

A new approach to better low-cost MEMS IMU performance using sensor arrays

Henry Martin, Paul Groves
University College London, United Kingdom
Mark Newman
BAE Systems, United Kingdom
Ramsey Faragher
University of Cambridge, United Kingdom

BIOGRAPHY

Henry Martin is a PhD student at University College London (UCL) in the Engineering Faculty's Space Geodesy and Navigation Laboratory (SGNL). His project is concerned with improving navigation performance from a low-cost MEMS IMU. He is interested in inertial navigation, IMU error modelling, multi-sensor integration and calibration algorithms. He holds a Master of Mathematics (MMath) degree from Trinity College at the University of Oxford and an MSc in Advanced Mechanical Engineering from Cranfield University. (henry.martin.11@ucl.ac.uk)

Dr Paul Groves is a Lecturer (academic faculty member) at UCL, where he leads a program of research into robust positioning and navigation and is a member of SGNL. He joined in 2009, after 12 years of navigation systems research at DERA and QinetiQ. He is interested in all aspects of navigation and positioning, including multi-sensor integrated navigation, improving GNSS performance under challenging reception conditions, and novel positioning techniques. He is an author of more than 50 technical publications, including the book Principles of GNSS, Inertial and Multi-Sensor Integrated Navigation Systems, now in its second edition. He is a Fellow of the Royal Institute of Navigation (RIN) and serves as an Associate Editor of both the RIN and ION journals. He holds a bachelor's degree and doctorate in physics from the University of Oxford. (p.groves@ucl.ac.uk)

Mark Newman graduated with a master's degree in electrical and electronic engineering (MEng) from the University of Aberdeen in 2001. Since then he has worked at BAE Systems and is now a Principal Design Engineer working on a variety of cutting edge technologies and future capabilities. His current field of work is focused on indoor and outdoor GPS denied positioning technologies for both ground and air based platforms. This forms part of BAE Systems' NAVSOP GPS-denied positioning system. (mark.newman3@baesystems.com)

Ramsey Faragher is a Senior Research Associate in the University of Cambridge working on GNSS-denied positioning, sensor fusion and machine learning schemes for consumer devices. Prior to this he was a Principal Scientist at the BAE Systems Advanced Technology Centre where he developed the NAVSOP GPS-denied

positioning system. He is also currently an Associate Editor for the journal of the RIN. (ramsey@cantab.net)

ABSTRACT

Over the past decade and a half, improvements in micro-electro-mechanical sensors (MEMS) technology and multisensor integration has enabled inertial sensors to be deployed over a much wider range of navigation applications [1]. However, factory calibration of these devices currently increases their cost by around \$1000 per unit, this means that adding additional sensor triads is economically viable if it means that this expensive calibration can be avoided.

This paper presents three advanced array-based techniques that could be applied to improve the performance of low-cost MEMS IMUs. These use knowledge of the sensor characteristics to obtain a more accurate estimate of specific force and/or angular rate than a simple average of the array's constituent sensors, which is the technique commonly used in the existing literature.

The three techniques presented are:

1. Arranging the sensors so that their sensitive axes are in opposing directions, so as to significantly reduce the effect systematic errors that are correlated across sensors of the same design.
2. Exploiting the performance differences between the in-plane and out-of-plane sensors on a sensor triad.
3. Combining the output of sensors with different operating ranges to increase the accuracy of the measurement during periods of relatively low dynamics without clipping and distorting under high dynamics.

In order to test the feasibility of these ideas we constructed a hardware platform that can record the output of an array containing multiple MEMS inertial sensors. Several experiments were conducted with this hardware to make characterisations of the systematic and stochastic errors of these sensors. The results are used to determine the feasibility of each technique for each sensor.

We find that for each technique proposed at least one of the models of inertial sensor tested could potentially benefit.

1. INTRODUCTION

Over the past 15 years, rapid advances have been made in micro-electro-mechanical sensors (MEMS) technology, which allows cheaper, smaller and lower power consumption inertial measurement units (IMUs) and inertial navigation systems (INSs). This has opened up new applications for IMUs across a range of navigation technologies [1]. Examples include pedestrian dead-reckoning using step detection technology [2, 3], aiding of GNSS signal tracking during jamming [4, 5], and simultaneous localisation and mapping (SLAM) using radio signals [6]. Inertial sensors can also be used for context detection enabling a navigation system to adapt to changes in the surrounding environment and host behaviour [7].

However, to achieve optimal performance, a MEMS IMU must be calibrated in the factory [15], which increases the cost from a few tens of dollars to more than \$1,000. In this paper, we present and investigate several techniques for using consumer-grade (i.e. low-cost) MEMS IMUs without prior factory calibration. The viability and effectiveness of the techniques will be assessed.

The process by which an INS-based navigation system works is illustrated in Figure 1. It shows how the various techniques used to improve navigation performance fit into the chain. For example, the integration of other sensors (such as GNSS) fits in between the IMU derived estimates of (changes in) user position and, by way of an integration algorithm, the information from the other sensors is used to come up with a combined estimate of position [1]. In this paper we are focussing on sensor level techniques, that is, techniques which use the output signals of the accelerometers and gyroscopes and improve the estimates of specific force and angular rate available from them.

This research is applicable to many different low-cost navigation and attitude determination applications. For micro air vehicles (MAVs), it can be used to improve the performance of both the attitude and heading reference system (AHRS) critical to flight control and the INS/GNSS integrated navigation system. It can improve the performance of foot-mounted inertial navigation for sports, dismounted soldier navigation and indoor positioning database generation. It can enhance robot and land vehicle navigation, where inertial sensors are often

integrated with odometry (wheel-based or visual). Finally, IMUs are increasingly being deployed in medicine to enable surgical instruments to be positioned more precisely.

The rest of this paper proceeds as follows. In Section 2 three advanced sensor array techniques will be proposed which could be used to improve navigation performance. In Section 3 the development of a hardware test-bed required to examine the feasibility of the techniques proposed is detailed. Section 4 outlines a series of experiments which were conducted with this hardware, to characterise various sensor errors. The results of these experiments are presented in Section 5 with both stochastic (Section 5.1) and systematic (Section 5.2) errors being examined. In Section 6 the implications of these results for the sensor level ideas proposed in Section 2 are examined. The conclusions are presented in Section 7 and future work in Section 8.

2. ADVANCED SENSOR ARRAY TECHNIQUES

In order to improve the performance of an inertial navigation system additional information is required. This additional information might be a characterisation of the motion profile or it might be information from another kind of sensor (e.g. GNSS). Here, we consider using extra inertial sensors to provide this additional information. Thus, we propose new ways of using arrays of two or more similar IMUs all measuring the same specific force and angular rate.

As discussed in Section 1, the cost of the laboratory calibration process is much higher than the hardware cost, so using multiple inertial sensors is economic. This has been tried before, although in almost all previous cases the approach has been simply to take an average specific force/angular rate signal from the array [8], or to track of the position solutions from two INSs separately and use these for fault detection and integrity monitoring [9]. In this paper we investigate more sophisticated techniques by examining: common-mode errors in sensors of the same design (Section 2.1); differences between the in-plane and out-of-plane sensors on an accelerometer or gyroscope triad (Section 2.2); and the complementary properties of sensors with different measurement ranges (Section 2.3). These techniques have the potential to give significantly better performance than simple averaging.

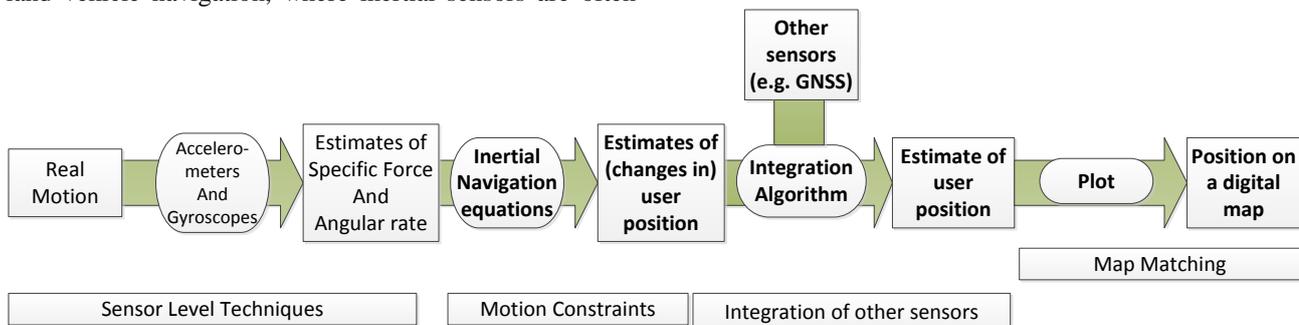


Figure 1: The Navigation system processing chain. The positions in which different techniques to improve navigation performance fit are underneath.

2.1 The common-mode errors of different sensors of the same design

There are many types of errors that an inertial sensor can be subject to, both systematic and stochastic. These include noise (stochastic), bias, scale factor errors, temperature sensitivity (systematic) and many others [1, 10]. The specifications of most inertial sensors state that particular types of error may be positive or negative by the same amount. One may naively assume that the error of a device is drawn randomly from this range, although a correlation between devices can be shown. Yuksel et al. [11] showed that for two ADIXRS150 gyroscopes the ‘zero-rate level change vs. temperature’ of one axis was in the same direction and similar magnitude for both sensors. Thus if the 3-axis sensors were arranged so their sensitive axes were facing in opposite directions when the output was combined their ‘bias drift with temperature’ would at least partially cancel out. This arrangement is shown in Figure 2.

In this paper we test the idea proposed by Yuksel et al. for temperature drift compensation and gyroscope g-dependent error for different MEMS sensors. We also aim to extend this idea to compensate other errors, which depend on an even power of the sensors input and are correlated between sensors of the same design. This includes all bias-like errors (0^{th} power) and components of non-linearity which depend on the output squared but it will not work for odd-order errors, such as scale factor errors, or errors which are not correlated between two sensors of the same model. Therefore one should expect this concept to work for some errors on some sensors, and it certainly will not work for every error on all sensors.

However, even when this idea does not help it does not produce any negative effects. For example, the noise averaged from two anti-parallel sensors will be no different to the noise from two sensors in the same orientation, similarly first order errors or uncorrelated errors will be unaffected. This means that there is no particular downside to using this technique.

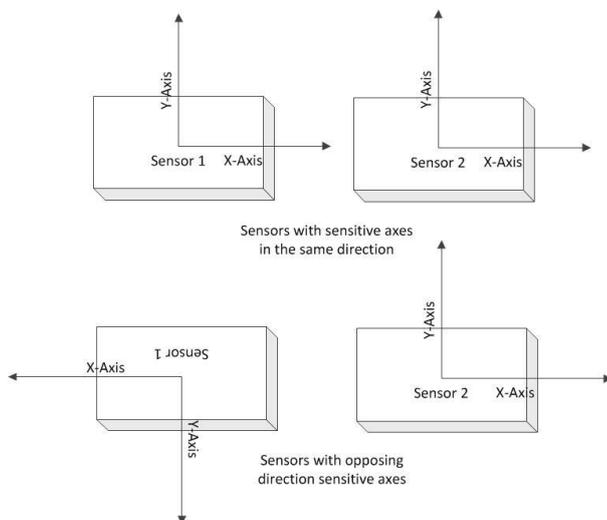


Figure 2: An illustration of sensors with aligned and opposing x and y sensitive axes.

2.2 The different characteristics of in-plane and out-of-plane sensors

MEMS inertial sensor fabrication is often conducted using technology developed for silicon micro-processors. This is based on building up layers of thin silicon wafers and thus there is not full freedom to create 3-D structures [10]. For 3-axis inertial sensors, usually all three sensors are constructed as a unit, so they must share the same construction plane. As a result, while the two sensors sensitive to motion along or about the in-plane axes (typically x and y) may be of an identical design rotated by 90 degrees, the out-of-plane sensor (typically z) must be of a different design.

In some cases the manufacturer specifies that the performance of the out-of-plane sensor is different e.g. for the ADXL345 [13]. In other cases the specifications are the same for in- and out-of-plane sensors, such as the BMA180 [14]. If the properties are significantly different the navigation performance of the system will become non-isotropic, perhaps with larger drift in one direction than the others. An array could be constructed, that combined in-plane and out-of-plane sensors in the same direction by having the MEMS sensors mounted on two perpendicular circuit boards. This would enable a sensor of the better performing design to be used for every axis.

2.3 The complementary properties of MEMS sensors with different measurement ranges.

The final idea we shall be examining in this paper concerns arrays of MEMS inertial sensors with different measurement ranges.

We define the measurement range or full scale of a MEMS sensor as the maximum specific force, or angular rate, that it can measure; this might be for example $\pm 4g$ or ± 500 degrees per second (dps), respectively. The fundamental assumption of this idea is that if one has two sensors of the same quality but different measurement range the lower-range sensor will measure small inertial forces more accurately.

The validity of this assumption is tested for three particular models of MEMS inertial sensor later in this paper. However one error that certainly will scale with measurement range is quantisation error. This is created when a continuous quantity is converted to a discrete number of levels by an analogue-to-digital converter (ADC). If two sensors both have the same resolution ADCs, the number of discrete levels they can measure will be the same. So the smallest increment that can be measured will scale with the measurement range, and thus the quantisation error will increase for higher measurement range sensors. Additionally we hypothesise that the magnitude of other errors such as noise might scale with dynamic range, as it is might be equally difficult to make a $\pm 16g$ accelerometer accurate to ± 160 milli-g as to make a $\pm 4g$ accelerometer accurate to ± 40 milli-g.

The technique proposed here to exploit this property is to combine the outputs of an array of inertial sensors according to the magnitude of the inertial forces sensed. While the system is experiencing low-dynamics the lower-range (assumed to be higher accuracy) sensor is used, but when the forces exceed the range of this sensor the high-range sensor is used, thus avoiding the signal being ‘clipped’. This is where any signal where the magnitude is outside of the measurement range is read as the maximum signal that can be sensed.

The simplest implementation of such a system would be a switch between two (or more) sensors, where at some threshold (close to the maximum measurement range of the low-range sensor) the array’s output switches from 100% low-range sensor to 100% high range. However for a number of reasons this is not optimal. For example, as a sensor gets close to the extremes of its measurement range one might assume that its reliability and accuracy become lower as it gets close to ‘maxing out’, as for instance the effect of nonlinearity increases [21]. This makes an abrupt switch undesirable. Also as there is considerable random noise on both sensors signals, the high-range sensor could still potentially provide useful information in the low-dynamics domain, this could be used to ‘average-out’ some of the noise (or other errors). However a weighted combination of the two outputs could be designed that is unaffected by these potential issues. The weighting system at very low dynamics would favour the low-range sensor, and its relative contribution would decrease as the dynamics increase reaching 0% at the end of its range. However, the precise weighting factor used should depend on the relative improvement in performance given by the decrease in measurement range, and its optimal composition will be the subject of future work

For all the three ideas proposed here, the key question is whether the performance improvement justifies the increased size, cost and power consumption of adding the additional sensors.

3. EXPERIMENTAL HARDWARE DEVELOPMENT

In order to assess the feasibility of the ideas presented in Section 2 it is necessary to have a hardware test-bed that can be used to record the output from an array of MEMS inertial sensors. In this section, the hardware and firmware used are described: the requirements for the design in Section 3.1 and the features selected to fulfil those requirements in Section 3.2.

3.1 Requirements

The main requirement for the hardware test-bed is that it should be able to log the output data of an array of MEMS inertial sensors, without a limit on the number of sensors of a particular type that could be connected. The second requirement is that sensors must be able to be turned on and off and re-configured in the on-board firmware without requiring any physical re-arrangement (or re-soldering) of the system.

It is also desirable that the test-bed be able to operate and log data free of a wired connection for several hours at a time. To perform some of the tests required the test-bed must be mounted in an orthogonal sided (precisely cuboid) frame. A frame, shown in Figure 3, was rapid-prototyped from nylon and the orthogonality of its sides checked.

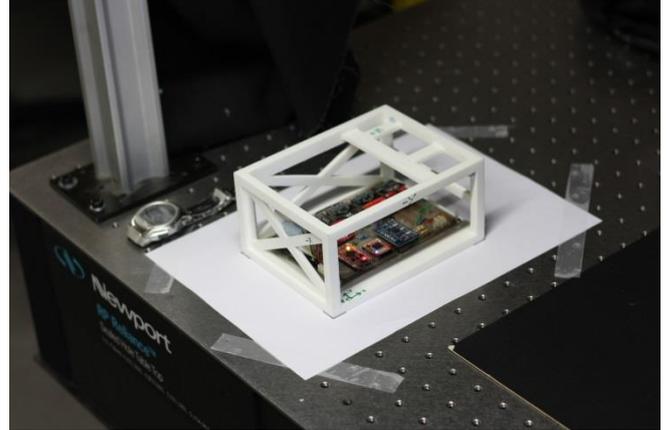


Figure 3: The Hardware test-bed. The box reference frame is visible.

3.2 Selection of features

The hardware built for this paper is based around the Arduino hardware and software platform [<http://www.arduino.cc/>]. This platform was primarily selected for ease of use and the simplicity of re-programming. The model used is an Arduino Pro Micro 3.3V. The Arduino both supplies all the sensors with a regulated power supply and acts as the master for communication with the sensors.

In this paper we test 3 models of MEMS inertial sensors. They are all ‘digital sensors’, i.e. sensors with analogue-to-digital converters (ADCs) built into the sensor chip. They are Bosch BMA180 accelerometers, Analogue Devices ADXL345 accelerometers and ST Microtronics L3G4200D gyroscopes [12, 13, 14]. The test-bed contains 4 BMA180s, 2 ADXL345s and 3 L3G4200Ds. These will be referred to in the results and figures as BMA1, BMA2, BMA3, BMA4, ADXL1, ADXL2, Gyro1, Gyro2, and Gyro3. The sensors are (deliberately) positioned on the circuit board so that their sensitive axes are in different directions. Their nominal orientations are given in Table 1. The array reference frame is defined relative to the sides of the box.

All the inertial sensors chosen have selectable measurement ranges, that is, they can be run at different operating ranges selectable through firmware. This was done because in order to test this assumption in Section 2.3, we need to pick sensors which are 'of the same quality' but have different measurement ranges. MEMS inertial technology has been increasing in performance rapidly [1, 10]. So, even if we took two sensor models from the same manufacturer, it is possible that one would use a design and/or a construction technique that is older and lower performance, despite similar prices.

However, using sensors that have a programmable measurement range enables us to test the factor of measurement range setting in isolation. We chose the three sensors used in the test-bed with this in mind. However, the ADXL345 accelerometer has a fixed conversion from LSB to milli-g (4mg/LSB) for all measurement ranges. This means that while it is a 13-bit sensor at 16g (although it is not really, see Section 5.1.2) it is only a 10-bit sensor at 2g, so we are not comparing like with like, as such we did not test it at multiple measurement ranges.

Sensor Name	Sensitive axis corresponding to box frame		
	+X	+Y	+Z
ADXL1	-X	-Y	+Z
ADXL2	-X	-Y	+Z
BMA1	+Y	-X	+Z
BMA2	-Y	+X	-Z
BMA3	+Y	-X	-Z
BMA4	-Y	+X	+Z
Gyro1	+X	+Y	+Z
Gyro2	+X	-Y	-Z
Gyro3	-X	-Y	+Z

Table 1: Approximate sensor sensitive axis orientation with respect to the body frame (also see Figure 4).

Additionally the hardware test-bed includes a Bosch BMP180 temperature and pressure sensor [16], which acts as an independent temperature sensor. Two Honeywell HMC5883 magnetometers [17] are also incorporated but are not used for this paper. The test-bed is shown in Figure 4.

The sensors communicate with the microcontroller by the inter-integrated-circuit (I2C, or I^2C) protocol, which allows a large number of devices to communicate with a master device using only two wires as each slave device is assigned a fixed ‘address’ [18]. However, as each of the MEMS inertial sensors used here can only be set to communicate on one of two hardwired addresses, communicating with more than two of a particular type requires a ‘bus-splitter’. In this case we use a dsscircuits.com I2C multiplexer board, containing an NXP Semiconductors PCA9544A 4-channel I2C-bus multiplexer.

The power, which is regulated by the Arduino, is supplied from an 850mAh lithium polymer battery. This gives a running time of in excess of 18 hours on a single charge.

The firmware running on the Arduino board is written in the Arduino language (similar to C++), using code adapted from the manufacturer supplied libraries. The firmware works in two parts. First, on start-up, the microcontroller configures each sensor, including setting programmable registers such as measurement range and output data rate. Second, the firmware continuously loops through recording a timestamp, reading all the sensors (unless only a subset of them are required) and then writing all the data to the Micro SD card

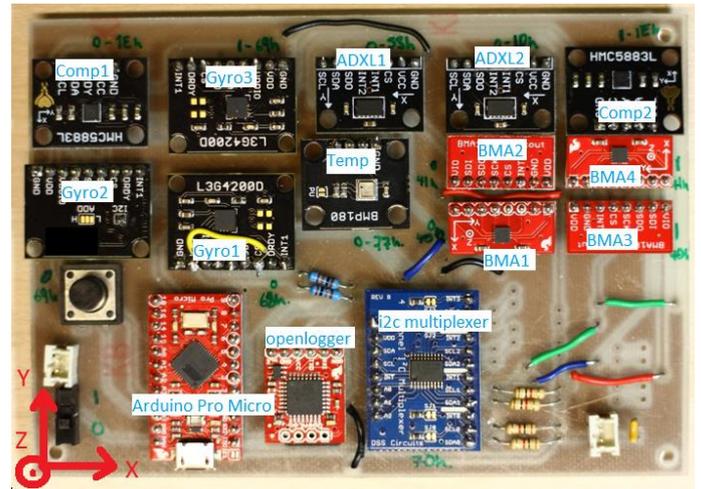


Figure 4: The hardware test bed, major components are labelled. The frame axes included in the bottom left hand corner.

4. EXPERIMENTS

Using the test-bed described in Section 3, several different types of experiment were conducted in order to analyse the characteristics of the sensors and thus determine the feasibility of the ideas presented in Section 2.

To conduct these experiments we did not have access to a sophisticated calibration rig or high performance reference IMU. However this is not a problem as we do not intend to replicate factory calibration, rather to see what is possible with only equipment that would be available to the end-user.

The main classes of experiments conducted were:

- static one-sensor experiments to maximise logging-rate;
- experiments where all the sensors were run and data was recorded for a specific time period on each of the faces of the ‘calibration cube’;
- static experiments where the test bed started at room temperature and was heated and then allowed to cool; and
- a final set of static experiments to compare the turn-on and in-run bias variation.

All of the experiments were conducted on a Newport RP reliance optical table, which is both damped and flat to eliminate contamination of the measurements from external vibration and improve repeatability.

The static one-sensor experiments were conducted separately for each of the sensors in order to maximise the logging rate allowing more detailed analysis of the sensor’s noise profile. For these tests the test bed was positioned on its side (with the face ‘+x’ uppermost), this was done in order to apply the majority of the specific force from the reaction to gravity to one of the in-plane sensors. This is intended to isolate any difference in the stochastic noise which might occur between the in- and out-of-plane sensors from that occurring when the sensor

is reading a specific force ($\sim 1g$ rather than $\sim 0g$) or g -dependent gyroscope behaviour. The results of these high rate individual samples are presented in Section 5.1.

The other classes of experiments conducted all used the entire sensor set and thus the logging rate was reduced by approximately a factor of 12 (as there were twelve sensor triads). There were two basic kinds of multi-sensor experiment conducted, the ‘tumbling cube’ and the heating experiment. Both of these experiments were mostly used for determining systematic errors that vary slowly with time and thus the low logging rate can be offset by taking longer samples.

In the ‘tumbling-cube’ experiment, static readings were taken on each of the six faces of the ‘calibration-cube’ in a particular order before returning to the first face to check that there had not been significant bias drift during the experiment. The analysis presented in this paper is based on a 15-minute sample on each side, although experiments were also conducted with two-minute samples. This type of experiment can be used to determine accelerometer biases, misalignments, cross-coupling and scale factors and to determine gyroscope biases and g -dependence.

In order to determine the sensitivity of the sensor’s biases to temperature, a static experiment was run where the sensors were heated with a hairdryer. The test bed was static on the optical table with the face ‘+Z’ uppermost, the hairdryer was a distance of approximately 40cm from the test-bed, from where its flow was sufficiently dispersed that it heated the area around the test-bed evenly. This is depicted in Figure 5.

The sensors were left to run at room-temperature for 5 minutes, they were then heated by the hairdryer at medium-heat for 5 minutes, then left to cool for a further 5 minutes then heated on high heat for 5 minutes followed by being left to cool for approximately 90 minutes.

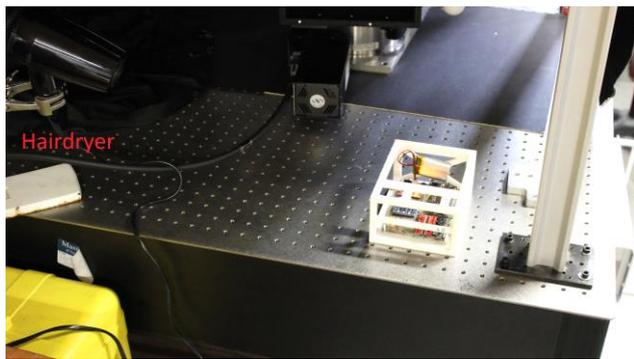


Figure 5: Heating experiment set-up. The hairdryer is in the top left corner of the picture.

The final type of experiment conducted for this study was to determine the turn-on-bias. For this experiment the hardware was in a single position and orientation. The test-bed was switched on, recorded 5 minutes of data, then switched off for at least 2 minutes before the process was repeated, collecting a total of 15 5 minute samples over a

2 hour period. This is then compared with 15 5 minute samples taken in post-processing from a single 100 minute collection. These experiments allow the switch-on and in-run bias variations to be determined and compared.

5. CHARACTERISATION OF SENSORS

In this section we present the properties of the sensors that were determined from the experiments detailed in Section 4. As different techniques are used to examine stochastic and systematic (approximately noise-like and bias-like) errors we examine each of these separately in sections 5.1 and 5.2, respectively.

5.1 Stochastic errors

When examining the stochastic errors of a sensor we are essentially asking the question “how white is the noise?”. This is because white noise is uncorrelated (in time) and so unpredictable, which means that its effect cannot be removed from a signal by modelling, so we are mostly interested in non-white components of the noise, which we could model. Despite this accurately knowing the magnitude of the white noise is needed for optimal sensor integration [1].

To examine the stochastic errors we primarily use the data from the single-sensor high-rate experiments described in Section 4. A standard investigation using Allan deviation revealed the sensors’ measurement noises to be white for the range 0.003 to 1400 seconds ($\sim 350Hz$ data). Except in the case of the L3G4200D gyroscope which showed some correlation at the shortest sample intervals, suggesting the true bandwidth is lower than specified in [12], and some small periodic components in the Y-sensor. A more detailed analysis of the stochastic errors is included in Appendix A.

A brief summary of the gyroscope high rate experiment is presented in Table 2. Note that the standard deviation of the output does not change (in physical units) with measurement range. The bias does not appear to be correlated between the two range settings despite both being experiments with the same sensor.

Output	Range	Mean (Bias)	Mean (Bias)	Standard Deviation of the data	
		LSB	dps*	LSB	dps*
Gyro1 x	250dps	14.174	0.1240	36.816	0.3221
Gyro1 y	250dps	-17.607	-0.1541	35.512	0.3107
Gyro1 z	250dps	-7.644	-0.0669	35.193	0.3079
Gyro1 x	2k dps	6.386	0.4470	4.654	0.3258
Gyro1 y	2k dps	-6.789	-0.4752	4.530	0.3171
Gyro1 z	2k dps	12.695	0.8887	4.683	0.3278

Table 2: Summary Statistics for the gyroscope static high-rate experiments. * The conversion to dps indicates using the “typical” scale factor from the data sheet [12].

5.2 Systematic errors

In this section we will examine the systematic errors, that is, the errors which are either fixed or depend on an environmental factor, such as temperature. Examples include fixed bias, scale factor error and bias variation with temperature. This type of error will vary significantly between sensors of the same model, due primarily to

random manufacturing differences, and thus the sensors will have their performance for these errors specified in terms of the distribution covering the whole population. However, for a particular sensor these errors will vary over time by a much smaller amount. Determining and compensating for these errors is the purpose of a laboratory calibration and the performance increase inferred in mostly because of it.

In order to remove the stochastic part of the signal so that the systematic errors can be determined, we need to take many data samples under a particular set of conditions and then take the mean of this output. This mean is still affected by the noise, however the standard error of the mean is reduced by a factor of \sqrt{n} , where n is the number of samples which it is averaged over. Thus the 1-sigma uncertainty of this estimate of the mean output for a sample of length n is (standard deviation of sample set) / \sqrt{n} . This figure is included in several tables so the reader can compare the quantity measured with the accuracy of the estimate of the mean.

The analyses of most of the systematic errors in this section are derived from a tumbling cube experiment. The output of this experiment is presented in Table 10, attached in the appendix. As this experiment takes more than 90 minutes and the sensors are very sensitive to changes in temperature the first orientation was repeated at the end of the experiment. The order in which the readings appear in Table 10 is that in which they were carried out, thus it is possible to observe a slight upward trend in temperature through the experiment. This seems to particularly affect the final ‘check’ reading, so it is likely that the apparent drift calculated using this might be overestimated.

The mean (across the six samples) of the 1-sigma uncertainty of the estimate of the mean (SD/\sqrt{n}), as mentioned above, is presented in Table 10 and also the mean of each sample’s standard deviation. These figures vary very little for each of the six samples as the standard deviation does not vary significantly with the sensor orientation relative to gravity and the samples are all the same length to within a few seconds, and as the 15 minutes represents more than 40,000 epochs this 1-sigma uncertainty is very small.

The rest of this section is structured as follows. In Sections 5.2.1 and 5.2.2 the systematic errors calculated using this six-sided experiment for the accelerometers and gyroscopes, respectively, are presented. In Sections 5.2.3 and 5.2.4 the in-run-stability and run-to-run stability are examined using a series of static experiments. In Section 5.2.5 the bias variation with temperature is examined using an experiment where the sensors were heated and then cooled.

5.2.1 Accelerometer Systematic Errors

It is possible to estimate several of the systematic errors of the output of the six-side experiment. Estimates of each

sensor’s bias, scale factor and attitude of its ‘apparent’ sensitive axis with respect to the can be calculated. The final two of which depend on the bias and also on each other. The method used for this is presented in Appendix B.

The systematic errors calculated from the ‘tumbling cube’ experiment are presented in Table 3. The ADXL345 datasheet [13] specifies that the bias should be $\pm 150\text{mg}$ for X and Y axes and $\pm 250\text{mg}$ for z, so the sensors are clearly well within specification. The scale factor is specified as 230 to 282 (256 ‘typical’) LSB/g, so this is also clearly within specification, although $\pm 10\%$ is a very loose specification.

The BMA180 datasheet [14] specifies the zero-offset as $\pm 60\text{mg}$, so both BMA1 and BMA2 (the two at 16g range) are apparently out-of-specification, although strictly the specification is for 2g measurement range. The specified scale factors for the two range settings tested are $5460 \pm 2\%$ (5351 to 5569) and $512 \pm 3\%$ (496.6 to 527.4) for 1.5g and 16g respectively. Thus, all 4 triads are within specification.

The differences between the sensors’ sensitive axes are noteworthy. Observe that the ADXL345 in-plane (x and y) sensors all have a positive bias and the scale factors are in the range of 261-264 LSB/g whereas the out-of-plane (z) sensors have a negative bias and both have scale factors of 255 LSB/g. There does not appear to be a pattern in the biases of the BMA180 sensors but for all four sensors the z-axis is the least sensitive (greatest LSB/g)

Only one measurement range of ADXL345 was tested as a reading of the datasheet and previous experiments led the authors to believe that the other measurement ranges were just created by cropping the MSB (see Section 6.3).

Sensor		Drift (LSB) (first minus check samples)	1-sigma of estimate of mean (LSB)	Bias (LSB)	Bias (mg)*	Scale Factor (LSB / g)
Name	Range Setting					
ADXL1 x	16g	-0.001	0.0280	7.734	30.2	-261.38
ADXL1 y	16g	0.175	0.0316	7.805	30.5	264.09
ADXL1 z	16g	0.206	0.0592	-8.260	-32.3	255.29
ADXL2 x	16g	0.035	0.0264	9.374	36.6	-263.42
ADXL2 y	16g	0.141	0.0309	12.366	48.3	263.45
ADXL2 z	16g	-0.006	0.0579	-19.170	-74.9	-255.21
BMA1 x	16g	-0.048	0.0392	-92.466	-180.6	511.44
BMA1 y	16g	0.143	0.0443	-127.755	-249.5	520.34
BMA1 z	16g	-0.280	0.0420	11.542	22.5	526.21
BMA2 x	16g	-0.385	0.0380	27.528	53.8	-521.86
BMA2 y	16g	0.219	0.0485	-94.598	-184.8	520.67
BMA2 z	16g	0.057	0.0400	68.060	132.9	522.03
BMA3 x	1.5g	6.066	0.1929	55.081	10.1	5427.45
BMA3 y	1.5g	1.134	0.2458	-11.481	-2.1	-5414.31
BMA3 z	1.5g	-0.852	0.2406	-52.891	-9.7	5497.69
BMA4 x	1.5g	0.850	0.2221	54.229	9.9	-5369.79
BMA4 y	1.5g	-0.365	0.3313	-42.010	-7.7	-5401.06
BMA4 z	1.5g	-3.673	0.2712	148.693	27.2	5481.16

Table 3: Accelerometer summary statistics. *Conversion of bias into milli-g is approximate based on the ‘typical’ scale factors in [13, 14].

Sensor		Average Bias		Difference in bias between plus and minus (box frame)						Drift from first (+z) to final (also +z) run	
Name	Range Setting			X		Y		Z			
		LSB	dps*	LSB	dps*	LSB	dps*	LSB	dps*	LSB	dps*
Gyro1 x	2k dps	8.389	0.587	0.006	0.000	-0.545	-0.038	-0.073	-0.005	0.0331	0.002
Gyro1 y	2k dps	-7.026	-0.492	-0.160	-0.011	-0.146	-0.010	-0.023	-0.002	-0.0682	-0.005
Gyro1 z	2k dps	12.888	0.902	-0.094	-0.007	-0.009	-0.001	-0.202	-0.014	-0.1195	-0.008
Gyro2 x	2k dps	18.930	1.325	0.007	0.000	0.385	0.027	-0.501	-0.035	-0.3706	-0.026
Gyro2 y	2k dps	-4.050	-0.284	0.236	0.017	0.209	0.015	-0.260	-0.018	0.0229	0.002
Gyro2 z	2k dps	7.526	0.527	-0.221	-0.015	0.097	0.007	0.007	0.000	-0.2091	-0.015
Gyro3 x	250dps	-10.822	-0.095	-0.508	-0.004	-3.315	-0.029	0.133	0.001	0.0796	0.001
Gyro3 y	250dps	80.196	0.702	1.109	0.010	0.701	0.006	0.694	0.006	2.4205	0.021
Gyro3 z	250dps	-6.116	-0.054	-0.674	-0.006	-1.410	-0.012	-0.206	-0.002	-0.3297	-0.003

Table 4: Gyroscope summary statistics, those entries where the apparent g-dependent bias is smaller than the bias drift are greyed out, because they may be due to in-run bias variation. *The conversion to dps is based on the ‘typical’ scale factor.

For the two ranges of BMA180 tested the bias, in terms of LSB is of similar magnitude. This means in terms of physical quantities (ms^{-2} or g) the magnitude of the bias is very significantly lower.

As the sensors are mounted on breakout boards which are not mounted completely flat on the main PCB and the main PCB is rotated slightly (< 5 degrees) relative to the frame, we shall not present the azimuth and elevation angles of the sensitive axes.

5.2.2 Gyroscope Bias and G-dependence

Using the methodology implemented here the scale factors and axis orientations of the gyroscopes cannot be determined. As such the only systematic errors we measure here are gyroscope bias and g-dependence.

G-dependence is the angular rate error created by an applied specific force. Nearly all consumer-grade MEMS gyroscopes are vibratory and thus work by measuring the Coriolis force as such they are inherently sensitive to specific force, and this sensitivity must be compensated for in the design [1]. Often this is achieved by having a pair of vibrating masses and measuring the difference in force between them [10]. As a result of this, one might expect any g-dependencies to be due to manufacturing errors, and so not exhibit any obvious pattern.

The L3G4200D datasheet [12] specifies that the bias (“digital zero-rate level”) should be ± 10 dps for 250dps range and ± 70 dps for 2000dps. The systematic errors calculated are presented in Table 4. They use the experimental results presented in Table 10. The bias, which is calculated by taking the mean of the results of each of the six orientations, is clearly well within specifications.

To calculate the g-dependence the differences between the biases with the ‘+x’ and ‘-x’ faces are presented, noting that this represents a difference on the x-axis of $2g$ (19.6ms^{-2}). Similarly, the differences along the ‘y’ and ‘z’ axes are presented.

In order to check that the observed g-dependence is real and not an error resulting from the temperature variation of the bias, the difference between the first measurement and last measurement is presented in Table 4 as ‘drift’,

because the sensor orientation is the same. As it happens these are the two measurements with the largest difference in temperature (see Table 10 in the Appendix), thus they are the two that might be expected to show the greatest difference. For this reason one might assume that any observed difference greater than the ‘drift’ is a real g-dependent error.

Unlike some of the other sensors the Gyros are mounted so that their sensitive axes are roughly on the same named axis of the box reference frame (See Table 1), although the signs differ. For all three triads the ‘x’ angular rate is most affected by specific force in the ‘Y’ direction, the ‘y’ angular rate is most affected by the ‘X’ direction specific force and the ‘z’ angular rate has no apparent pattern, although it has the lowest magnitude of g-dependence.

5.2.3 Bias stability

To examine the bias stability, a 90 minute static experiment was conducted and then the data was split afterwards into 5 minute sections. The full results of this are presented in Table 11 in the appendix, on the right hand side. A summary is presented in Table 5. There was more variation in the z-axis sensors than the x and y for the two ADXL345s and the BMA180s that were at $\pm 1.5g$ setting, but not the other BMA180s or the gyroscopes.

Sensor name and axis (range)	Mean over the sensors and axes of the standard deviation of the 15 5-minute sample means (LSB)
ADXL X&Y	0.04125
ADXL Z	0.0945
BMA (16g) X&Y	0.1045
BMA (16g) Z	0.1065
BMA (1.5g) X&Y	0.9645
BMA (1.5g) Z	1.505
Gyro (2kdps) X&Y	0.098
Gyro (2kdps) Z	0.077
Gyro (250dps) X&Y	0.3745
Gyro (250dps) Z	0.442

Table 5: Summary results of the in-run bias variation

In comparing the high- and low-range BMA180s and gyroscopes the bias variation scales approximately as would be expected given the scale factor. An illustration of the in-run variation is in Figure 6.

These results make it clear that there is no point calibrating the units’ biases with a higher precision than 1LSB for the 1.5g accelerometers or 250dps gyroscopes or 0.1-0.2 LSB for all the other sensors.

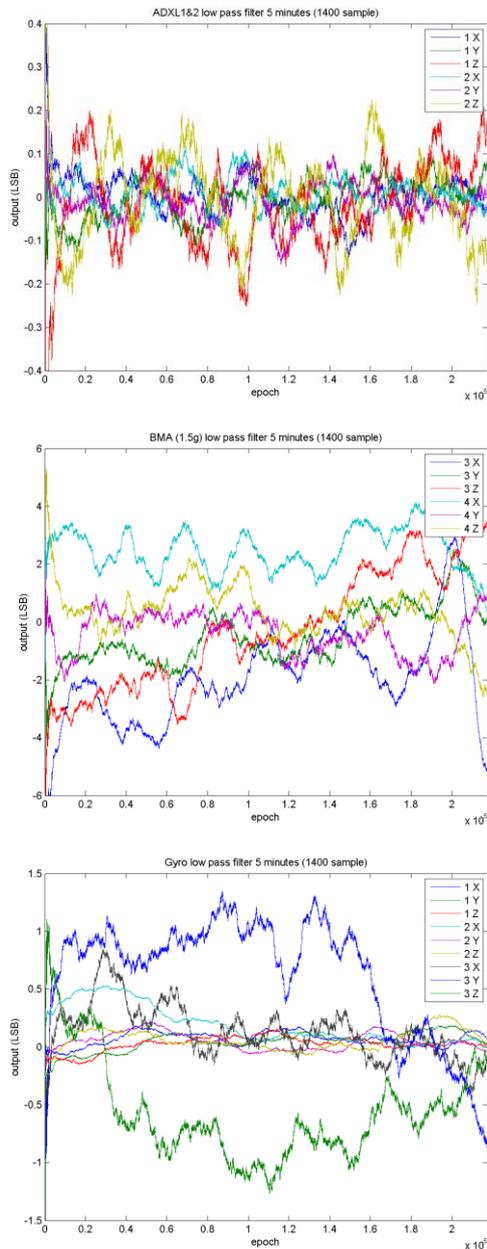


Figure 6: Static output minus mean value of the three types of sensors ADXL1&2 (top), BMA3&4 (middle), Gyro1&2&3 (bottom) smoothed over 1400 samples (5 minutes) to show the in-run bias variation.

5.2.4 Turn-on-bias variation

An experiment was run to determine the magnitude of the variation of the bias between runs. In this experiment 15 5-minute data samples were collected with a 2-4 minute gap between each sample, during which the test-bed was switched off. These tests were run sequentially over about 2 hours. This represents a situation when the run-to-run variation should be the lowest as the temperature is nearly the same and the actual time between samples is low, run-to-run variation might be higher if the readings were on separate days.

Table 11 in the appendix presents the results of the turn-on-bias determination experiments. BMA1 and Gyro3 are

greyed out and, for reasons given below, are not considered further in this analysis. Also presented are the in-run variations discussed in the previous section and the ratio between the in-run and between run “standard deviation of the 15 means”.

Each of the sensor designs will be considered separately. The ADXL accelerometers only show an inter-run variation (standard deviation) of approximately the same magnitude as the uncertainty in estimates of the means, so we could conclude that there is not a significant turn-on-bias, because it is also only 1.2 to 1.8 times the in-run variation. The BMA180 accelerometers show run-to-run standard deviation that is significant, often several LSB. This reinforces the idea presented in the last section that this should be viewed as a limit on the accuracy of any pre-run calibration.

The gyroscopes’ run-to-run variation is even more significant at 2.4 to 7.7 times the in-run variation. This is enough that it may be worth having an extra modelling term for this.

Two sensors exhibited unusual behaviour. The Bosch accelerometer known as “BMA1” had very consistent bias values except in a one of the samples (run 10), when both the specific force and temperature reading were markedly different, after which the values returned to the previous values. This is illustrated in Table 6. Additionally the third gyroscope, known as “Gyro3” appears to flip between two different sets of fixed biases (0, 69, 2) and (-13, 75, 0) for x, y and z (units are LSB). This is presented in Table 7, with one group highlighted in blue and the other in red. This behaviour might be due to communications errors in the start-up process, perhaps leading to the incorrect value being set for the measurement range. This would explain the unusual output for run 10 with BMA1 as this z-output would be expected if the measurement range was the default +/- 2g rather the +/- 16g that the firmware was meant to set.

These two sensors suggest that it might be worth having an algorithm that could store two possible values for the fixed bias and switch between the two of them depending on sensed output, or alternatively search for booting errors and re-initialise the sensors in that case.

Run	x-output (LSB)	y- output (LSB)	z- output (LSB)	temp int
1	-100.491	-134.406	536.7756	-4.43686
2	-100.437	-134.349	536.8192	-4.80532
3	-100.125	-134.394	536.5554	-5.18888
4	-100.058	-134.185	536.5377	-5.2798
5	-100.037	-134.197	536.525	-5.21758
6	-100.115	-134.337	536.405	-5.2359
7	-99.8368	-134.134	536.5055	-5.248
8	-100.026	-134.119	536.3266	-5.12513
9	-99.9712	-134.064	536.5216	-5.09975
10	31.79535	-124.464	4308.719	-3.14497
11	-100.104	-134.096	536.4102	-5.01496
12	-100.131	-134.173	536.2859	-4.9306
13	-100.076	-134.415	536.2423	-4.90121
14	-99.9396	-134.165	536.1224	-4.93264
15	-99.8322	-134.406	536.3374	-4.7673

Table 6: Mean 3-axis outputs of “BMA1” over each of the 15 runs. Contrast run 10 (in red) with the others.

Run	x- output (LSB)	y- output (LSB)	z- output (LSB)	temp int
1	0.088086	71.83647	0.259345	21.19651
2	0.247328	71.01841	1.897628	21.42149
3	-12.9493	77.17322	0.23046	21.91818
4	-13.1201	76.19548	0.591934	21.97628
5	-12.1075	74.64321	0.741931	21.87275
6	-0.33707	69.68126	2.61271	21.90305
7	-13.0202	75.05754	1.725069	21.99339
8	-0.34012	69.09461	2.625957	21.89698
9	0.07932	69.01431	3.153997	21.92783
10	-12.2896	73.89692	0.454047	21.88836
11	-12.2955	74.03204	0.621792	21.8028
12	0.04686	69.56259	2.379449	21.71733
13	-12.0067	74.54049	0.117055	21.82495
14	-0.39189	69.6091	2.201399	21.82179
15	-13.0925	75.8296	1.496677	21.77527

Table 7: Mean outputs of “Gyro3” over each of 15 5-minute runs. Note the two groups of output biases, coloured blue and red.

5.2.5 Temperature variation of bias

In order to determine the effect of temperature on the sensors’ biases, an experiment was conducted to determine the effect of heating on the various MEMS inertial sensors tested. The temperature profile of the experiment is presented in Figure 7.

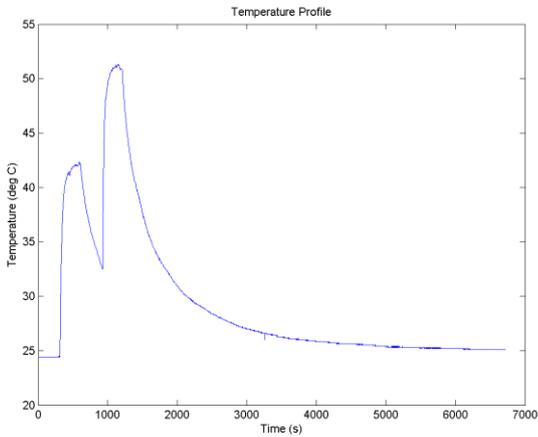


Figure 7: The temperature profile provided by the BMP180 sensor for the experiment.

As discussed in Section 4, the sensors were left to run at room-temperature for 5 minutes (around 24°C), they were then heated by the hairdryer at medium-heat for 5 minutes (reaching around 42-43°C), then left to cool for a further 5 minutes then heated on high heat for 5 minutes (reaching around 52°C) followed by being left to cool for approximately 90 minutes, not quite returning to the original temperature possibly due to a change in the ambient temperature.

Several of the sensors have their own internal temperature sensors (the BMA180s and the L3G4200D gyroscopes, but not the ADXL345s). However, for this experiment we have used the temperature output of the independent temperature/pressure sensor (BMP180). This is to keep a consistent reference between all the sensors, and allow the higher precision and accuracy reading of temperature to be used.

The results of a linear fit between the sensor output and temperature are presented in Table 8. Higher order (e.g. quadratic) fits were also tried however this did not significantly improve the fit, and so there would be little justification modelling the higher order temperature component of the biases. Examples of how these coefficients fit to the data are presented in Figures 8 and 9, as can be seen there is a significant bias variation with temperature particularly for the gyroscopes.

Name	Range	Slope (LSB/°C)	Intercept (at 0 °C) (LSB)
ADXL1 x	16g	-0.024	12.99
ADXL1 y	16g	0.019	21.90
ADXL1 z	16g	-0.352	256.53
ADXL2 x	16g	-0.011	7.42
ADXL2 y	16g	-0.067	16.21
ADXL2 z	16g	-0.194	241.49
BMA1 x	16g	-0.034	-99.87
BMA1 y	16g	-0.199	-129.67
BMA1 z	16g	0.192	532.54
BMA2 x	16g	0.205	47.07
BMA2 y	16g	-0.258	-79.90
BMA2 z	16g	0.496	-463.79
BMA3 x	1.5g	-0.548	530.01
BMA3 y	1.5g	-0.121	-4.28
BMA3 z	1.5g	4.516	-5635.33
BMA4 x	1.5g	0.946	289.72
BMA4 y	1.5g	-1.714	-12.27
BMA4 z	1.5g	-2.235	5656.97
Gyro1 x	2k dps	0.235	1.79
Gyro1 y	2k dps	0.265	-13.71
Gyro1 z	2k dps	0.221	7.30
Gyro2 x	2k dps	0.521	5.67
Gyro2 y	2k dps	0.261	-11.83
Gyro2 z	2k dps	0.187	2.50
Gyro3 x	250dps	0.904	-23.05
Gyro3 y	250dps	2.133	14.69
Gyro3 z	250dps	-2.782	69.02

Table 8: The results of a linear best fit between sensor output and temperature (derived from BMP180).

The ADXL345 accelerometers show only a very small temperature drift for the in-plane sensors (with three out of four being slightly negative) and a much more significant (5-10 times greater) negative temperature drift for the out-of-plane sensor.

The BMA180 accelerometers do not show an obvious pattern in their temperature-drift behaviour. In general, the z-axis sensor shows the greatest variation (and 3 of 4 times this is positive), but as the reaction to gravity is acting on this sensor, this may be also influenced by a change of sensitivity with temperature. The slope of the temperature variation in LSB/°C is around 10x steeper for the 1.5g sensors than the 16g sensors so the range setting does not seem to affect the magnitude of the temperature drift in physical units significantly.

In 8 out of 9 cases, the gyroscope bias temperature drift is positive. This means that the opposite-direction sensitive axis combination (Section 2.1) may be feasible. For both 2000dps gyroscopes (Gyro1 and Gyro2) the drift is always positive and in every case but one 1.9-2.6 LSB/°C. If we were to pick the two y-axis sensors which are in opposing directions already (See Table 1) we would have a combined bias drift of 0.004 LSB/°C. This is illustrated in Figure 10.

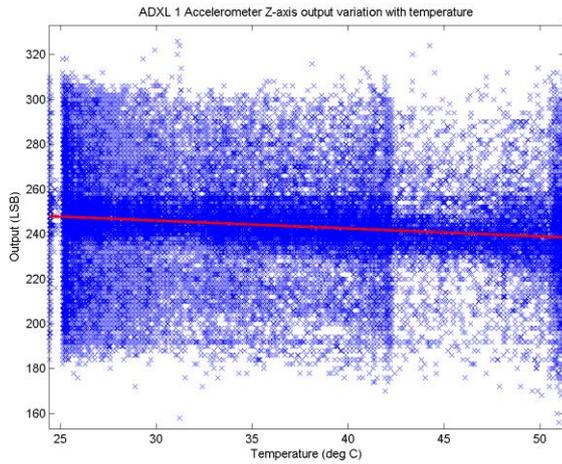


Figure 8: Scatter plot of ADXL1 Z-axis output versus temperature. Best fit line in red.

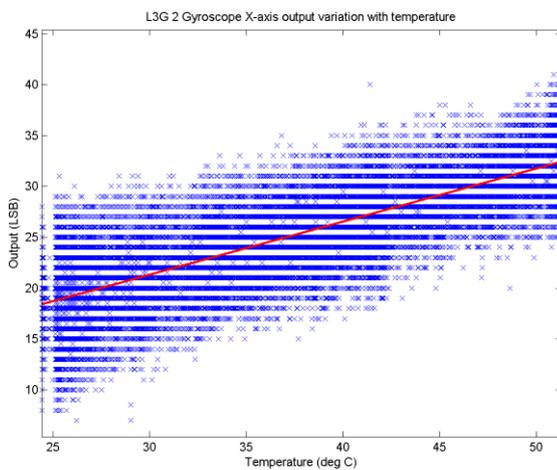


Figure 9: Scatter plot of Gyro2 X-axis output versus temperature. Best fit line in red.

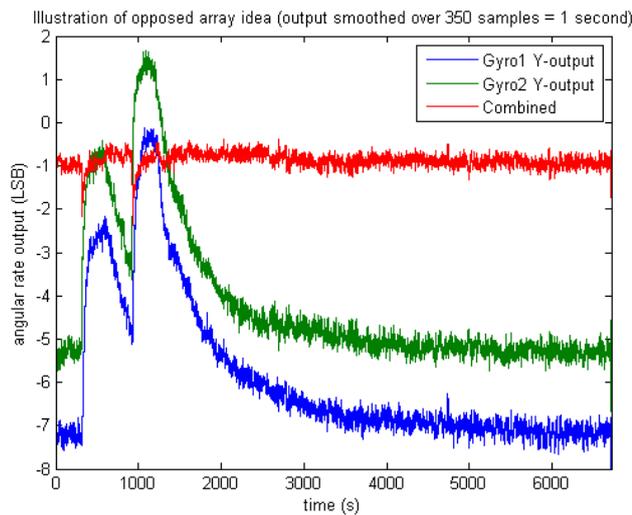


Figure 10: Comparison of y-output of Gyro1, Gyro2 and half of the difference between the outputs.

6. FEASIBILITY OF ADVANCED SENSOR ARRAY TECHNIQUES

In Section 2 we proposed three techniques to improve the navigation performance of a low-cost IMU using an array

of multiple triads of MEMS inertial sensors. We performed characterisation tests on three specific models of low-cost MEMS inertial sensors namely Analogue Devices ADXL345 accelerometers, Bosch BMA180 accelerometers and ST Microtronics L3G4200D gyroscopes, the results of which were presented in Section 5. The tests were intended to determine which of the ideas presented in Section 2 would work for those three sensors. The rest of this section is presented idea-by-idea.

6.1 Common mode errors of sensors of the same design

In Section 2.1 we presented the idea of using a pair of oppositely orientated sensors of the same type to mitigate an error coming from the environment. This builds on the idea, presented by Yuksel et al. [11] to mitigate the effect of temperature on gyroscope bias

In this paper we investigate the feasibility of this technique for different sensors and for other errors. This technique would theoretically work for other correlated even-order errors. For example if the fixed bias tended to always be positive then this technique would reduce the total bias relative to the alternative of combining sensors with their sensitive axes in the same direction, see Figure 2.

The experiments we carried out were able to show that this idea would work for some of the errors on some of the sensors, as we expected.

The gyroscopes' temperature drifts, despite the fact that they are specified as being ± 0.04 dps/ $^{\circ}$ C was in 8 out of 9 cases positive, meaning that this idea would likely apply to this model of gyroscope. Additionally both the ADXL345 z-axis sensors showed significant negative temperature drift. The other axes of the ADXL345 showed much less significant temperature drift. The BMA180 showed significant temperature drift but there was no apparent pattern to whether it was positive or negative, so this idea would not help.

6.2 Different characteristics of in-plane and out-of-plane sensors of the same triad

In Section 2.2, we observed that as on many inertial sensor triads the in-plane and out-of-plane sensors were of a different design, this could lead to differences in behaviour that might impact navigation performance.

The gyroscope's z sensor apparently has very slightly better performance, but not significantly. Each BMA180 has a turn on bias is about twice as high for the z-axis sensors (see Table 11).

The ADXL345 sensor has the most significant difference between the in- and out-of-plane sensors. The standard deviation of the noise is approximately double that of the x- and y-sensors (about 12 LSB rather than about 6 LSB). The z-sensor also has much larger temperature drift and significantly lower bias stability.

We recommend mounting the BMA180s with the z-axis vertical because the vertical accelerometer biases are easier to observe using integration, alignment and zero update algorithms than the horizontal biases [1].

In the case of the ADXL345s the best performance might be achieved by not using the z-axis sensors at all. This might be by putting a triad on each of two perpendicular PCBs or arranging three triads on three faces of a cube so that there are two in-plane sensors facing in each direction. The z-axis sensors could perhaps still be used for fault detection and integrity monitoring.

6.3 Multiple sensors with different measurement ranges

In Section 2.3, we presented the idea of combining sensors with different measurement ranges to increase navigation performance by weighting the outputs of the two different range sensors according to the sensed dynamics. We proposed using mostly a low-range sensor when the dynamics of the system are low, and when the dynamics are higher using mostly the higher range sensor.

However the one major assumption that we had to make was that if we had two sensors of the same quality but different measurement ranges the one with the lower range would be more accurate. We sought to test this assumption through the experiments described in Section 4 and analysed in Section 5.

We did not test the ADXL345 at different dynamic ranges as it does not appear to be an equivalent sensor at all the dynamic ranges as it has a fixed conversion from LSB to milli-g (4mg/LSB) for all measurement ranges.

We tested the L3G4200D gyroscope at 250 and 2000 degrees sec⁻¹ (dps) measurement ranges. The noise level, was not affected by the measurement range setting at all. Of the systematic errors measured only the fixed bias showed any significant improvement. However this is only helpful in the case that no calibration is carried out at all, which is unusual as gyroscope bias calibration is very simple to do, simply placing the IMU still on a table. Thus we do not expect any significant performance improvements from combining L3G4200D gyroscopes of different measurement ranges. However, this approach may still be viable for other types of gyroscope.

We tested the BMA180 accelerometer at 1.5g and 16g. The temperature drift was slightly lower in physical units, similarly the biases were slightly more stable. However the fixed bias was significantly lower in physical units. Most importantly the noise was much lower, the standard deviation being an average of around 50 LSB rather than around 10 LSB (see Table 10), which means that the standard deviation of the low-range sensors was half as much as the high range sensor. Thus, significant performance benefit could be achieved through the weighted combination of low-range and high-range sensors. The optimal pair of ranges for a specific

application would depend on how much the noise performance varies between the 7 possible settings

7. CONCLUSIONS

This paper has presented three advanced array-based techniques that could be applied to improve the performance of low-cost MEMS IMUs. They are expected to provide a more accurate estimate of specific force and/or angular rate than a simple average of the array's constituent sensors. They are:

1. Arranging the sensors so that their sensitive axes are in opposing directions, so as to significantly reduce the effect of systematic errors that are correlated across sensors of the same design.
2. Exploiting the performance differences between the in-plane and out-of-plane sensors on a sensor triad.
3. Combining the output of sensors with different operating ranges to increase the accuracy of the measurement during periods of relatively low dynamics without clipping and distorting under high dynamics.

In order to examine the feasibility of these ideas we constructed a hardware platform to characterise the performance of three models of MEMS inertial sensors: Bosch BMA180 accelerometers, Analogue Devices ADXL345 accelerometers and ST Microtronics L3G4200D gyroscopes, examining both the systematic and random errors.

The sensors' characteristics led to the conclusion that for all three sensor level ideas proposed at least one of the three sensors tested could potentially benefit. The L3G4200D gyroscopes could benefit from the opposing direction sensitive axis arrangement. The ADXL345 could benefit from a proper consideration of the differences between its in-plane and out-of-plane sensors. The BMA180 could show considerable performance benefit from a weighted combination of low- and high-range sensors.

8. FUTURE WORK

The next stage of this research is to implement each of the proposed sensor array techniques with the appropriate sensor type and quantify their performance impacts. In addition, testing will be extended to dynamic conditions, using a manufacturer-calibrated IMU to provide a reference.

Another way of improving the navigation performance of an integrated navigation system which includes low-cost inertial sensors is to have the end-user conduct a calibration routine, similar to that conducted by the manufacturer for more expensive sensors. As low-cost MEMS sensors are supplied essentially without any calibration, even an approximate calibration would be a big improvement and allow more precise estimation to be made by an integrated navigation system much more quickly.

There are two main requirements for a user-conducted calibration-on-purchase. First, it is essential that the physical movements required of the sensor are very simple and easily understood and completed, even if the underlying method is complex. Second no sophisticated or expensive equipment should be required, for instance a reference IMU is out of the question.

We have shown that by simply mounting the IMU in a cuboid box and then taking readings on each face, we can make an estimate of the many of the systematic errors of the accelerometers and the gyroscope biases and g-dependant errors. However it is desirable to use a technique that offers both greater flexibility in terms of IMU manoeuvring and calibrates additional errors such as gyroscope scale factor and cross-coupling errors. Kalman-filter based alignment, sensor integration and zero-update algorithms [1] already offer these capabilities. However they are designed to estimate the turn-on and in-run contributions to systematic errors that have already been calibrated by the manufacturer. Applying these algorithms to raw sensors could lead to large linearization errors. Therefore further research will be conducted to determine the best approach.

In the long term, the sensor array techniques proposed here could form part of a new generation of multi-sensor integrated navigation system alongside techniques such as GNSS shadow matching [22], environmental feature matching [23], opportunistic radio navigation [6] and context adaptivity [7].

ACKNOWLEDGEMENTS

Henry Martin is jointly funded by the Engineering and Physical Sciences Research Council (EPSRC) and BAE Systems Advanced Technology Centre.

The authors would like to thank Santosh Bhattari (of UCL) and John Davis (of the National Physical Laboratory) for their help interpreting the Allan variance graphs. The authors would like to thank Tuula Eriksson (UCL) for the loan of the hairdryer and Ian Seaton, Warren Gaynor and Les Irwin (all UCL) for their help with the hardware development.

REFERENCES

[1] Groves, P. D., Principles of GNSS, inertial, and multi-sensor integrated navigation systems, Second Edition, Artech House, 2013.

[2] Gustafson, D., J. Dowdle, and K. Flueckiger, "A Deeply Integrated Adaptive GPS-Based Navigator with Extended Range Code Tracking," Proc. IEEE PLANS 2000.

[3] Groves, P. D., C. J. Mather and A. A. Macaulay, "Demonstration of Non-Coherent Deep INS/GPS Integration for Optimized Signal to Noise Performance," Proc. ION GNSS 2007.

[4] Ma, Y., W. Soehren, W. Hawkinson, and J. Syrstad, "An Enhanced Prototype Personal Inertial Navigation System," Proc. ION GNSS 2012.

[5] Groves, P. D., et al., "Inertial Navigation Versus Pedestrian Dead Reckoning: Optimizing the Integration," Proc. ION GNSS 2007.

[6] Faragher, R., Sarno, C., and Newman, M. "Opportunistic Radio SLAM for Indoor Navigation using Smartphone Sensors," Proc. IEEE/ION PLANS 2012. pp 120-128

[7] Groves, P. D., Martin, H., Voutsis, K., Walter, D. & Wang, L. Context Detection, Categorization and Connectivity for Advanced Adaptive Integrated Navigation. ION GNSS+ 2013. Nashville, Tennessee.

[8] Bancroft, J. B., and G. Lachapelle, "Data fusion algorithms for multiple inertial measurement units," Sensors, Vol. 11, No. 7, 2011, pp. 6771-6798.

[9] M. Becker, U. Bestmann, A. Schwithal, P. Hecker R. Bott, P. Kohl, S. Knedlik, H. Scheyer, E. v. Hinüber, "Evaluation of a double equipped MEMS IMU based on real flight trial scenarios" 2010 IEEE pp.202-213

[10] Kempe, V. Inertial MEMS: Principles and Practice. Cambridge University Press, 2011.

[11] Yuksel, Y., N. El-Sheimy, N., and A. Noureldin, "Error modelling and characterization of environmental effects for low cost inertial MEMS units," Proc. IEEE/ION PLANS 2010 , pp. 598-612.

[12] STMicroelectronics. L3G4200D MEMS motion sensor: ultra-stable three-axis digital output gyroscope datasheet, Doc ID 17116 Rev. 3 ed., December 2010.

[13] Analog Devices. ADXL345 3-Axis Digital Accelerometer Datasheet, rev C ed. One Technology Way, P.O. Box 9106, Norwood, MA 02062-9106, U.S.A., 2009-2011.

[14] Bosch Sensortec, BMA180 digital, triaxial acceleration sensor data sheet, v2.5, December 2010

[15] Titterton, D., and Weston, J. Strapdown Inertial Navigation Technology. Institution of Electrical Engineers 2004.

[16] Bosch Sensortec, BMP180 digital pressure datasheet, v.2.4, January 2012

[17] Honeywell. 3-Axis Digital Compass IC HMC5883 datasheet, Form 900405 Rev A, 12001 Highway 55, Plymouth, MN 55441, March 2010

[18] UM10204 I2C-bus specification and user manual Rev. 03, June 2007

[19] Allan, D. W. Statistics of atomic frequency standards. Proceedings of the IEEE, 1966.

[20] IEEE Standard Specification Format Guide and Test Procedure for Single Axis Interferometric Fiber Optic Gyros, IEEE Std 952-1997 Annex C, 02/1998.

[21] Le Traon, O. et al. "The VIA Vibrating Beam Accelerometer: Concept and Performances" Proc. IEEE/ION PLANS, Palm Springs, CA, 1998, pp. 25-37

[22] Wang, L., P. D. Groves, and M. K. Ziebart, "Urban Positioning on a Smartphone: Real-time Shadow Matching Using GNSS and 3D City Models," Proc. ION GNSS+ 2013.

[23] Walter, D. J., P. D. Groves, R. J. Mason, J. Harrison, J. Woodward, and P. Wright, "Novel Environmental Features for Robust Multisensor Navigation," Proc. ION GNSS+ 2013.

APPENDIX A

In this appendix more detailed analysis of the sensors stochastic errors are presented. The stochastic analysis of each sensor model examined in separate sub-sections.

A.1 L3G4200D Gyroscope

The data from two long (~1400 sec) high rate (~350Hz) static experiments was used to make a first assessment of the Gyroscope’s noise performance. The registry settings were set for an output data rate (ODR) of 400Hz and cut-off of 110Hz. For one of these experiments the measurement range was set to 2000 degrees per second (dps) and the other to 250dps. After converting the raw output from LSB to dps (according to the typical specification, i.e. not calibrated), the six outputs (3 axes 2 experiments) were plotted as Allan deviation curves [19, 20], presented in Figure 11.

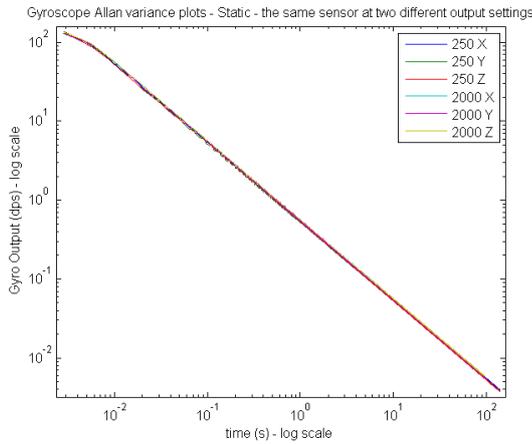


Figure 11: Allan Deviation curves for the 3 axes of output for two static experiments with “Gyro1” one with 250dps measurement range and one with 2000dps.

In Figure 11, it can be seen that the output of the sensor is close to white across most of time interval sampled as it resembles a straight line of slope -1, and that the noise has approximately the same order of magnitude in dps for all axes and both measurement ranges, and examining Table 9 one can see that there is no significant difference in the standard deviations of the output so the noise level is the same for 250dps and 2000dps measurement range. However, it differs in two important ways, at the two points where an enlarged view of the graph is presented in Figure 12. First for the shortest time intervals the slope slackens off suggesting some memory in the process, and also there is some periodic behaviour visible on both Y-axis signals.

In order to further investigate the possible memory in the process we compute an autocorrelation, presented in Figure 13, for the 250dps measurement range. The 2000dps range is very similar, but not shown. It can be seen that there is a positive correlation of around 0.4 for all three axes between one sample and the following sample (lag=1), this implies that the actual bandwidth of the sensor is lower than the ~200Hz Nyquist rate. Also the periodic component in the Y-axis can be clearly seen, at around 6 samples.

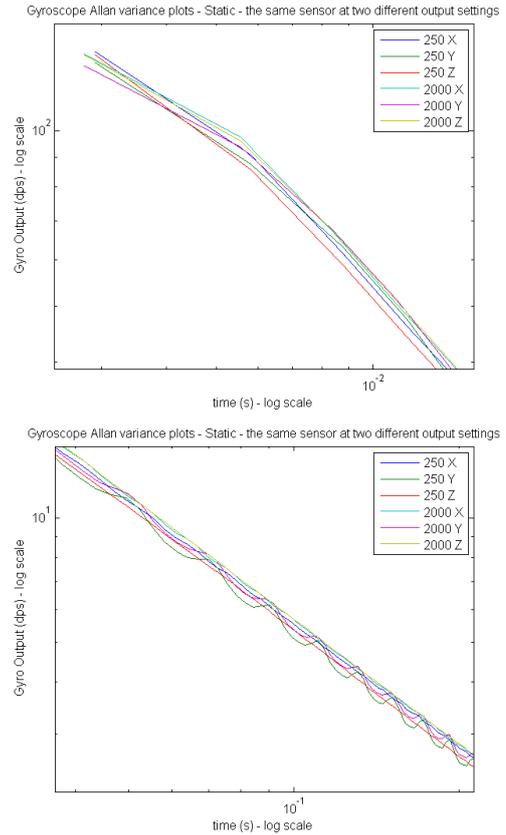


Figure 12: Zoomed in view of two parts of Figure 11.

Output	Range	Mean (Bias)	Standard error of the mean	Mean (Bias)	Standard error of the mean	Standard of the data	
		Output integers/LSB		Approx* conv. to DPS		Output LSB	DPS*
Gyro1 x	250dps	14.174	0.0530	0.1240	0.00046	36.816	0.3221
Gyro1 y	250dps	-17.607	0.0511	-0.1541	0.00045	35.512	0.3107
Gyro1 z	250dps	-7.644	0.0507	-0.0669	0.00044	35.193	0.3079
Gyro1 t	int	24.008	0.0006			0.398	
Gyro1 x	2k dps	6.386	0.0079	0.4470	0.00055	4.654	0.3258
Gyro1 y	2k dps	-6.789	0.0077	-0.4752	0.00054	4.530	0.3171
Gyro1 z	2k dps	12.695	0.0080	0.8887	0.00056	4.683	0.3278
Gyro1 t	int	24.083	0.0006			0.347	

Table 9: Summary Statistics for the static high-rate gyro experiments. The approximate conversion to dps (denoted by *) indicates using the “typical” conversion from the data sheet.

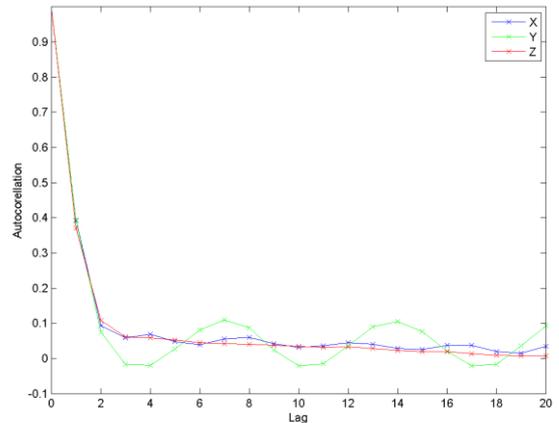


Figure 13: the autocorrelation of a L3G4200D gyroscope at 250dps measurement range.

To show the periodic behaviour a spectral density plot was calculated for the 3 axes of the gyroscope's output. This is presented in Figure 14, for the 250dps measurement range. The 2000dps range is again very similar, but not shown. Both ranges show a distinct peak at ~50Hz for the y-axis only. The experiment will be repeated to determine whether this is a genuine artefact of the sensor or mains hum picked up from elsewhere in the room.

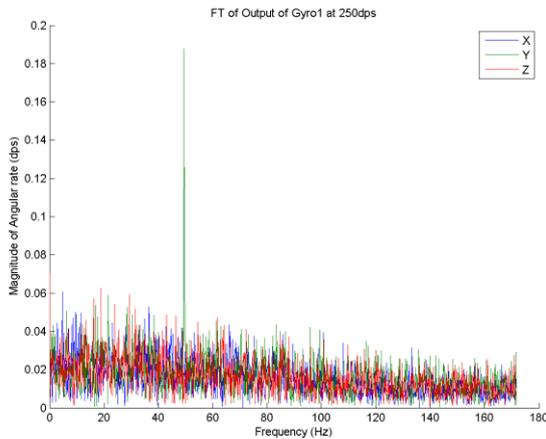


Figure 14: A fast Fourier transform of the output of a L3G4200D Gyroscope at 250dps.

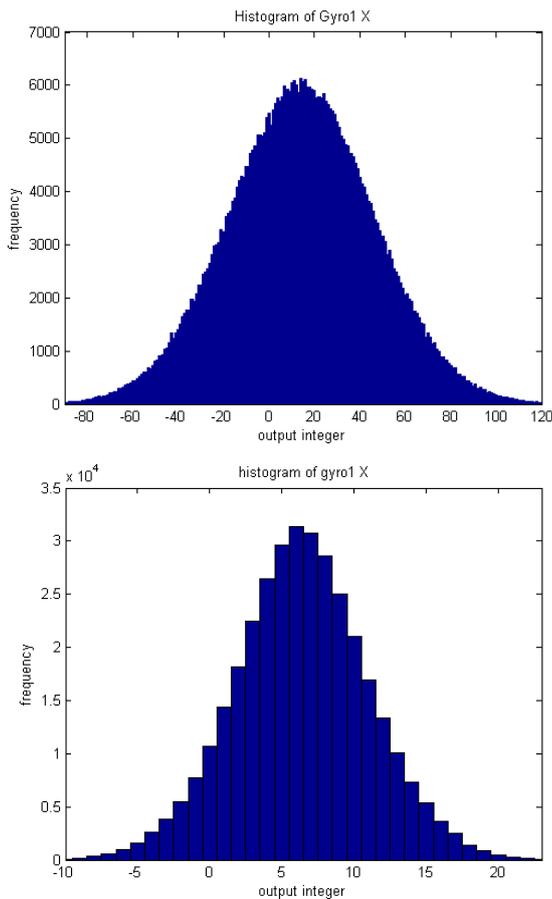


Figure 15: Histogram of Gyro1 X-output showing the approximately Gaussian distribution of the sensor noise. Bin size is 1 LSB. The 250dps range is above and 2000dps range below

In Figure 15 a histogram of the gyroscopes output is shown, one can observe two significant things from this. One, the noise is distributed approximately normally. Two, the quantisation level is fairly insignificant compared to the noise, which explains why the more sensitive range setting does not have any noticeable impact on the noise level.

A.2 ADXL345 Accelerometer

To characterise the stochastic errors of the ADXL345 accelerometer we took a long high-rate data sample from the sensor. The Allan variance of this is plotted in Figure 16. It can be observed that the lines for each axis straight throughout the sample, implying white noise. Additionally the z-axis noise can be observed to be of a considerably higher magnitude than the other axes. These inferences are supported by the sample's spectral density, in Figure 17. The specification for noise density is 1.1 LSB rms (for z) and 0.75 LSB rms (x and y) [13]. So the sensor appears to be out of specification for z, but strictly the specification is for 100Hz output data rate, not the 400Hz used here.

As the sensor noise is essentially white it is reasonable to characterise it by standard deviation. The standard deviation has been shown in multiple experiments (see Table 10 and 11) to be approximately 6 LSB for the X and Y sensors and 12 LSB for the Z.

According to [13], the ADXL345 accelerometer, at the 'full resolution' setting, has a fixed scale factor of 4mg/LSB for all measurement ranges. This means that while it is a 13-bit sensor at 16g it is only a 10-bit sensor at 2g. As we are not comparing like with like, we did not test it at multiple measurement ranges. It is, apparently, just cropping the MSB to achieve programmable dynamic ranges. Additionally the sensor, throughout this and all the other tests only outputs even numbers. This suggests that the sensor actually has a 12-bit ADC rather than the 13-bit ADC specified.

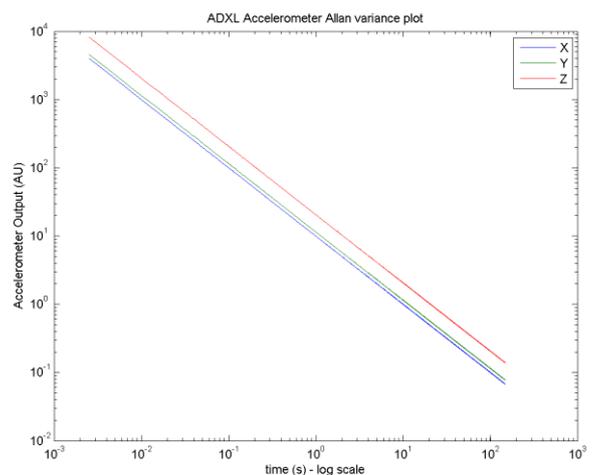


Figure 16: Allan deviation of ADXL345 accelerometer. The higher noise on the z-axis sensor is clearly visible.

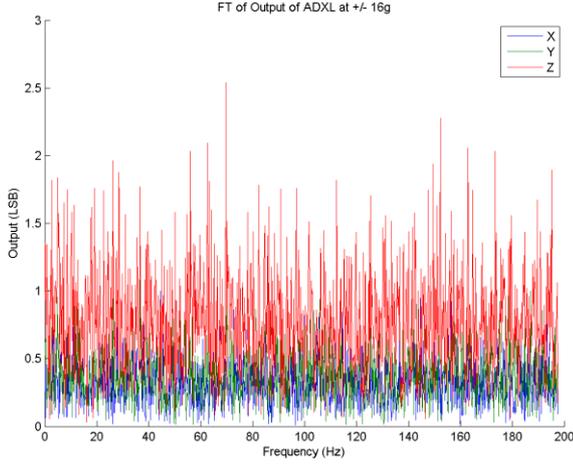


Figure 17: Spectral Density of ADXL345 accelerometer.

A.3 BMA180 Accelerometer

A long static experiment was also conducted for a BMA180 accelerometer to characterise its stochastic noise.

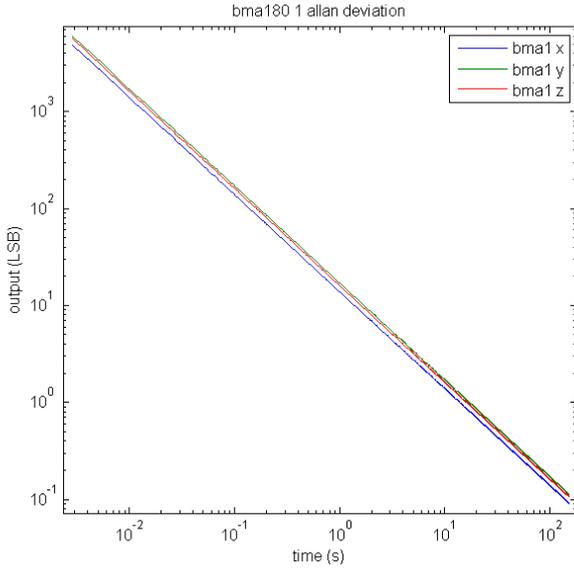


Figure 18: Allan deviation of BMA180 accelerometer.

The Allan deviation of the 3 axes of BMA180 output is shown in Figure 18. This shows that the noise is approximately white for the whole time scale examined. No significant resonances (see Figure 19) or significant autocorrelation are present.

The standard deviations for the x, y and z-axes, in LSB, are 8.0201, 9.8857 and 9.3002, respectively (15.6642, 19.3081 and 18.1645 milli-g).

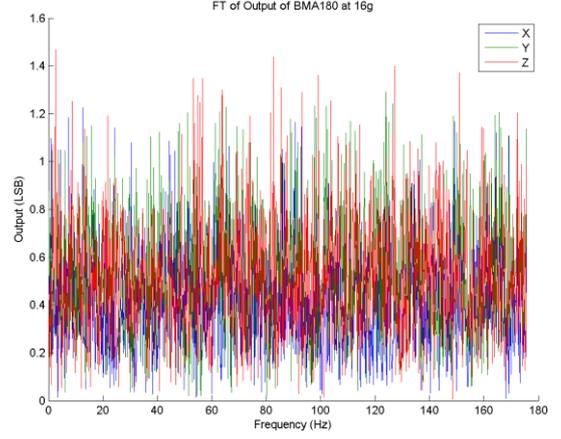


Figure 19: Spectral density of a BMA180 Accelerometer.

APPENDIX B

An outline of the method used to calculate the accelerometers' systematic errors is presented below. Consider average specific force readings taken statically from opposite faces of the cube where f_x^{yup} refers to the specific force reading of the x-axis sensor taken with the '+y' face uppermost, in each the specific force due to gravity is applying on the sensor in opposite directions. Thus an estimate of the bias can be calculated as:

$$bias_x = \frac{1}{2}(f_x^{xup} + f_x^{xdown})$$

Or considering all the six readings a more accurate estimate is,

$$bias_x = \frac{1}{6}(f_x^{xup} + f_x^{xdown} + f_x^{yup} + f_x^{ydown} + f_x^{zup} + f_x^{zdown})$$

Then remove this bias estimate from the six readings, e.g.,

$$f_{x_{bf}}^{xup} = f_x^{xup} - bias_x$$

To improve accuracy average the 'up' and 'down' for each axis, e.g.,

$$f_x^{xa} = \frac{1}{2}(f_{x_{bf}}^{xup} - f_{x_{bf}}^{xdown})$$

Then the azimuth of x, Ψ_x . This is the angle between the x-axis and the projection the sensitive axis onto the XY-plane in the direction so that the Y-axis is 90 degrees. This is calculated by

$$\Psi_x = \arctan\left(\frac{f_x^{xa}}{f_x^{ya}}\right)$$

And, the elevation of x, θ_x , defined as the angle between the XY plane and the sensitive axis, positive so that the z-axis is 90 degrees.

$$\theta_x = \arctan\left(\frac{2f_x^{za}}{f_x^{xa}/\cos(az_x) + f_x^{ya}/\sin(az_x)}\right)$$

Then the scale factor could be calculated 6 ways, and thus the average of these six is used, so,

$$SF_x = \frac{1}{6} \left(\frac{f_{x_{bf}}^{x_{up}}}{\cos(\theta_x)} + \frac{f_{x_{bf}}^{y_{up}}}{\cos(\theta_x) * \sin(\Psi_x)} + \frac{f_{x_{bf}}^{z_{up}}}{\sin(\theta_x) * \cos(\Psi_x)} + \frac{-f_{x_{bf}}^{x_{down}}}{\cos(\theta_x) * \cos(\Psi_x)} + \frac{-f_{x_{bf}}^{y_{down}}}{\cos(\theta_x) * \sin(\Psi_x)} + \frac{-f_{x_{bf}}^{z_{down}}}{\sin(\theta_x)} \right)$$

Alternatively, the bias-free estimates could be used to calculate a 3x3 cross-coupling and scale factor matrix, see [1, p160].

In a working system the transformation matrix would be preferred for computational simplicity. However, in order to compare to the specification the “scale factor and sensitive axis” realisation is required.

APPENDIX C

This appendix contains the full results of the six-position cube experiment (Table 10) and the turn-on-bias experiment (Table 11).

Sensor		Mean Outputs for each cube face (LSB)							SD/root(n) average uncertainty in estimate of mean outputs	Average Standard Deviation of data
Name	Range Setting	Z up	X up	Y up	X down	Y down	Z down	Z up (check)		
Indep. Temp	°C	26.23803	26.45391	26.47834	26.42805	26.49729	26.49944	26.68422	0.001211	0.251412
ADXL1 x	16g	12.33293	-253.497	2.469027	269.075	13.12513	2.899611	12.33366	0.027971	5.811277
ADXL1 y	16g	22.35989	13.01353	-256.075	2.25853	271.2409	-5.9701	22.18462	0.031592	6.560084
ADXL1 z	16g	247.2157	-3.95661	4.159566	-11.9417	-22.4309	-262.606	247.0093	0.059165	12.28892
ADXL2 x	16g	7.175946	-253.973	5.220675	272.7924	13.62696	11.40376	7.141308	0.026397	5.481664
ADXL2 y	16g	14.30768	17.24602	-251.402	8.269124	275.401	10.37481	14.16685	0.030856	6.40681
ADXL2 z	16g	236.3945	-22.7733	-18.1663	-15.774	-20.7224	-273.978	236.4002	0.057945	12.03305
BMA1 x	16g	-100.816	-81.7408	-604.112	-102.087	418.4423	-84.4824	-100.769	0.039215	8.14338
BMA1 y	16g	-134.907	391.6477	-115.322	-648.648	-139.155	-120.145	-135.05	0.044294	9.199038
BMA1 z	16g	537.0402	13.09122	7.244381	10.84351	16.37739	-515.345	537.3207	0.042002	8.723506
BMA1 t	int	-1.652	-0.8761	-0.83958	-0.69925	-0.72451	-0.87959	-0.90098	0.005439	1.130422
BMA2 x	16g	52.25275	21.47985	549.6836	31.50906	-492.828	3.071563	52.63792	0.038041	7.901824
BMA2 y	16g	-86.6883	425.7175	-91.8819	-615.501	-97.1331	-102.099	-86.9078	0.048531	10.07948
BMA2 z	16g	-451.826	69.95953	90.44258	64.45256	44.14392	591.1872	-451.883	0.040044	8.315939
BMA2 t	int	5.789189	5.91244	6.021608	6.343173	6.156546	5.796374	6.463376	0.006815	1.414752
BMA3 x	1.5g	515.6351	144.0243	-5351.68	-24.0278	5461.916	-415.383	509.5691	0.192855	40.04875
BMA3 y	1.5g	-9.89211	-5422.59	-77.3823	5405.179	57.58821	-21.7895	-11.0259	0.245779	51.06081
BMA3 z	1.5g	-5522.18	20.71929	-451.205	-138.035	329.0204	5444.336	-5521.33	0.240589	49.97539
BMA3 t	int	3.311734	3.746916	3.574619	3.288311	3.498656	3.695056	3.938489	0.003748	0.778485
BMA4 x	1.5g	312.1222	-39.5078	5419.983	136.5785	-5306.03	-197.77	311.2725	0.222086	46.12221
BMA4 y	1.5g	-66.5468	-5438.93	-120.47	5361.847	47.06203	-35.0201	-66.1819	0.331269	68.80461
BMA4 z	1.5g	5603.323	183.5848	-161.923	129.4264	477.9277	-5340.18	5606.997	0.271194	56.33199
BMA4 t	int	4.636847	4.929013	4.64173	5.160408	5.017566	4.923241	5.331828	0.006028	1.252147
Gyro1 x	2k dps	8.320732	8.428642	8.161185	8.422897	8.705721	8.394012	8.287603	0.012642	2.626539
Gyro1 y	2k dps	-7.01099	-7.1224	-7.14924	-6.9622	-7.00326	-6.98778	-6.94282	0.012121	2.517964
Gyro1 z	2k dps	12.79765	12.80977	12.88774	12.90363	12.8971	13.00009	12.91716	0.013115	2.72343
Gyro1 t	int	21.08948	20.99158	21.0095	20.53514	20.72901	20.72739	20.60719	0.001901	0.393757
Gyro2 x	2k dps	18.61264	18.9369	19.16082	18.92977	18.77578	19.11366	18.98328	0.011828	2.456826
Gyro2 y	2k dps	-4.21055	-3.83984	-3.91548	-4.0759	-4.12492	-3.95023	-4.2334	0.011711	2.432673
Gyro2 z	2k dps	7.491982	7.383834	7.554506	7.604883	7.457798	7.485404	7.701132	0.014489	3.008718
Gyro2 t	int	26.40961	26	26	26	26	26	26	0.000479	0.100982
Gyro3 x	250dps	-10.8525	-10.8913	-12.5124	-10.3831	-9.19768	-10.9854	-10.9321	0.095807	19.90578
Gyro3 y	250dps	80.72398	81.41718	80.64504	80.30853	79.94402	80.03002	78.30344	0.0975	20.25934
Gyro3 z	250dps	-5.92837	-6.55747	-7.26471	-5.88309	-5.85468	-5.72226	-5.59863	0.09484	19.69882
Gyro3 t	int	19.29718	19.1129	18.78583	18.93536	18.96216	19.02831	18.88469	0.001698	0.353635

Table 10: Summary results of the six-position cube experiment. The coloured columns group pairs of outputs.

Sensor		Mean Properties of the 15 separate samples			Properties of turn-on-bias variation		Mean Properties of the 15 consecutive 5-min samples			Properties of in-run-bias variation		ratio
Name and axis	Range setting	Sample Mean Output	sample Std Dev. / sqrt (sample size)	sample Std Dev.	std dev. of the 15 sample means	range of the 15 sample means	Sample Mean Output	sample Std Dev. / sqrt (sample size)	sample Std Dev.	std dev. of the 15 sample means	range of the 15 sample means	
Indep. Temp	°C	24.505	0.000	0.030	0.131	0.473	23.817	0.000	0.019	0.086	0.193	2.33
ADXL1 x	16g	12.453	0.047	5.652	0.060	0.206	12.226	0.046	5.559	0.044	0.162	1.36
ADXL1 y	16g	22.372	0.056	6.770	0.058	0.198	22.478	0.054	6.495	0.043	0.162	1.34
ADXL1 z	16g	247.324	0.100	12.180	0.126	0.436	247.227	0.099	11.905	0.105	0.349	1.20
ADXL2 x	16g	7.425	0.046	5.531	0.059	0.220	7.112	0.045	5.404	0.040	0.152	1.50
ADXL2 y	16g	14.637	0.052	6.338	0.063	0.249	14.435	0.052	6.208	0.038	0.115	1.65
ADXL2 z	16g	235.952	0.099	11.986	0.163	0.497	236.311	0.099	11.875	0.084	0.257	1.95
BMA1 x	16g	-91.292	0.068	8.246	34.052	132.287	-99.405	0.067	8.077	0.117	0.391	290.41
BMA1 y	16g	-133.593	0.077	9.292	2.529	9.951	-133.935	0.075	9.047	0.146	0.464	17.34
BMA1 z	16g	787.939	0.073	8.891	973.994	3772.596	537.047	0.072	8.618	0.095	0.363	10285.99
BMA1 t	int	-4.889	0.008	0.924	0.574	2.135	-6.527	0.008	0.970	0.090	0.108	17.86
BMA2 x	16g	53.256	0.065	7.934	0.148	0.555	53.596	0.065	7.777	0.074	0.247	2.00
BMA2 y	16g	-86.178	0.084	10.174	0.167	0.510	-85.792	0.084	10.080	0.081	0.299	2.07
BMA2 z	16g	-453.149	0.068	8.308	0.293	1.102	-451.918	0.068	8.198	0.118	0.367	2.48
BMA2 t	int	2.570	0.011	1.347	0.575	1.034	0.953	0.010	1.233	0.096	0.163	6.03
BMA3 x	1.5g	527.875	0.322	39.076	1.058	3.519	499.246	0.346	41.688	1.446	5.435	0.73
BMA3 y	1.5g	-8.390	0.397	48.161	0.650	2.441	-14.840	0.429	51.630	1.000	2.978	0.65
BMA3 z	1.5g	-5539.339	0.386	46.853	4.008	13.852	-5527.056	0.408	49.151	2.178	6.894	1.84
BMA3 t	int	0.156	0.006	0.687	0.306	0.641	-1.685	0.006	0.768	0.052	0.217	3.32
BMA4 x	1.5g	326.603	0.364	44.117	1.111	3.623	325.311	0.374	45.023	0.695	2.301	1.60
BMA4 y	1.5g	-61.749	0.508	61.681	1.573	5.070	-56.896	0.518	62.432	0.717	1.945	2.19
BMA4 z	1.5g	5616.692	0.465	56.380	3.678	12.983	5617.841	0.466	56.120	0.832	3.103	4.42
BMA4 t	int	1.460	0.010	1.154	0.532	0.758	-0.259	0.009	1.134	0.073	0.234	3.18
Gyro1 x	2k dps	6.853	0.021	2.602	0.513	1.341	6.225	0.020	2.424	0.067	0.217	7.71
Gyro1 y	2k dps	-7.097	0.020	2.418	0.520	1.429	-6.558	0.018	2.114	0.083	0.285	6.27
Gyro1 z	2k dps	12.446	0.023	2.794	0.186	0.606	12.503	0.023	2.774	0.067	0.241	2.79
Gyro1 t	int	24.060	0.003	0.408	0.123	0.513	24.150	0.003	0.330	0.171	0.699	0.72
Gyro2 x	2k dps	17.925	0.021	2.489	0.410	1.147	17.913	0.020	2.360	0.170	0.493	2.41
Gyro2 y	2k dps	-5.057	0.019	2.352	0.301	1.051	-5.730	0.018	2.137	0.072	0.248	4.16
Gyro2 z	2k dps	7.238	0.025	3.024	0.213	0.681	7.015	0.025	2.965	0.087	0.309	2.46
Gyro2 t	int	28.998	0.004	0.535	0.152	0.596	29.109	0.003	0.343	0.198	0.704	0.90
Gyro3 x	250dps	-6.766	0.166	20.157	6.478	13.367	-1.270	0.151	18.168	0.246	0.981	26.35
Gyro3 y	250dps	72.746	0.168	20.388	2.886	8.159	68.448	0.152	18.361	0.503	1.673	5.74
Gyro3 z	250dps	1.407	0.162	19.627	1.032	3.037	3.184	0.157	18.920	0.442	1.468	2.33
Gyro3 t	int	21.796	0.004	0.467	0.213	0.797	22.083	0.002	0.255	0.111	0.461	1.95

Table 11: Summary statistics for the turn-on-bias experiments. BMA1 and Gyro3 are greyed out for reasons discussed in Section 5.2.4. All outputs are raw LSBs, temperature outputs are in blue.