{38}$$

$$Q_k = \int_{t_k}^{t_{k+1}} \int_{t_k}^{t_{k+1}} \Phi(t_{k+1}, \xi)\, G(\xi)\, E\{u(\xi)\, u^T(\eta)\} \; G^T(\eta)\Phi^T(t_{k+1}, \eta)\, d\xi\, d\eta \tag{39}$$

with:

$$E\{u(\xi)u^T(\eta)\} = \delta(\xi-\eta). \tag{40}$$

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

$$Q_k = \int_0^{\Delta t} \Phi(t_{k+1}, \xi) q_k \Phi^T(t_{k+1}, \eta) d\xi, \tag{41}$$

where: q_k is a diagonal matrix of spectral densities. Evaluating with the a, b, c and d parameters given above, we obtain:

$$Q(1,1) = q(1,1)\Delta t + \frac{a^2}{3} q(4,4)\Delta t^3 + \frac{b^2}{3} q(5,5)\Delta t^3, \tag{42}$$

$$Q(2,1) = Q(1,2) = \frac{ac}{3} q(4,4)\Delta t^2 + \frac{bd}{3} q(5,5)\Delta t^2, \tag{43}$$

$$Q(4,1) = Q(1,4) = \frac{a}{2} q(4,4)\Delta t^2, \tag{44}$$

$$Q(5,1) = Q(1,5) = \frac{b}{2} q(5,5)\Delta t^2, \tag{45}$$

$$Q(2,2) = q(2,2)\Delta t + \frac{c^2}{3} q(4,4)\Delta t^3 + \frac{d^2}{3} q(5,5)\Delta t^3, \tag{46}$$

$$Q(4,2) = Q(2,4) = \frac{c}{2} q(4,4)\Delta t^2, \tag{47}$$

$$Q(5,2) = Q(2,5) = \frac{d}{2} q(5,5)\Delta t^2, \tag{48}$$

$$Q(3,3) = q(4,4)\Delta t + \frac{1}{3} q(3,3)\Delta t^3, \tag{49}$$

$$Q(4,4) = q(5,5)\Delta t. \tag{50}$$

https://doi.org/10.1017/S0373463301001357 Published online by Cambridge University Press

Not all of the five possible spectral densities, $q(n,n)$, were employed. The driving noise values for the dynamics corresponded 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:

$$q(3,3) = (5°/\text{s})^2\text{s (heading rate) and}$$
$$q(5,5) = (2\ \text{m/s})^2\text{s (speed).}$$

The dynamics model was found to work well while moving, but heading and position would occasionally drift while the vehicle was stationary. To fix this anomaly, the prediction phase was not performed if neither of the wheels detected motion during the last three measurements.

To minimize jitter about the trajectory, yet maintain the ability of the filter to respond to sharp turns, it was also decided that two different process noise values would be used for the heading rate, depending on whether a turn was in progress or not. This procedure has been called smart stochastic modelling 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 $(20°/\text{s})^2\text{s}$ when a turn was detected.

Because of the non-linear dynamics model, strong correlations tended to develop between some of the parameters. Several times these covariances exceeded the diagonals in the covariance matrix, and matrix inversions would fail. This was found to occur after periods of stationary vehicle operation or long travel times in a cardinal direction. The problem was also controlled with the addition of the x adaptive features described above.

Blunder GPS headings and positions were found to be a major source of error in the filtered solutions, so a statistical test was implemented to detect them before corrupting the filter. With the residual given by:

$$r = z - Hx, \tag{51}$$

the residual covariance can be calculated as:

$$P_r = HP_xH^T + R. \tag{52}$$

Each residual can then be tested against a pre-determined threshold, such as:

$$\frac{r_i}{P_{rii}} < 4{\cdot}57. \tag{53}$$

A decision process was also implemented whereby the first two consecutive blunder observations would not be included, but the next observation following two blunders was used, regardless of whether it passed the test. This allowed a biased trajectory to be snapped back after a GPS outage, but lowered the chance of a blunder entering the solution, since multiple consecutive blunders would have to occur for this to happen.

7. SYSTEM TESTING. A test route through Calgary's downtown core was devised that would create long-duration GPS signal outages and vehicle dynamics variations. Previous implementations of the filter were tested on open roads with simulated signal masking (Stephen & Lachapelle, 2000). These tests could simulate

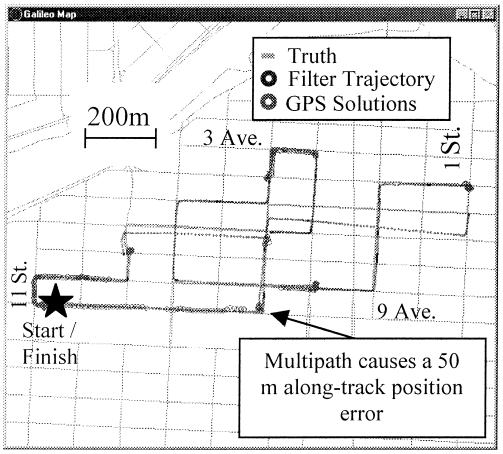

Figure 10. April 4, Run 1 – Gyro and odometry with wide-correlator GPS.

typical signal masking problems adequately, but not the high multi-path that also plagues GPS users in urban canyons. A 6-km test route running through Calgary's downtown core and shown in Figure 10 was selected. High-rise buildings of up of 50 storeys line the route selected. A good measure of the signal masking occurring along the route is the availability of GPS stand-alone solutions, shown in Figure 10; availability is less than 20%.

The test route was driven three times on two separate occasions. Each run took about 20 minutes, and the three runs were performed nearly consecutively. Some correlation of the GPS multi-path between successive runs is evident, due to the similar geometry between tests. There is little correlation between the GPS data on separate dates since the tests were taken three hours later on the second test. Results are presented from two modes of DR augmentation: both use ABS odometry for speed estimation, but heading information was drawn from ABS differential odometry or the gyro. The availability of the augmented solutions is nearly 100%.

During the first test (April 4, 2000), the lead to one of the wheel sensor inputs was not secure and large voltage spikes were sensed. As a result, only results using one wheel for speed and gyro for heading are presented for the first test. A second test was also carried out for this case in which a second GPS receiver was run in Narrow Correlator mode for comparison purpose. Results using both the narrow and wide correlator receivers are found in Figure 10 and Figure 11, respectively. The true trajectory refers to the street centreline. The GPS solutions created by C^3NAVG^2 are shown as large circles. The filtered combination of dead reckoning sensors and GPS solutions is shown as small circles.

The route was designed so that some calibration data would be available at the beginning of each test in order to establish an appropriate initial heading. This portion begins on the lower left at the corner of 9th Avenue and 11th Street heading East. There are few buildings to the South on this road, and being at higher latitude (51°), the number of satellites visible to the South tends to be high. Very few GPS solutions were obtained elsewhere, except at a few intersections, due to extreme masking.

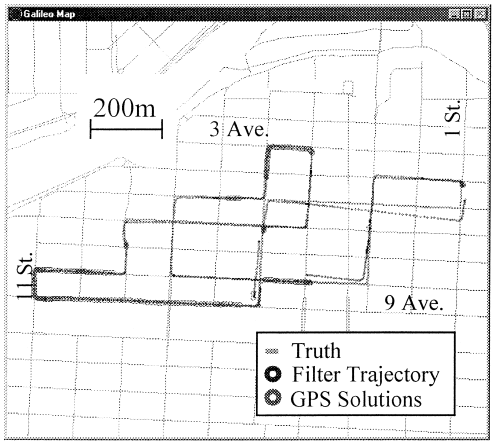

Figure 11. April 4, Run 1 – Gyro and odometry with narrow-correlator GPS.

Table 6. April 4 Tests – Stand-alone GPS position statistics.

Correlator spacing	Max error (m)	RMS error (m)
Wide	28	7
Narrow	24	5

Table 7. April 4 Tests – filter position errors.

Correlator spacing	Max error (m)	RMS error (m)
Wide	81	13
Narrow	44	9

Since there were not enough GPS measurements to create an accurate carrier-phase reference trajectory, across-track errors were computed as the distance to the centre-line derived from the digital map. The limitation of this method is that the vehicle often did not drive close to the centre-line, especially during turns. In addition, the shortest computed distance is not always the distance to the road on which the vehicle was driving, e.g. a biased trajectory will show zero error every time it crosses another road on the map. Also, along-track errors are not recognized. An example of this is a large multi-path error, which causes a position error along the direction of the road, as shown in Figure 10. The bias is not apparent in the distance errors until the vehicle makes a right hand turn about three blocks later.

Table 6 summarizes the quality of the GPS solutions for both the wide and narrow correlator spacing. The narrow correlator spacing reduces multi-path error, but these are still significant at a few locations.

Table 7 summarizes the maximum and RMS errors for the filtered data, i.e. GPS solutions augmented with gyro and odometry dead reckoning. The filtered data errors

https://doi.org/10.1017/S0373463301001357 Published online by Cambridge University Press

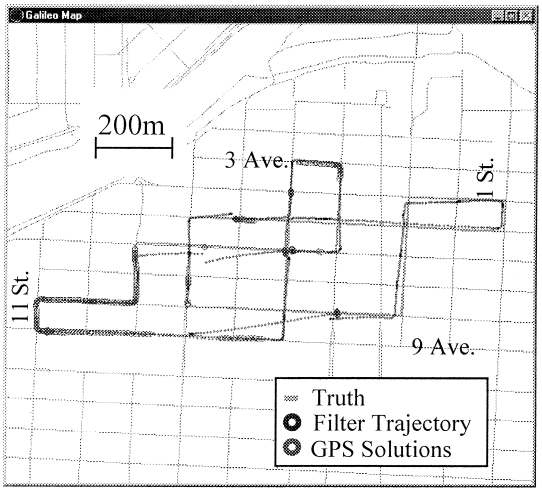

Figure 12. April 21, Run 1 – Differential ABS DR.

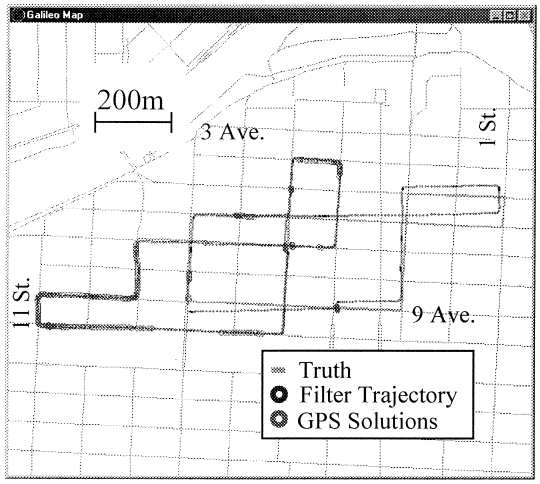

Figure 13. April 21, Run 1 – Gyro 1 and odometry DR.

are related closely to the GPS position and velocity errors. The figures show very little evidence of un-modelled gyro drift (such as a slow constant change of heading) or scale factor errors (consistently recording too much or too little heading change on corners) while coasting in periods of GPS masking. The primary causes of errors are multi-path induced position errors, and poor GPS headings due to low speeds, multi-path, and noisy Doppler measurements.

In addition, the assumption that multiple consecutive blunders would rarely occur proved to be optimistic in such urban canyons. Also, after long GPS outages, the heading and position variances increase, allowing larger errors to enter the solution without being detected as blunders. For a blunder detection algorithm to perform well under such conditions, it would need to determine whether the present trajectory was in error, or whether the GPS measurements contain blunders. Using a digital

https://doi.org/10.1017/S0373463301001357 Published online by Cambridge University Press

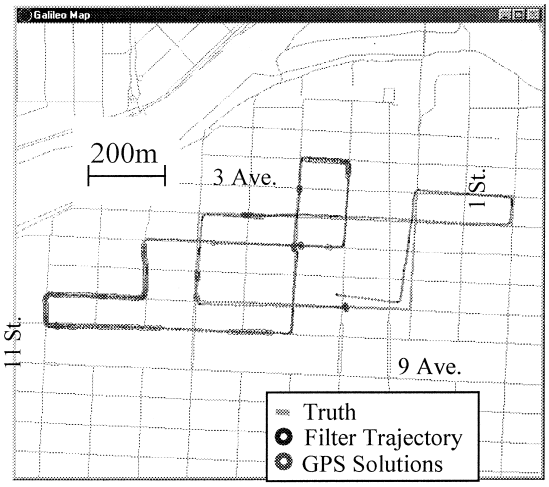

Figure 14. April 21, Run 1 – Gyro 2 and odometry DR.

map, the distinction can usually be made as to whether the trajectory, GPS positions, or both are in error, but it is difficult to impart these decision-making criteria to a map-aiding algorithm.

The satellite visibility was higher during the second set of tests (April 21, 2000), resulting in more GPS position fixes, mainly at intersections. In addition, a second gyro unit was tested at the same time to validate the results obtained with the primary gyro. Results from the first test run are shown in Figures 12 to 14. Again, the gyro outperforms ABS since the zero-rotation bias can be calibrated every time the vehicle is forced to stop, while error due to wheel-scale factor estimation and wheel slippage are apparent in the ABS results. The across-track errors are summarized in Table 8

Table 8. April 21 Tests – filter position errors.

Dead Reckoning	Max (m)	RMS (m)
Differential ABS	115	18
Gyro 1 & Odometry	67	14
Gyro 2 & Odometry	69	14

for the combination of all three test runs. The maximum stand-alone GPS position error was 47 m, while the RMS error was 7 m. This position error and the GPS heading error accounted for a large portion of the integrated filter position errors.

The difference in observation variance for the gyro and differential ABS heading rates results in larger errors in the differential ABS heading. The difficulty in estimating the wheel-scale factors and frequent occurrence of wheel slippage require that a higher variance be used on the ABS heading rate than the gyro heading rate measurements. As a result, GPS headings receive more relative weight, and large GPS heading errors have a greater effect on the filtered heading. The second gyro unit appears to have been better calibrated for the zero-offset voltage than the first. This

is indicated by the constant heading drift to the left that is evident in the right hand side of Figure 13.

8. CONCLUSIONS. The GPS-augmented multi-sensor navigation method presented in this paper is capable of near-continuous coverage in urban canyons of the type tested, with little added-cost to the production of a vehicle. ABS, which is now standard on most personal vehicles, used in conjunction with a low-cost rate gyro and a GPS receiver, are sufficient to maintain a position accuracy of the order of 20 m or better most of the time, although larger errors can occur due to the presence of multi-path.

When ABS is compared to the low-cost rate gyro, the latter is found to outperform ABS for several reasons. Firstly, the gyro drift can be calibrated very quickly if the vehicle makes a stop. Wheel slippage plagues odometry 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 tyres are variable with speed, temperature, and other factors, making calibration difficult. 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 city cores where DR is most needed. Because of the poorer performance of differential odometry, no significant improvement was found by using it in conjunction with the gyro for heading information.

Performance could probably be further improved by adding map information. Headings could be constrained by road direction, and biases could be removed by shifting the position to the most likely road, where appropriate. GPS headings are a major error source in many of the tests. In addition, 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 continual calibration. Finally, the use of emerging low-cost MEMS inertial measuring units (IMUs) could decrease this error source while also assisting in the detection of GPS multi-path.

ACKNOWLEDGEMENTS

The authors would like to thank NSERC for financial support. Mr. Jeff Geier, Motorola, provided the Murata gyros. Also, we would also like to thank many of graduate students of the Department of Geomatics Engineering who have assisted with this research project.

REFERENCES

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.

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.

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.

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 Engineerin*. 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.

Julier, S. and Durrant-Whyte, H. (1995). Navigation and parameter estimation of high speed road vehicles. *Proceedings of the 1995 IEEE Conference on Robotics and Automation*, Pp.101–105.

Kealy, A., Tsakiri, M. and Stewart, M. (1999). Land navigation in the urban canyon – a Kalman filter solution using integrated GPS, GLONASS and Dead Reckoning. *Proceedings ION GPS-99*.

Lan, H. (1996). Development of a real-time kinematic GPS System. *M.Sc. Thesis*, *UCGE Report 20107*, *Department of Geomatics Engineering*. University of Calgary.

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

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

St. Lawrence, W. (1993). Understanding tilt vectors and processing tiltmeter data. *Applied Geomechanics*, Santa Cruz.

Stephen, J. (2000). Development of a GNSS-based multi-sensor vehicle navigation system. *M.Sc. Thesis UCGE Report 20140*, *Department of Geomatics Engineering*. University of Calgary. [http://www.ensu.ucalgary.ca/Papers/Thesis/MGS/00.20140.JStephen.pdf].

Stephen, J., and Lachapelle, G. (2000). Development of a GNSS-based multi-sensor vehicular navigation system. *Proceedings of National Technical Meeting*, The Institute of Navigation (Anaheim, CA, Jan 20–22), 268–278.