



An introduction to inertial navigation

Oliver J. Woodman

August 2007

© 2007 Oliver J. Woodman

Technical reports published by the University of Cambridge
Computer Laboratory are freely available via the Internet:

<http://www.cl.cam.ac.uk/techreports/>

ISSN 1476-2986

An introduction to inertial navigation

Oliver J. Woodman

Abstract

Until recently the weight and size of inertial sensors has prohibited their use in domains such as human motion capture. Recent improvements in the performance of small and lightweight micro-machined electromechanical systems (MEMS) inertial sensors have made the application of inertial techniques to such problems possible. This has resulted in an increased interest in the topic of inertial navigation, however current introductions to the subject fail to sufficiently describe the error characteristics of inertial systems.

We introduce inertial navigation, focusing on strapdown systems based on MEMS devices. A combination of measurement and simulation is used to explore the error characteristics of such systems. For a simple inertial navigation system (INS) based on the Xsens Mtx inertial measurement unit (IMU), we show that the average error in position grows to over 150 m after 60 seconds of operation. The propagation of orientation errors caused by noise perturbing gyroscope signals is identified as the critical cause of such drift. By simulation we examine the significance of individual noise processes perturbing the gyroscope signals, identifying white noise as the process which contributes most to the overall drift of the system.

Sensor fusion and domain specific constraints can be used to reduce drift in INSs. For an example INS we show that sensor fusion using magnetometers can reduce the average error in position obtained by the system after 60 seconds from over 150 m to around 5 m. We conclude that whilst MEMS IMU technology is rapidly improving, it is not yet possible to build a MEMS based INS which gives sub-meter position accuracy for more than one minute of operation.

1 Introduction

This report provides an introduction to inertial navigation and the error characteristics of inertial navigation systems. Its purpose is to address the lack of a readable introduction into the subject which does not oversimplify or ignore the error properties of inertial navigation systems. The report also aims to explain the meaning and significance of performance specification measurements such as ‘noise density’ and ‘bias stability’ which are often stated by manufacturers.

The reader should note that whilst this report aims to provide a broad introduction to the subject of inertial navigation, the latter chapters focus mainly on strapdown type inertial navigation systems using micro-machined electromechanical systems (MEMS) devices. MEMS technology is of particular interest at the current time since it offers rugged, low cost, small and lightweight inertial sensors relative to the other available technologies. The performance of MEMS inertial devices is also improving rapidly. Throughout the report a simple inertial navigation system (INS) is developed based on an Xsens Mtx device. The report is structured as follows:

- Section 2 introduces the reader to inertial navigation, its uses, and the two main varieties of inertial navigation system.
- Sections 3 and 4 describe gyroscopes and accelerometers in detail. Both sections contain an overview of the different types of sensors available, as well a description of error sources.
- Section 5 introduces Allan Variance, a technique which can be used to detect and measure the noise characteristics of gyroscope and accelerometer signals.

- Section 6 describes strapdown inertial navigation in more detail and explains how errors in individual gyroscopes and accelerometers propagate through the navigation system as a whole. The performance of a simple INS is analysed in order to illustrate the relative importance of noise perturbing the gyroscope and accelerometer signals.
- Section 7 describes how simulation can be used to analyse the relative importance of different noise sources. A simple simulator is constructed and verified against the real system developed in Section 6.
- Section 8 introduces several methods for reducing drift in inertial systems.




Figure 1: The body and global frames of reference.

2 Inertial Navigation

Inertial navigation is a self-contained navigation technique in which measurements provided by accelerometers and gyroscopes are used to track the position and orientation of an object relative to a known starting point, orientation and velocity. Inertial measurement units (IMUs) typically contain three orthogonal rate-gyroscopes and three orthogonal accelerometers, measuring angular velocity and linear acceleration respectively. By processing signals from these devices it is possible to track the position and orientation of a device, as described in Section 2.1.

Inertial navigation is used in a wide range of applications including the navigation of aircraft, tactical and strategic missiles, spacecraft, submarines and ships. Recent advances in the construction of MEMS devices have made it possible to manufacture small and light inertial navigation systems. These advances have widened the range of possible applications to include areas such as human and animal motion capture.

2.1 Inertial System Configurations

Nearly all IMUs fall into one of the two categories outlined below. The difference between the two categories is the frame of reference in which the rate-gyroscopes and accelerometers operate. Throughout this report we will refer to the navigation system's frame of reference as the body frame and to the frame of reference in which we are navigating as the global frame, as shown in Figure 1.

2.1.1 Stable Platform Systems

In stable platform type systems the inertial sensors are mounted on a platform which is isolated from any external rotational motion. In other words the platform is held in alignment with the global frame. This is achieved by mounting the platform using gimbals (frames) which allow the platform freedom in all three axes, as shown in Figure 2. The platform mounted gyroscopes detect any platform rotations. These signals are fed back to torque motors which rotate the gimbals in order to cancel out such rotations, hence keeping the platform aligned with the global frame.

To track the orientation of the device the angles between adjacent gimbals can be read using angle pick-offs. To calculate the position of the device the signals from the platform mounted accelerometers are double integrated. Note that it is necessary to subtract acceleration due to gravity from the vertical channel before performing the integration. The stable platform inertial navigation algorithm is shown in Figure 3.




Figure 2: A stable platform IMU.




Figure 3: Stable platform inertial navigation algorithm.

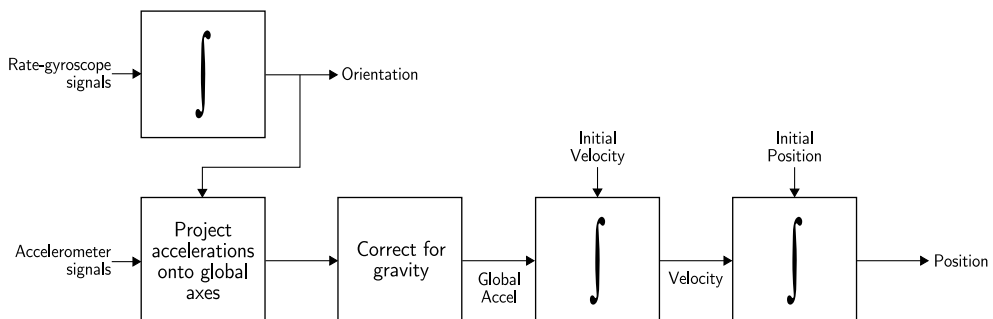


Figure 4: Strapdown inertial navigation algorithm.

2.1.2 Strapdown Systems

In strapdown systems the inertial sensors are mounted rigidly onto the device, and therefore output quantities measured in the body frame rather than the global frame. To keep track of orientation the signals from the rate gyroscopes are ‘integrated’, as described in Section 6. To track position the three accelerometer signals are resolved into global coordinates using the known orientation, as determined by the integration of the gyro signals. The global acceleration signals are then integrated as in the stable platform algorithm. This procedure is shown in Figure 4.

Stable platform and strapdown systems are both based on the same underlying principles. Strapdown systems have reduced mechanical complexity and tend to be physically smaller than stable platform systems. These benefits are achieved at the cost of increased computational complexity. As the cost of computation has decreased strapdown systems have become the dominant type of INS.




Figure 5: A conventional mechanical gyroscope (source: [1]).

3 Gyroscopes

3.1 Types of Gyroscope

In this section the main types of gyroscope are presented. Note that this is far from an exhaustive list. In particular there are many different varieties of mechanical gyroscope which are not described. A more comprehensive survey can be found in [1].

3.1.1 Mechanical

A conventional gyroscope consists of a spinning wheel mounted on two gimbals which allow it to rotate in all three axes, as shown in Figure 5. An effect of the conservation of angular momentum is that the spinning wheel will resist changes in orientation. Hence when a mechanical gyroscope is subjected to a rotation the wheel will remain at a constant global orientation and the angles between adjacent gimbals will change. To measure the orientation of the device the angles between adjacent gimbals can be read using angle pick-offs. Note that a conventional gyroscope measures orientation. In contrast nearly all modern gyroscopes (including the optical and MEMS types outlined in Sections 3.1.2 and 3.1.3) are rate-gyros, which measure angular velocity.

The main disadvantage of mechanical gyroscopes is that they contain moving parts. Moving parts cause friction, which in turn causes the output to drift over time. To minimise friction high-precision bearings and special lubricants are used, adding to the cost of the device. Mechanical gyroscopes also require a few minutes to warm up, which is not ideal in many situations.

3.1.2 Optical

A fibre optic gyroscope (FOG) uses the interference of light to measure angular velocity. A FOG consists of a large coil of optical fibre. To measure rotation two light beams are fired into the coil in opposite directions. If the sensor is undergoing a rotation then the beam travelling in the direction of rotation will experience a longer path to the other end of the fibre than the beam travelling against the rotation, as illustrated in Figure 6. This is known as the Sagnac effect. When the beams exit the fibre they are combined. The phase shift introduced due to the Sagnac effect causes the beams to interfere, resulting in a combined beam whose intensity depends on the angular velocity. It is therefore possible to measure the angular velocity by measuring the intensity of the combined beam.

Ring laser gyroscopes (RLGs) are also based on the Sagnac effect. The difference between a FOG and RLG is that in a RLG laser beams are directed around a closed path using mirrors rather than optical




Figure 6: The Sagnac effect. The dashed line is the path taken by the beam travelling in the direction of rotation. The solid line is the beam travelling against the rotation. θ is the angle through which the gyro turns whilst the beams are in flight.

fibre.

Unlike mechanical gyroscopes, optical gyros contain no moving parts and require only a few seconds to start-up. The accuracy of an optical gyro is largely dependent on the length of the light transmission path (larger is better), which is constrained by the size of the device.

3.1.3 MEMS Gyroscopes

Despite years of development, mechanical and optical gyroscopes still have high part counts and a requirement for parts with high-precision tolerances and intricate assembly techniques. As a result they remain expensive. In contrast MEMS sensors built using silicon micro-machining techniques have low part counts (a MEMS gyroscope can consist of as few as three parts) and are relatively cheap to manufacture.

MEMS gyroscopes make use of the Coriolis effect, which states that in a frame of reference rotating at angular velocity ω , a mass m moving with velocity v experiences a force:

$$F_c = -2m(\omega \times v) \quad (1)$$

MEMS gyroscopes contain vibrating elements to measure the Coriolis effect. Many vibrating element geometries exist, such as vibrating wheel and tuning fork gyroscopes. The simplest geometry consists of a single mass which is driven to vibrate along a drive axis, as shown in Figure 7. When the gyroscope is rotated a secondary vibration is induced along the perpendicular sense axis due to the Coriolis force. The angular velocity can be calculated by measuring this secondary rotation.

At present MEMS sensors cannot match the accuracy of optical devices, however they are expected to do so in the future. Below is a list of the advantageous properties of MEMS sensors, taken from [1].

- small size;
- low weight;
- rugged construction;
- low power consumption;
- short start-up time;
- inexpensive to produce (in high volume);




Figure 7: A vibrating mass gyroscope (source: [2]).

	GG1320AN (Laser Gyro)	GG5300 (MEMS 3xGyro)
Size	88 mm × 88 mm × 45 mm	50 mm × 50 mm × 30 mm
Weight	454 g	136 g
Start-Up Time	< 4 s	< 1 s
Power	15 Vdc, 1.6 watts nominal 5 Vdc, 0.375 watts nominal	5 Vdc, < 800 mA
Operating Temperature Range	−54 °C to 85 °C	−45 °C to 85 °C
Angular Random Walk	0.0035°/√h	0.2°/√h
Bias Stability	0.0035°/h	< 70°/h

Table 1: Specifications for the Honeywell GG1320AN and GG5300 gyroscopes.

- high reliability;
- low maintenance;
- compatible with operations in hostile environments;

As previously mentioned, the major disadvantage of MEMS gyroscopes is that they are currently far less accurate than optical devices. Table 1 illustrates the advantages and disadvantages of MEMS technology by comparing the specifications of two gyroscopes manufactured by Honeywell¹. The GG1320AN is a single axis digital laser gyro. The GG5300 is a three axis MEMS rate gyro. Note that the MEMS package not only contains three gyros as opposed to one, but also has a lower power consumption, a shorter start-up time, and is both smaller and lighter than the optical device. The main disadvantage of the MEMS device is that it is far less accurate, as indicated by the bias stability and angular random walk measurements. These measurement types are explained in Section 3.2.

3.2 MEMS Gyro Error Characteristics

In this section we examine the errors which arise in MEMS gyros, and their effect on the integrated (orientation) signal.

¹<http://www.honeywell.com>

3.2.1 Constant Bias

The bias of a rate gyro is the average output from the gyroscope when it is not undergoing any rotation (i.e: the offset of the output from the true value), in $^\circ/\text{h}$. A constant bias error of ϵ , when integrated, causes an angular error which grows linearly with time $\theta(t) = \epsilon \cdot t$.

The constant bias error of a rate gyro can be estimated by taking a long term average of the gyro's output whilst it is not undergoing any rotation. Once the bias is known it is trivial to compensate for it by simply subtracting the bias from the output.

3.2.2 Thermo-Mechanical White Noise / Angle Random Walk

The output of a MEMS gyro will be perturbed by some thermo-mechanical noise which fluctuates at a rate much greater than the sampling rate of the sensor. As a result the samples obtained from the sensor are perturbed by a white noise sequence, which is simply a sequence of zero-mean uncorrelated random variables. In this case each random variable is identically distributed and has a finite variance σ^2 .

To see what effect this noise has on the integrated signal we can do a simple analysis in which it is assumed that the rectangular rule is used to perform the integration. Let N_i be the i^{th} random variable in the white noise sequence. Each N_i is identically distributed with mean $E(N_i) = E(N) = 0$ and finite variance $\text{Var}(N_i) = \text{Var}(N) = \sigma^2$. By the definition of a white sequence $\text{Cov}(N_i, N_j) = 0$ for all $i \neq j$. The result of using the rectangular rule to integrate the white noise signal $\epsilon(t)$ over a timespan $t = n \cdot \delta t$ is

$$\int_0^t \epsilon(\tau) d\tau = \delta t \sum_{i=1}^n N_i \quad (2)$$

where n is the number of samples received from the device during the period and δt is the time between successive samples. Using the standard formulae $E(aX + bY) = aE(X) + bE(y)$ and $\text{Var}(aX + bY) = a^2\text{Var}(X) + b^2\text{Var}(Y) + 2ab\text{Cov}(X, Y)$ (where a and b are constants and X and Y are random variables) it follows that

$$E\left(\int_0^t \epsilon(\tau) d\tau\right) = \delta t \cdot n \cdot E(N) = 0 \quad (3)$$

$$\text{Var}\left(\int_0^t \epsilon(\tau) d\tau\right) = \delta t^2 \cdot n \cdot \text{Var}(N) = \delta t \cdot t \cdot \sigma^2. \quad (4)$$

Hence the noise introduces a zero-mean random walk² error into the integrated signal, whose standard deviation

$$\sigma_\theta(t) = \sigma \cdot \sqrt{\delta t \cdot t} \quad (5)$$

grows proportionally to the square root of time.

Since we are usually interested in how the noise affects the integrated signal it is common for manufacturers to specify noise using an angle random walk (ARW) measurement

$$\text{ARW} = \sigma_\theta(1) \quad (6)$$

with units $^\circ/\sqrt{\text{h}}$. For example the Honeywell GG5300 has an ARW measurement of $0.2^\circ/\sqrt{\text{h}}$. This means that after 1 hour the standard deviation of the orientation error will be 0.2° , after 2 hours it will be $\sqrt{2} \cdot 0.2 = 0.28^\circ$ and so on. Other measurements used to specify noise are power spectral density (units $(^\circ/\text{h})^2/\text{Hz}$) and FFT noise density (units $^\circ/\text{h}/\sqrt{\text{Hz}}$). It is possible to convert between the various different noise specifications using the equations:

$$\text{ARW } (^\circ/\sqrt{\text{h}}) = \frac{1}{60} \cdot \sqrt{\text{PSD } ((^\circ/\text{h})^2/\text{Hz})} \quad (7)$$

$$\text{ARW } (^\circ/\sqrt{\text{h}}) = \frac{1}{60} \cdot \text{FFT } (^\circ/\text{h}/\sqrt{\text{Hz}}) \quad (8)$$

For more information about angle random walk and noise specifications see [3].

²Here a random walk is defined as a process consisting of a series of steps, in which the direction and size of each step is randomly determined.

3.2.3 Flicker Noise / Bias Stability

The bias of a MEMS gyroscope wanders over time due to flicker noise in the electronics and in other components susceptible to random flickering. Flicker noise is noise with a $1/f$ spectrum, the effects of which are usually observed at low frequencies in electronic components. At high frequencies flicker noise tends to be overshadowed by white noise. Bias fluctuations which arise due to flicker noise are usually modelled as a random walk.

A bias stability measurement describes how the bias of a device may change over a specified period of time, typically around 100 seconds, in fixed conditions (usually including constant temperature). Bias stability is usually specified as a 1σ value with units $^{\circ}/h$, or $^{\circ}/s$ for less accurate devices. Under the random walk model bias stability can be interpreted as follows; If B_t is the known bias at time t , then a 1σ bias stability of $0.01^{\circ}/h$ over 100 seconds means that the bias at time $(t + 100)$ seconds is a random variable with expected value B_t and standard deviation $0.01^{\circ}/h$. Over time this property creates a random walk in the gyro bias, whose standard deviation grows proportionally to the square root of time. For this reason bias stability is occasionally specified by a bias random walk measurement

$$\text{BRW } (^{\circ}/\sqrt{h}) = \frac{\text{BS } (^{\circ}/h)}{\sqrt{t} \text{ (h)}} \quad (9)$$

where t is the timespan for which the bias stability is defined.

As usual we are interested in how this error affects the orientation obtained from integrating the rate-gyro signal. If we assume the bias random walk model, then the result of integrating the bias fluctuations is a second-order³ random walk in angle. In reality bias fluctuations do not really behave as a random walk. If they did then the uncertainty in the bias of a device would grow without bound as the timespan increased. In practice the bias is constrained to be within some range, and therefore the random walk model is only a good approximation to the true process for short periods of time.

3.2.4 Temperature Effects

Temperature fluctuations due to changes in the environment and sensor self heating induce movement in the bias. Note that such movements are not included in bias stability measurements which are taken under fixed conditions.

Any residual bias introduced due to a change in temperature will cause an error in orientation which grows linearly with time, as described in Section 3.2.1. The relationship between bias and temperature is often highly nonlinear for MEMs sensors. Most inertial measurement units (IMUs) contain internal temperature sensors which make it possible to correct for temperature induced bias effects. Some IMUs such as the Xsens⁴ Mtx perform such corrections internally.

3.2.5 Calibration Errors

The term ‘calibration errors’ refers collectively to errors in the scale factors, alignments, and linearities of the gyros. Such errors tend to produce bias errors that are only observed whilst the device is turning. Such errors lead to the accumulation of additional drift in the integrated signal, the magnitude of which is proportional to the rate and duration of the motions [4]. It is usually possible to measure and correct calibration errors. Some new IMUs such as the Xsens Mtx apply internal corrections for calibration errors.

3.2.6 Summary

The main sources of error outlined in this section are summarised in Table 2. The relative importance of each error source varies across different gyroscopes. For MEMS gyroscopes angle random walk (noise) errors and uncorrected bias errors either due to uncompensated temperature fluctuations or an error in the initial bias estimation are usually the most important sources of error. Angle random walk can be used as a lower bound for uncertainty in the orientation obtained from integrating a rate-gyroscope’s signal.

³The integral of a first-order random walk.

⁴<http://www.xsens.com>

Error Type	Description	Result of Integration
Bias	A constant bias ϵ	A steadily growing angular error $\theta(t) = \epsilon \cdot t$
White Noise	White noise with some standard deviation σ	An angle random walk, whose standard deviation $\sigma_\theta(t) = \sigma \cdot \sqrt{\delta t \cdot t}$ grows with the square root of time
Temperature Effects	Temperature dependent residual bias	Any residual bias is integrated into the orientation, causing an orientation error which grows linearly with time
Calibration	Deterministic errors in scale factors, alignments and gyro linearities	Orientation drift proportional to the rate and duration of motion
Bias Instability	Bias fluctuations, usually modelled as a bias random walk	A second-order random walk

Table 2: A Summary of Gyro Error Sources.




Figure 8: A mechanical accelerometer (source: [1]).

4 Linear Accelerometers

4.1 Types of Accelerometer

An accelerometer can be broadly classified as either a mechanical or solid state device. In this section these two types of accelerometer are described, as are MEMS accelerometers. For a complete survey the reader should refer to [1].

4.1.1 Mechanical

A mechanical accelerometer consists of a mass suspended by springs, as shown in Figure 8. The displacement of the mass is measured using a displacement pick-off, giving a signal that is proportional to the force F acting on the mass in the direction of the input axis. Newton's second law $F = ma$ is then used to calculate the acceleration acting on the device.

4.1.2 Solid State

Solid-state accelerometers can be broken into various sub-groups, including surface acoustic wave, vibratory, silicon and quartz devices. Solid state accelerometers are small, reliable and rugged.

An example of a solid-state accelerometer is the surface acoustic wave (SAW) accelerometer. A SAW accelerometer consists of a cantilever beam which is resonated at a particular frequency, as shown in Figure 9. A mass is attached to one end of the beam which is free to move. The other end is rigidly attached to the case. When an acceleration is applied along the input axis the beam bends. This causes the frequency of the surface acoustic wave to change proportionally to the applied strain. By measuring this change in frequency the acceleration can be determined.

4.1.3 MEMS Accelerometers

Micro-machined silicon accelerometers use the same principles as mechanical and solid state sensors. There are two main classes of MEMS accelerometer. The first class consists of mechanical accelerometers (i.e: devices which measure the displacement of a supported mass) manufactured using MEMS techniques. The second class consists of devices which measure the change in frequency of a vibrating element caused by a change of tension, as in SAW accelerometers.

The advantages of MEMS devices listed in Section 3.1.3 apply equally to accelerometers as they do to gyroscopes. They are small, light and have low power consumption and start-up times. Their main disadvantage is that they are not currently as accurate as accelerometers manufactured using traditional techniques, although the performance of MEMS devices is improving rapidly.




Figure 9: A surface acoustic wave accelerometer (source: [1]).

4.2 MEMS Accelerometer Error Characteristics

In this section we examine the errors which arise in MEMS accelerometers. The types of error described below are analogous to those in Section 3.2 for gyroscopes, and are therefore presented in less detail here. The important difference between errors arising from accelerometers is that they are integrated twice in order to track position, where as rate-gyro signals are only integrated once to track orientation.

4.2.1 Constant Bias

The bias of an accelerometer is the offset of its output signal from the true value, in m/s^2 . A constant bias error of ϵ , when double integrated, causes an error in position which grows quadratically with time. The accumulated error in position is

$$s(t) = \epsilon \cdot \frac{t^2}{2} \quad (10)$$

where t is the time of integration.

It is possible to estimate the bias by measuring the long term average of the accelerometer's output when it is not undergoing any acceleration. Unfortunately this is complicated by gravity, since a component of gravity acting on the accelerometer will appear as a bias. It is therefore necessary to know the precise orientation of the device with respect to the gravitational field in order to measure the bias. In practice this can be achieved by calibration routines in which the device is mounted on a turntable, whose orientation can be controlled extremely accurately.

4.2.2 Thermo-Mechanical White Noise / Velocity Random Walk

The output samples obtained from a MEMS accelerometer are perturbed by a white noise sequence. In Section 3.2.2 we showed that integrating white noise produces a random walk whose standard deviation grows proportionally to \sqrt{t} . Hence white noise on the output of an accelerometer creates a velocity random walk, usually specified with units $\text{m/s}/\sqrt{\text{h}}$. To see what effect accelerometer white noise has on the calculated position we can do a similar analysis in which we double integrate the samples obtained from an accelerometer. Let N_i be the i^{th} random variable in the white noise sequence, with $E(N_i) = E(N) = 0$ and $\text{Var}(N_i) = \text{Var}(N) = \sigma^2$. The result of double integrating the white noise signal $\epsilon(t)$ over a timespan $t = n \cdot \delta t$ is

$$\int_0^t \int_0^\tau \epsilon(\tau) d\tau d\tau = \delta t \sum_{i=1}^n \delta t \sum_{j=1}^i N_j \quad (11)$$

$$= \delta t^2 \sum_{i=1}^n (n-i+1) N_i \quad (12)$$

where n is the number of samples received from the device during the period, and δt is the time between successive samples. The expected error in position is

$$\mathbb{E} \left(\int_0^t \int_0^t \epsilon(\tau) d\tau d\tau \right) = \delta t^2 \sum_{i=1}^n (n-i+1) \mathbb{E}(N_i) \quad (13)$$

$$= 0 \quad (14)$$

and the variance is

$$\text{Var} \left(\int_0^t \int_0^t \epsilon(\tau) d\tau d\tau \right) = \delta t^4 \sum_{i=1}^n (n-i+1)^2 \text{Var}(N_i) \quad (15)$$

$$= \frac{\delta t^4 n(n+1)(2n+1)}{6} \text{Var}(N) \quad (16)$$

$$\approx \frac{1}{3} \cdot \delta t \cdot t^3 \cdot \sigma^2 \quad (17)$$

where the approximation assumes that δt is small (i.e. that the sampling frequency is large, which is a valid assumption for modern MEMS accelerometers). This analysis shows that accelerometer white noise creates a second order random walk in position, with zero mean and a standard deviation

$$\sigma_s(t) \approx \sigma \cdot t^{3/2} \cdot \sqrt{\frac{\delta t}{3}} \quad (18)$$

which grows proportionally to $t^{3/2}$.

4.2.3 Flicker Noise / Bias Stability

MEMS accelerometers are subject to flicker noise, which causes the bias to wander over time. Such fluctuations are usually modelled as a bias random walk, as described in Section 3.2.3. Using this model, flicker noise creates a second order random walk in velocity whose uncertainty grows proportionally to $t^{3/2}$, and a third order random walk in position which grows proportionally to $t^{5/2}$.

4.2.4 Temperature Effects

As with gyroscopes, temperature changes cause fluctuations in the bias of the output signal. The relationship between bias and temperature depends on the specific device, however it is often highly nonlinear. Any residual bias introduced causes an error in position which grows quadratically with time, as described in Section 4.2.1. If the IMU contains a temperature sensor then it is possible to apply corrections to the output signals in order to compensate for temperature dependent effects.

4.2.5 Calibration Errors

Calibration errors (errors in scale factors, alignments and output linearities) appear as bias errors which are only visible whilst the device is undergoing acceleration. Note that these ‘temporary’ bias errors may be observed even when the device is stationary due to gravitational acceleration.

4.2.6 Summary

The main sources of error for MEMS accelerometers are similar to those for gyroscopes. Table 3 summarises the error sources presented in this section. For MEMS devices angle random walk (noise) and uncorrected bias errors are typically the error sources which limit the performance of the device, however the relative importance of each error source depends on the specific device being used.

Error Type	Description	Result of Double Integration
Bias	A constant bias ϵ in the accelerometer's output signal	A quadratically growing position error $s(t) = \epsilon \cdot \frac{t^2}{2}$
White Noise	White noise with some standard deviation σ	A second-order random walk. The standard deviation of the position error grows as $\sigma_s(t) = \sigma \cdot t^{3/2} \cdot \sqrt{\frac{\delta t}{3}}$
Temperature Effects	Temperature dependent residual bias	Any residual bias causes an error in position which grows quadratically with time
Calibration	Deterministic errors in scale factors, alignments and accelerometer linearities	Position drift proportional to the squared rate and duration of acceleration
Bias Instability	Bias fluctuations, usually modelled as a bias random walk	A third-order random walk in position

Table 3: A Summary of Accelerometer Error Sources.




Figure 10: A possible log-log plot of Allan Deviation analysis results (source: [5]).

5 Signal Noise Analysis

In the previous sections we described a number of noise processes which arise in accelerometer and gyroscope signals. In this section we describe a technique known as Allan Variance, which can be used to detect and determine the properties of such processes. We then apply this technique to the accelerometer and gyroscope signals emitted from an Xsens Mtx device. The Mtx is a MEMS IMU containing three orthogonal gyroscopes, accelerometers and magnetometers. The device measures 38 mm × 53 mm × 21 mm, has a mass of 30 g, and is targeted mainly towards human motion capture applications.

5.1 Allan Variance

Allan Variance is a time domain analysis technique originally designed for characterising noise and stability in clock systems. The technique can be applied to any signal to determine the character of the underlying noise processes. The Allan Variance of a signal is a function of averaging time. For an averaging time t , the Allan Variance is computed as follows:

1. Take a long sequence of data and divide it into bins of length t . There must be enough data for at least 9 bins (otherwise the results obtained begin to lose their significance).
2. Average the data in each bin to obtain a list of averages $(a(t)_1, a(t)_2, \dots, a(t)_n)$, where n is the number of bins.
3. The Allan Variance is then given by

$$\text{AVAR}(t) = \frac{1}{2 \cdot (n-1)} \sum_i (a(t)_{i+1} - a(t)_i)^2 \quad (19)$$

To determine the characteristics of the underlying noise processes, Allan Deviation

$$\text{AD}(t) = \sqrt{\text{AVAR}(t)} \quad (20)$$

is plotted as a function of t on a log-log scale. Different types of random process cause slopes with different gradients to appear on the plot, as shown in Figure 10. Furthermore different processes usually appear in different regions of t , allowing their presence to be easily identified. Having identified a process it is then possible to read its numerical parameters directly from the plot. For a MEMS device such as the Mtx the important processes that we want to measure are random walk and bias instability, which can be identified and read as follows:




Figure 11: Allan Deviation plot for the Mtx gyroscopes.

	Bias Instability	Angle Random Walk
X Axis	$0.010^\circ/\text{s} = 36^\circ/\text{h}$ (at 620 s)	$0.075^\circ/\sqrt{\text{s}} = 4.6^\circ/\sqrt{\text{h}}$
Y Axis	$0.009^\circ/\text{s} = 32^\circ/\text{h}$ (at 530 s)	$0.078^\circ/\sqrt{\text{s}} = 4.8^\circ/\sqrt{\text{h}}$
Z Axis	$0.012^\circ/\text{s} = 43^\circ/\text{h}$ (at 270 s)	$0.079^\circ/\sqrt{\text{s}} = 4.8^\circ/\sqrt{\text{h}}$

Table 4: Gyroscope Noise Measurements.

- White noise appears on the Allan Deviation plot as a slope with gradient -0.5 . The random walk measurement for this noise (ARW for a rate-gyro, VRW for an accelerometer) is obtained by fitting a straight line through the slope and reading its value at $t = 1$.
- Bias Instability appears on the plot as a flat region around the minimum. The numerical value is the minimum value on the Allan Deviation curve.

For a full description of the Allan Variance technique see [5], which explains why each type of random process yields a slope with a particular gradient and derives the techniques for measuring each random process directly from the plot.

5.2 Xsens Mtx Analysis

The Allan Variance technique was applied to a 12 hour log of raw output data gathered from a stationary Mtx device with a 100 Hz sampling rate. Figure 11 shows the Allan Deviation curves for the x,y and z axis gyroscopes. The gradient of the slope on the left of the plot is almost exactly -0.5 , indicating the presence of white noise as expected. Table 4 lists the numerical noise values read from the graph for each gyroscope.




Figure 12: Allan Deviation plot for the Mtx accelerometers.

	Bias Instability	Velocity Random Walk
X Axis	$3.0 \cdot 10^{-4} \text{ m/s}^2 = 1.1 \text{ m/h}^2$ (at 670 s)	$0.0011 \text{ m/s}^2/\sqrt{s} = 0.066 \text{ m/s}^2/\sqrt{h}$
Y Axis	$2.8 \cdot 10^{-4} \text{ m/s}^2 = 1.0 \text{ m/h}^2$ (at 110 s)	$0.0011 \text{ m/s}^2/\sqrt{s} = 0.066 \text{ m/s}^2/\sqrt{h}$
Z Axis	$2.6 \cdot 10^{-4} \text{ m/s}^2 = 0.94 \text{ m/h}^2$ (at 620 s)	$0.0020 \text{ m/s}^2/\sqrt{s} = 0.12 \text{ m/s}^2/\sqrt{h}$

Table 5: Accelerometer Noise Measurements.

Figure 12 shows the Allan Deviation curves for the three accelerometers in the Mtx device. This graph shows that the z-axis accelerometer has a velocity random walk that is almost twice that of the other two accelerometers. Initially it was thought that this may be due to a gravitational effect (since the log used was recorded with the z-axis pointing upwards) or due to an error when measuring the scale factor of the device. We double checked that the scale factor was correct and performed the analysis on another log made with the z-axis aligned horizontally. The results were similar to those in Figure 12, hence we concluded that for the sample device tested the z-axis accelerometer was for some reason noisier than the x and y-axis devices. Table 5 lists the numerical noise values read from the Allan Deviation plot for each accelerometer.




Figure 13: Strapdown inertial navigation algorithm.

6 Strapdown Inertial Navigation

The strapdown navigation algorithm is shown in Figure 13. This section describes the algorithm in detail and outlines how errors which arise in the individual accelerometers and gyroscopes propagate through the algorithm. Throughout this section the subscripts b (body frame) and g (global frame) are used to indicate the frame of reference in which vector quantities are measured.

6.1 Tracking Orientation

6.1.1 Theory

The orientation, or attitude, of an INS relative to the global frame of reference is tracked by ‘integrating’ the angular velocity signal $\boldsymbol{\omega}_b(t) = (\omega_{bx}(t), \omega_{by}(t), \omega_{bz}(t))^T$ obtained from the system’s rate-gyroscopes. In order to specify the orientation of an INS one of several attitude representations must be used. Common representations include Euler angles, quaternions and direction cosines. In this section the direction cosines representation will be used to derive an algorithm for tracking the attitude. Similar derivations using both Euler angles and quaternions can be found in [1].

In the direction cosines representation the attitude of the body frame relative to the global frame is specified by a 3×3 rotation matrix \mathbf{C} , in which each column is a unit vector along one of the body axes specified in terms of the global axes. A vector quantity \mathbf{v}_b defined in the body frame is equivalent to the vector

$$\mathbf{v}_g = \mathbf{C}\mathbf{v}_b \quad (21)$$

defined in the global frame. The inverse transformation is given by

$$\mathbf{v}_b = \mathbf{C}^T \mathbf{v}_g \quad (22)$$

since the inverse of a rotation matrix is equal to its transpose.

To track the attitude of an INS we must track \mathbf{C} through time. If the attitude at time t is given by $\mathbf{C}(t)$ then the rate of change of \mathbf{C} at t is given by

$$\dot{\mathbf{C}}(t) = \lim_{\delta t \rightarrow 0} \frac{\mathbf{C}(t + \delta t) - \mathbf{C}(t)}{\delta t} \quad (23)$$

in which $\mathbf{C}(t + \delta t)$ can be written as the product of two matrices

$$\mathbf{C}(t + \delta t) = \mathbf{C}(t)\mathbf{A}(t) \quad (24)$$

where $\mathbf{A}(t)$ is the rotation matrix which relates the body frame at time t to the body frame at time $t + \delta t$. If $\delta\phi$, $\delta\theta$ and $\delta\psi$ are the small rotations through which the body frame has rotated between times t and $t + \delta t$ about its x, y and z axes respectively, then using the small angle approximation (see Appendix A) we can write $\mathbf{A}(t)$ as

$$\mathbf{A}(t) = \mathbf{I} + \delta\boldsymbol{\Psi} \quad (25)$$

where

$$\delta\mathbf{\Psi} = \begin{pmatrix} 0 & -\delta\psi & \delta\theta \\ \delta\psi & 0 & -\delta\phi \\ -\delta\theta & \delta\phi & 0 \end{pmatrix} \quad (26)$$

Hence by substitution:

$$\dot{\mathbf{C}}(t) = \lim_{\delta t \rightarrow 0} \frac{\mathbf{C}(t + \delta t) - \mathbf{C}(t)}{\delta t} \quad (27)$$

$$= \lim_{\delta t \rightarrow 0} \frac{\mathbf{C}(t)\mathbf{A}(t) - \mathbf{C}(t)}{\delta t} \quad (28)$$

$$= \lim_{\delta t \rightarrow 0} \frac{\mathbf{C}(t)(\mathbf{I} + \delta\mathbf{\Psi}) - \mathbf{C}(t)}{\delta t} \quad (29)$$

$$= \mathbf{C}(t) \lim_{\delta t \rightarrow 0} \frac{\delta\mathbf{\Psi}}{\delta t} \quad (30)$$

In the limit $\delta t \rightarrow 0$ the small angle approximation is valid, and

$$\lim_{\delta t \rightarrow 0} \frac{\delta\mathbf{\Psi}}{\delta t} = \mathbf{\Omega}(t) \quad (31)$$

where

$$\mathbf{\Omega}(t) = \begin{pmatrix} 0 & -\omega_{bz}(t) & \omega_{by}(t) \\ \omega_{bz}(t) & 0 & -\omega_{bx}(t) \\ -\omega_{by}(t) & \omega_{bx}(t) & 0 \end{pmatrix} \quad (32)$$

is the skew symmetric form of the angular rate vector $\boldsymbol{\omega}_b(t)$. Hence in order to track the orientation the attitude algorithm must solve the differential equation

$$\dot{\mathbf{C}}(t) = \mathbf{C}(t)\mathbf{\Omega}(t) \quad (33)$$

which has the solution

$$\mathbf{C}(t) = \mathbf{C}(0) \cdot \exp\left(\int_0^t \mathbf{\Omega}(t) dt\right) \quad (34)$$

where $\mathbf{C}(0)$ is the initial attitude of the device.

6.1.2 Implementation

Rather than providing a continuous signal $\boldsymbol{\omega}_b(t)$, an IMU actually provides samples of the angular velocity, usually at a fixed frequency. An integration scheme must be used to integrate the sampled signal. The choice of scheme is application dependent. For short timespan and low accuracy applications a low order scheme such as the rectangular rule may be sufficient. For more demanding applications a third or fourth-order scheme may be more appropriate. In this section the rectangular rule solution is presented.

Let the period between successive angular velocity samples be δt . For a single period $[t, t + \delta t]$ the solution to Equation 33 can be written as

$$\mathbf{C}(t + \delta t) = \mathbf{C}(t) \cdot \exp\left(\int_t^{t+\delta t} \mathbf{\Omega}(t) dt\right) \quad (35)$$

Using the rectangular rule we can write

$$\int_t^{t+\delta t} \mathbf{\Omega}(t) dt = \mathbf{B} \quad (36)$$

where

$$\mathbf{B} = \begin{pmatrix} 0 & -\omega_{bz}\delta t & \omega_{by}\delta t \\ \omega_{bz}\delta t & 0 & -\omega_{bx}\delta t \\ -\omega_{by}\delta t & \omega_{bx}\delta t & 0 \end{pmatrix} \quad (37)$$

and $\boldsymbol{\omega}_b = (\omega_{bx}, \omega_{by}, \omega_{bz})^T$ is the angular velocity sample corresponding to the update period. Letting $\sigma = |\boldsymbol{\omega}_b \delta t|$, substituting Equation 36 into Equation 35 and finally performing the Taylor expansion of the exponential term, we get

$$\mathbf{C}(t + \delta t) = \mathbf{C}(t) \left(\mathbf{I} + \mathbf{B} + \frac{\mathbf{B}^2}{2!} + \frac{\mathbf{B}^3}{3!} + \frac{\mathbf{B}^4}{4!} + \dots \right) \quad (38)$$

$$= \mathbf{C}(t) \left(\mathbf{I} + \mathbf{B} + \frac{\mathbf{B}^2}{2!} - \frac{\sigma^2 \mathbf{B}}{3!} - \frac{\sigma^2 \mathbf{B}^2}{4!} + \dots \right) \quad (39)$$

$$= \mathbf{C}(t) \left(\mathbf{I} + \left(1 - \frac{\sigma^2}{3!} + \frac{\sigma^4}{5!} \dots \right) \mathbf{B} + \left(\frac{1}{2!} - \frac{\sigma^2}{4!} + \frac{\sigma^4}{6!} \dots \right) \mathbf{B}^2 \right) \quad (40)$$

$$= \mathbf{C}(t) \left(\mathbf{I} + \frac{\sin \sigma}{\sigma} \mathbf{B} + \frac{1 - \cos \sigma}{\sigma^2} \mathbf{B}^2 \right) \quad (41)$$

which is the attitude update equation used to update \mathbf{C} as each new sample becomes available.

6.1.3 Propagation of Errors

The angular velocity signals obtained from the gyroscopes are ‘integrated’ by the standard INS attitude algorithm, therefore errors in the gyroscope signals propagate through to the calculated orientation. The isolated effects of each gyroscope error source on the integrated signal (in a single axis) were described in Section 3.2. For most MEMS devices white noise and uncorrected bias errors are the main causes of an error in the orientation. White noise causes an angle random walk whose standard deviation grows proportionally to the square root of time. An uncorrected bias causes an error in orientation which grows linearly with time. Quantisation errors also arise in the calculated attitude due to the quantisation of the angular velocity samples and due to the integration scheme used to update \mathbf{C} .

6.2 Tracking Position

6.2.1 Theory

To track the position of an INS the acceleration signal $\mathbf{a}_b(t) = (a_{bx}(t), a_{by}(t), a_{bz}(t))^T$ obtained from the accelerometers is projected into the global frame of reference:

$$\mathbf{a}_g(t) = \mathbf{C}(t) \mathbf{a}_b(t) \quad (42)$$

Acceleration due to gravity is then subtracted and the remaining acceleration is integrated once to obtain velocity, and again to obtain displacement:

$$\mathbf{v}_g(t) = \mathbf{v}_g(0) + \int_0^t \mathbf{a}_g(t) - \mathbf{g}_g dt \quad (43)$$

$$\mathbf{s}_g(t) = \mathbf{s}_g(0) + \int_0^t \mathbf{v}_g(t) dt \quad (44)$$

where $\mathbf{v}_g(0)$ is the initial velocity of the device, $\mathbf{s}_g(0)$ is the initial displacement and \mathbf{g}_g is the acceleration due to gravity in the global frame.

6.2.2 Implementation

As with the attitude algorithm, an integration scheme must be used to integrate the samples arriving from the accelerometers. A simple implementation using the rectangular rule uses the following update equations

$$\mathbf{v}_g(t + \delta t) = \mathbf{v}_g(t) + \delta t \cdot (\mathbf{a}_g(t + \delta t) - \mathbf{g}_g) \quad (45)$$

$$\mathbf{s}_g(t + \delta t) = \mathbf{s}_g(t) + \delta t \cdot \mathbf{v}_g(t + \delta t) \quad (46)$$

6.2.3 Propagation of Errors

Errors which arise in the accelerometers propagate through the double integration, as described in Section 4.2. This is the obvious cause of drift in the tracked position. Errors in the angular velocity signals also cause drift in the calculated position, since the rotation matrix \mathbf{C} obtained from the attitude algorithm is used to project the acceleration signals into global coordinates. An error in orientation causes an incorrect projection of the acceleration signals onto the global axes. This causes several problems. Firstly, the accelerations of the device are integrated in the wrong direction. Secondly, acceleration due to gravity can no longer be correctly removed.

In the strapdown algorithm $1g$ is subtracted from the (globally) vertical acceleration signal to remove acceleration due to gravity before the signal is integrated. A tilt error ϵ will cause a component of the acceleration due to gravity with magnitude $g \cdot \sin(\epsilon)$ to be projected onto the horizontal axes. This causes a residual bias due to gravity with magnitude $g \cdot \sin(\epsilon)$ to remain in the globally horizontal acceleration signals. There will also be a residual bias of magnitude $g \cdot (1 - \cos(\epsilon))$ in globally vertical axis, however this is much less severe since for small ϵ we have $\cos(\epsilon) \rightarrow 1$ and $\sin(\epsilon) \rightarrow \epsilon$. Hence the error in position caused by a small tilt error will occur mainly in the global xy-plane.

The propagation of gyroscope errors through to the calculated position is the critical error path in nearly all INS systems. In most applications the magnitude of g is much greater than the mean absolute acceleration of the IMU itself. In such cases the critical problem is that a component of the acceleration due to gravity is projected onto the globally horizontal axes. As a concrete example consider a tilt error of just 0.05° . This error will cause a component of the acceleration due to gravity with magnitude 0.0086 m/s^2 to be projected onto the horizontal axes. This residual bias causes an error in the horizontal position which grows quadratically to 7.7 m after only 30 seconds.

6.3 Mtx INS Example

In this section we analyse the performance of an example strapdown INS which uses the accelerometers and gyroscopes contained within an Xsens Mtx device. The INS attitude and position algorithms described in Sections 6.1.2 and 6.2.2 were implemented within an event based framework written in Java. A Java interface to the Xsens Mtx was also written within this framework to allow the INS algorithms to be applied to real data, either in real time or from a log file.

To test the performance of the system it was applied to 500 streams of data obtained from a stationary Mtx device. Each stream was 60 seconds in length. The device was aligned with its z-axis pointing vertically upwards⁵ and the initial position of the device was set to be the global origin $\mathbf{s}_g(0) = (0, 0, 0)^T$. The device's sampling frequency was set to 100 Hz . Since the device was stationary whilst the data was obtained, the total 'drift' for a particular run at some time t can be measured as

$$\text{drift}(t) = |\mathbf{s}_g(t)| \quad (47)$$

which is the Euclidian distance between the origin to the position returned by the INS algorithm at that time. Figure 14 shows how the average drift ($\text{drift}(t)$ averaged over all 500 runs of the algorithm) increases over time. The error bars indicate $\pm\sigma$ values.

After 60 seconds the average recorded drift was 152.67 m . Interestingly the average drift along the z-axis $\text{drift}_z(t) = |s_{gz}(t)|$ was only 1.76 m . Nearly all of the drift occurred in the global xy-plane. This is typical of an error in position caused by the propagation of a small tilt error in the orientation, as described in Section 6.2.3. This suggests that nearly all of the drift in the system can be attributed to an error in the orientation. To test this we re-ran the algorithm on each stream of data with accelerometer and gyroscope noise selectively removed. It is possible to remove noise from the signals because it was known that the device was stationary and aligned with the z-axis pointing vertically upwards. Hence to remove any noise in the accelerometer signals each sample in the stream is simply set to the true value $\mathbf{a}_b(t) = (0, 0, -g)^T$, where g is acceleration due to gravity. Similarly gyroscope noise is removed by setting each angular velocity sample to $\boldsymbol{\omega}_b(t) = (0, 0, 0)^T$. Figure 15 shows the results obtained by selectively removing noise in 10 sample runs. Figure 16 shows the mean drift suffered when gyroscope and accelerometer noise is selectively removed. Removing accelerometer noise has a negligible effect in

⁵Due to a lack of specialist equipment, this could not be done precisely





Figure 14: A plot showing how the average drift in position incurred the system increases over time.

reducing the drift, proving that errors in orientation propagating through the algorithm are by far the main source of error in the system.


Further insight into the short term performance of the system can be gained by examining the log-log version of Figure 16, as shown in Figure 17. The graph shows that initially accelerometer noise is the main cause of drift in the system, but only for 0.3 seconds. After this time the error in orientation caused by noise in the gyroscope signals becomes the main cause of drift. After the algorithm has been running for more than a few seconds the drift which can be attributed to accelerometer noise is insignificant.

6.4 Summary

In this section INS attitude and position algorithms have been described. The propagation of errors arising in the individual gyroscopes and accelerometers has also been discussed. In practice it is the accuracy of the gyroscopes and not the accelerometers which limit the overall accuracy of most INSs. An error in the computed orientation causes the body-frame acceleration signals obtained from the INS to be projected incorrectly onto the global axes. This results in a rapidly accumulating error in position when the signals are subsequently integrated. A simple INS was implemented based on the Xsens Mtx device. The stationary performance of the system was shown to be was poor, suffering an average drift in position of 152.67 m after 60 seconds. Nearly all of this drift can be attributed to orientation errors propagating through the INS algorithm. Drift which occurs due to accelerometer noise is only significant in the first few seconds of the algorithm's application.



(a) Zoomed out. The paths obtained after removing gyroscope noise are not visible.



(b) Zoomed in, showing the paths obtained after all gyroscope noise was removed.

Figure 15: Ten paths obtained by applying the strapdown algorithm to a stationary Xsens device for 60 second periods. The solid (red) paths are those obtained by applying the algorithm to the raw gyroscope and accelerometer data. The dashed (blue) paths were obtained by applying the algorithm to data stream after removing all accelerometer noise. The solid (light green) paths were obtained after removing all gyroscope noise. The scale is in meters.




Figure 16: The average drift with selectively removed noise sources.

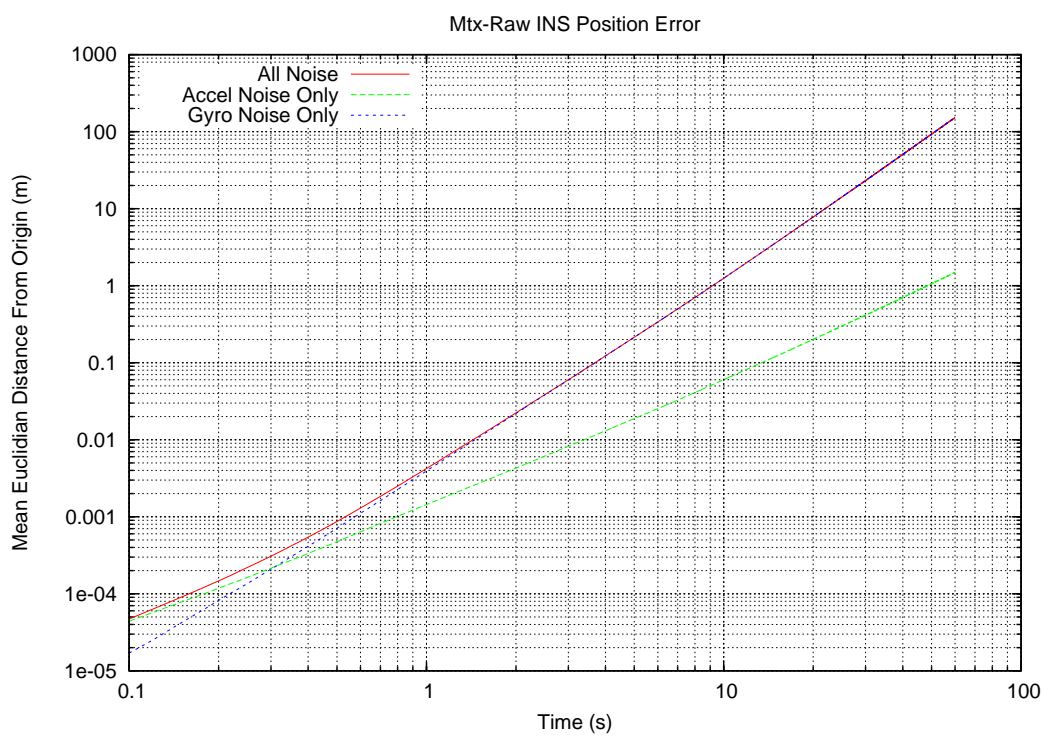


Figure 17: A log-log plot of the average drift with selectively removed noise sources.

7 Strapdown INS Simulation

Simulation can be used to further investigate the error characteristics of inertial navigation systems. The potential uses of INS simulation include analysing the performance of hypothetical systems, analysing the effect of individual noise processes, and establishing bounds for the performance of INS systems.

An INS simulation can be implemented by simply swapping the real IMU in an existing INS for a virtual measurement unit. This approach maximises code reuse by allowing the same INS algorithms to be applied to both real and virtual devices. Hence the problem of simulating an INS can be reduced to the problem of simulating the output signals of an IMU.

7.1 IMU Simulation

In order to simulate an IMU it is necessary to simulate the noise processes which perturb the gyroscope and accelerometer signals of a real device. We implemented a virtual IMU as part of the Java framework introduced in Section 6.3. The virtual device outputs sampled accelerometer and gyroscope signals. A given signal whose perfect samples are $(s_1, s_2, s_3...)$ can be perturbed by the following additive noise processes:

- A fixed bias error ϵ .
- White Noise. The samples are perturbed by an additive Gaussian white noise sequence, resulting in an angle or velocity random walk when applied to gyroscope or accelerometer signals respectively. A Gaussian white noise sequence is defined to be a sequence of random variables $(N_1, N_2, N_3...)$ in which each N_i is an identically distributed Gaussian random variable with a mean of zero and a finite variance σ^2 . Note that it is standard practice to assume that random errors follow a Gaussian distribution when modelling errors in INSs [1]. A white noise sequence is specified by its variance σ^2 .
- Bias instability, modelled as a random walk. The samples are perturbed by a random walk sequence, defined to be a sequence of random variables $(R_1, R_2, R_3...)$ where

$$R_i = \sum_{j=1}^i N_j \quad (48)$$

in which $(N_1, N_2, N_3...)$ is a Gaussian white noise sequence. A random walk is specified by the variance σ^2 of the underlying white noise sequence.

The output of a single channel of the virtual IMU is a sequence $(s_1 + e_1, s_2 + e_2, s_3 + e_3...)$ in which each s_i is a perfect sample and

$$e_i = \epsilon + N_i + R_i \quad (49)$$

is the total additive error perturbing it. Figure 18 shows the first 5 seconds of example white noise and random walk sequences (in this case perturbing a gyroscope signal).

7.2 Simulation Validation

To attempt to validate our simulation we repeated the experiment of applying the INS algorithm to 500 streams of data gathered from a stationary device, as described in Section 6.3. This time the streams of data were generated by a virtual IMU, whose white noise and bias instability parameters were set to correspond to those measured for the Mtx device using the Allan Variance technique, as described in Section 5.

It was necessary to convert the noise parameters measured using Allan Variance to σ values. Random walk measurements are converted to white noise σ values using the formula

$$\sigma = \frac{RW}{\sqrt{\delta t}} \quad (50)$$




Figure 18: The first 5 seconds of example random noise sequences.

	Bias Instability (σ value of the underlying white noise sequence)	White Noise (σ value of the white noise sequence)
X-Gyro ($^{\circ}/s$)	4.0×10^{-5}	0.75
Y-Gyro ($^{\circ}/s$)	3.9×10^{-5}	0.78
Z-Gyro ($^{\circ}/s$)	7.3×10^{-5}	0.79
X-Accel (m/s^2)	1.2×10^{-6}	0.011
Y-Accel (m/s^2)	2.7×10^{-6}	0.011
Z-Accel (m/s^2)	1.0×10^{-6}	0.020

Table 6: σ simulation noise parameters corresponding to the Mtx device.

where $\delta t = 0.01$ s is the sampling period of the device. For modelling bias instability with a random walk sequence, the σ value of the underlying white noise process is obtained using

$$\sigma = \sqrt{\frac{\delta t}{t}} \cdot \text{BS} \quad (51)$$

where t is the averaging time at which the bias stability measurement was made. The converted σ values are shown in Table 6.

Figure 19 shows how the mean drift results obtained from the real device in Section 6.3 compare to those obtained from the simulation. The drift grows significantly faster when using the real device, indicating that our simulation does not fully model all of the error sources which caused drift in the real system. This is not unexpected, since our simulation makes a number of assumptions. In particular it was assumed that the initial accelerometer and gyroscope biases were known exactly, which was not true for the tests conducted with the real device. A bias error of $0.03^{\circ}/\sqrt{s}$ in either the x or y gyroscope signals would account for the extra drift. Calibration errors and temperature dependent effects could also have caused errors to arise in the real device. These error sources are not modelled by the simulation.

Figure 20 shows how the simulated drift grows over time when the gyroscope and accelerometer noise processes are selectively removed. Reassuringly the plot looks almost identical to the equivalent plot generated using the real device (Figure 17). This shows that the relative importance of gyroscope and accelerometer errors is correctly modelled by the simulation. As in the real case accelerometer noise is the dominant cause of drift for around the first 0.3 seconds, after which the gyroscope noise becomes the dominant factor. Drift caused by the accelerometers becomes insignificant after a few seconds.

In conclusion we have shown that although our simulation does not fully model all sources of error in a real IMU, it does correctly simulate the relative importance of gyroscope and accelerometer errors. Results obtained from INS simulations can be used as bounds which the performance of equivalent real systems cannot exceed.




Figure 19: Average position drifts measured for the simulated and real devices.




Figure 20: A log-log plot of the simulated average drift with selectively removed noise sources.




Figure 21: A plot of the simulated average drifts when the gyro signals are perturbed by selective noise processes.

7.3 Gyroscope Errors

When analysing a real device it is possible to selectively remove acceleration and gyroscope noise, given that the true path followed by the IMU is known. It is not possible to selectively remove the effects of individual noise processes perturbing the inertial signals, since in general they cannot be separated. One of the benefits of an INS simulation is that it is possible to remove the effects of individual processes, allowing critical noise processes to be identified.

Figure 21 shows the average drift incurred by the simulated INS when the gyroscope signals were perturbed by both white noise and bias instability processes, just white noise and just bias instability. The graph shows that over 60 seconds white noise is the underlying process which accounts for most of the drift in the system.

7.4 Summary

In this section we have described how to implement a simple INS simulation based on the construction of a virtual IMU device. We implemented an INS simulation and verified it against results obtained using a real device. As expected it was found that our simulation does not fully model the errors which arise in a real device, although the results obtained can be used as bounds on the performance of a real system. Using the simulation it was determined that over a 60 second period white noise perturbing the gyroscope signals is the random process which accounts for most of the drift in an INS system based on the Xsens Mtx device. A toy version (shown in Figure 22) of the INS simulator can be downloaded from <http://www.cl.cam.ac.uk/~ojw28/ins.html>.

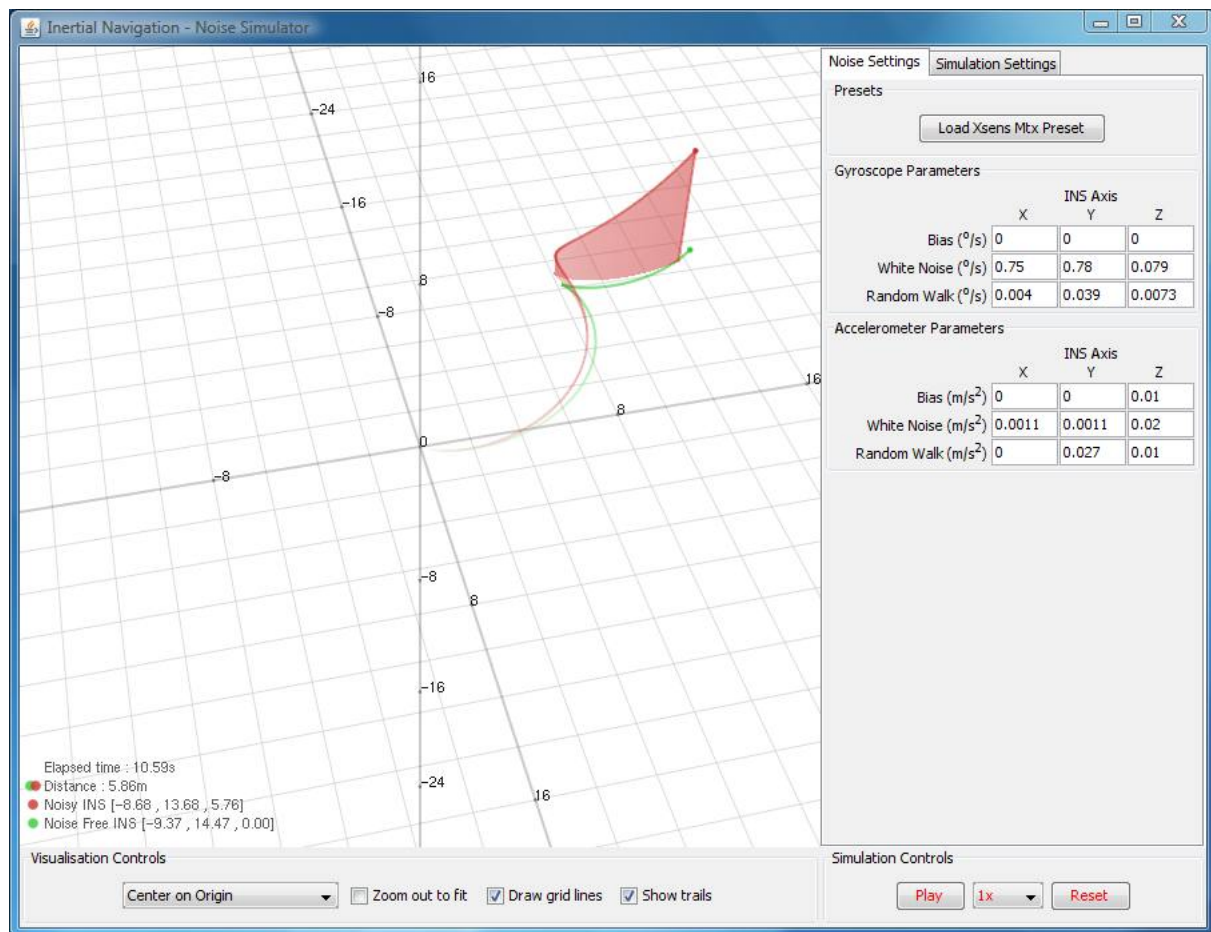


Figure 22: The simulation user interface.

8 Reducing Drift In Inertial Navigation Systems

In the previous sections we have demonstrated that ‘drift’, caused mainly by errors in attitude computation, is a major problem affecting pure inertial navigation systems. This section provides a brief overview of methods which can be used to reduce or correct drift. Such methods generally fall into one of two categories, the use of sensor fusion and the application of domain specific assumptions.

8.1 Sensor Fusion

Sensor fusion refers to processes in which signals from two or more types of sensor are used to update or maintain the state of a system. In the case of inertial navigation systems the state generally consists of the orientation, velocity and displacement of the device measured in a global frame of reference. A sensor fusion algorithm maintains this state using IMU accelerometer and gyroscope signals together with signals from additional sensors or sensor systems. There are many techniques for performing sensor fusion, the most popular of which are Kalman and particle filters. The implementation of such algorithms is beyond the scope of this report.

8.1.1 Fusion With Absolute Positioning Systems

One common approach is to periodically correct drift using position data from an absolute positioning system, such as GPS. Such systems try to combine the drift free nature of positions obtained from absolute positioning systems with the high sampling frequency of the accelerometers and gyroscopes which form the IMU. Roughly speaking the accelerometer and gyroscope signals are used to ‘fill in the gaps’ between successive updates from the absolute positioning system. GPS data is commonly used, however GPS positions can usually only be obtained outdoors and will typically provide absolute positions which are only accurate to around 15 m. This makes fusion with GPS data unsuitable for indoor use and for applications such as human motion capture where high accuracy is required. For indoor applications positioning systems such as the Ubisense⁶ system can be used, however such systems are usually expensive and require fixed infrastructure to be installed.

8.1.2 Fusion With Magnetometers

Another type of sensor commonly used to reduce drift is the vector magnetometer, which measures magnetic field strength in a given direction. Inertial measurement units often contain three orthogonal magnetometers in addition to the orthogonal gyroscopes and accelerometers. The magnetometers measure the strength and direction of the local magnetic field, allowing the north direction to be found.

Magnetometers are not accurate enough to replace gyroscopes in INSs. In particular they are affected by local disturbances in the earth’s magnetic field caused by nearby magnetic objects. Their data can however be fused with gyroscope data to improve the accuracy of the calculated orientation. The main advantage of this approach over the use of absolute positioning systems is that the navigation system remains self contained (it does not rely on any external infrastructure). The main disadvantage is that it only allows the rate at which position errors grow to be reduced, rather than allowing absolute corrections to be applied.

The Xsens Mtx device uses a proprietary sensor fusion algorithm to combine magnetometer and gyroscope data, outputting the orientation of the device relative to a global frame of reference. Figure 23 shows the improvement which can be gained from using magnetometer data. The solid line shows the average drift incurred when applying the INS algorithm to the gyroscope and accelerometer signals obtained from a stationary Mtx device, as described in Section 6.3. The dashed line shows the average drift when the attitude portion of the INS algorithm is replaced with the output of the Xsens sensor fusion algorithm. The average drift after 60 seconds when using sensor fusion is 5.21 m.

8.2 Domain Specific Assumptions

In some applications it is possible to make assumptions about the movement of the body to which the IMU is attached. Such assumptions can be used to minimise drift. One example in which domain specific

⁶<http://www.ubisense.net>




Figure 23: The performance gain obtained by using the Xsens sensor fusion algorithm rather than the INS attitude algorithm to calculate the orientation of the device.

assumptions are exploited is NavShoe [6], in which a shoe mounted IMU is used to track the location of a pedestrian. The assumption that a pedestrian’s foot has zero velocity when in contact with the ground is used to provide zero-velocity updates, allowing drift in velocity to be periodically corrected. By measuring the acceleration due to gravity when the device is stationary it is also possible to estimate and make adjustments for the tilt of the device.

The main disadvantage of using domain specific assumptions is that the assumptions must hold for the results to be valid. For instance NavShoe would fail should a pedestrian use an escalator. The benefits obtained from using assumptions must be weighed against the risk that they may be broken.

9 Conclusion

This report has provided an introduction to inertial navigation, including an analysis of the error characteristics of inertial navigation systems. We have identified the propagation of orientation errors caused by noise perturbing the gyroscope signals as the critical cause of drift in strapdown INS systems. A small tilt error in the calculated orientation causes a component of acceleration due to gravity to be projected onto the globally horizontal axes. This residual error is then integrated twice, causing a rapidly growing error in the calculated position.

We have implemented a simple INS algorithm and applied it to real and simulated IMU data. For the Xsens Mtx device we showed that errors arising in the gyroscope signals accounted for nearly all of the drift incurred by the system after it had been running for more than a few seconds. By simulation we showed that white noise was the most important noise process perturbing the gyroscope signals.

In practice it is necessary to reduce the drift of MEMS based INSs such as the one presented in this report. This can be achieved with sensor fusion techniques, often using additional data from absolute positioning systems and magnetometers. Drift can also be reduced by exploiting constraints which are known to apply to the movement of the IMU, such as known points in time at which the device must have a zero velocity.

The current state of the art MEMS inertial sensors (which include magnetometers to reduce drift in orientation) can be used to construct inertial navigation systems which suffer average drifts of around 5m after 60 seconds of stationary operation. It is not currently possible to construct an INS which maintains sub-meter accuracy for more than 60 seconds using MEMS devices.

In the future this report could be extended by analysing the performance of inertial navigation systems under dynamic motion, rather than just looking at the stationary case. In particular quantisation errors become much more significant when the device is undergoing rotations and accelerations, and the choice of integration scheme can have a big effect on the performance of the system. Other possible future work includes extending the simulator described in this report to model more sources of error.

References

- [1] D. Titterton and J. Weston. *Strapdown Inertial Navigation Technology*. The American Institute of Aeronautics and Astronautics, second edition, 2004.
- [2] Daniel Roetenberg. *Inertial and Magnetic Sensing of Human Motion*. PhD thesis, Universiteit Twente, 2006.
- [3] W Stockwell. *Angle Random Walk*. <http://www.xbow.com>.
- [4] E. Foxlin. *Handbook of Virtual Environment Technologies*, chapter Motion Tracking Technologies and Requirements, pages 163–210. Lawrence Erlbaum Publishers, 2002.
- [5] *IEEE Std 962-1997 (R2003) Standard Specification Format Guide and Test Procedure for Single-Axis Interferometric Fiber Optic Gyros, Annex C*. IEEE, 2003.
- [6] Eric Foxlin. Pedestrian tracking with shoe-mounted inertial sensors. *IEEE Comput. Graph. Appl.*, 25(6):38–46, 2005.

A Rotation Matrices

A rotation matrix is any orthogonal matrix whose determinant is equal to 1. The effect of multiplying a vector by a rotation matrix is to change its direction but not its magnitude. In 3-dimensional space rotations about the x, y and z axes are achieved using the following rotation matrices:

$$\mathbf{R}_x = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \phi & \sin \phi \\ 0 & -\sin \phi & \cos \phi \end{pmatrix} \quad (52)$$

$$\mathbf{R}_y = \begin{pmatrix} \cos \theta & 0 & -\sin \theta \\ 0 & 1 & 0 \\ \sin \theta & 0 & \cos \theta \end{pmatrix} \quad (53)$$

$$\mathbf{R}_z = \begin{pmatrix} \cos \psi & \sin \psi & 0 \\ -\sin \psi & \cos \psi & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (54)$$

Any rotation can be expressed as the product of the three matrices above

$$\mathbf{R} = \mathbf{R}_x \mathbf{R}_y \mathbf{R}_z \quad (55)$$

however it is important to note that the order in which the rotations are applied is important, since in general rotations in 3 dimensions are non-commutative.

A.1 Small Angle Approximation

If ϕ , θ and ψ are small, then $\sin \phi \rightarrow \phi$, $\sin \theta \rightarrow \theta$, $\sin \psi \rightarrow \psi$ and the cosines of ϕ , θ and ψ tend to 1. Using these approximations and ignoring the products of angles (since they are small), the order in which rotations around the individual axes are applied becomes unimportant, and the rotation matrix \mathbf{R} becomes:

$$\mathbf{R} = \mathbf{R}_x \mathbf{R}_y \mathbf{R}_z \approx \begin{pmatrix} 1 & -\psi & \theta \\ \psi & 1 & -\phi \\ -\theta & \phi & 1 \end{pmatrix} \quad (56)$$

This small angle approximation is used when tracking the attitude of an INS relative to a global frame of reference.