Magnetometer-Free Orientation Estimation Using Multiple Inertial Measurement Units

T. I. Edridge

Master of Science Thesis

TU Delft Delft University of Technology

Delft Center for Systems and Control

Magnetometer-Free Orientation Estimation Using Multiple Inertial Measurement Units

Master of Science Thesis

For the degree of Master of Science in Systems and Control at Delft University of Technology

T. I. Edridge

October 9, 2020

Faculty of Mechanical, Maritime and Materials Engineering (3mE) · Delft University of Technology

Abstract

Higher-grade Inertial Measurement Units (IMUs) provide a better estimation of the orientation but are more expensive than lower-grade IMUs. The increase in availability and affordability of lower-grade IMUs over the last few years has provided an alternate solution: combine multiple lower-grade IMUs to increase sensor accuracy at low cost. Existing literature has already shown the advantages of multiple IMUs. One advantage is a better estimation of linear and rotational motion. Another advantage is Fault Detection and Isolation (FDI) from component redundancy. However, existing literature is confined to a set of configurations that are impractical (platonic solids) or small (arrays). In this thesis, we focus on combining FDI with estimation of linear and rotational motion and further improve FDI by changing the IMU configuration from an array to scattered placement of IMUs. First, we show in a Monte Carlo simulation that because of the array configuration, the angular velocity and angular acceleration covariance decrease proportionally with the number of accelerometers used rather than quadratically, which is theoretically possible. We also reveal that for human motion capture, the angular velocity and the specific covariance cannot practically compete with a higher-grade IMU. Second, we propose a novel FDI method, based on the parity space method, to reject disturbances which can be applied to bigger configurations and high rotational motion. As a proof of concept for real-life implementation, we tested FDI on a simulated human arm with three configurations. The first with lower-grade IMUs placed evenly around the arm. The second and third are placed in the middle of the arm where the second is a lower-grade array and the third is a single higher-grade IMU. A disturbance was applied, which resulted in a 4°, 21° and 13° maximum yaw angle error respectively, illustrating the superior disturbance rejection capability of the larger configuration. In conclusion, we show that combining multiple IMUs improves accuracy and when used in a larger configuration, may be beneficial over a higher-grade IMU. The practical limitations are discussed in this thesis.

Table of Contents

List of Figures

List of Tables

Acknowledgements

I would like to thank my supervisor dr. M. Kok for her time and assistance during the thesis. I would also like to thank Frida, Evan, Bart, Daan and Stefan for your helpful comments and sharing your insights during the bi-weekly meetings.

Delft, University of Technology
October 9, 2020

T. I. Edridge

Chapter 1

Introduction

In Section 1-1 of this chapter the motivation of the research is shown, as well as the main goal that guides this thesis. Afterwards, in Sections 1-2 and 1-3 the existing work regarding multiple IMU fusion and Fault Detection and Isolation are discussed. This is followed by Section 1-4 the research questions of this thesis are introduced. The contributions of this thesis is shown thereafter in Section 1-5. This chapter ends on introducing the outline of the thesis in Section 1-6.

1-1 Background Research Motivation

Inertial Measurement Units (IMUs) are electronic devices used for measuring the movement of an object or person, also known as motion capture [6]. Over the last decades, Micro-Electromechanical System (MEMS) technology has revolutionized the IMU industry [25], making them ubiquitous [29]. MEMS IMUs (see Figure 1-1 for an example) are lightweight, cheap and small-sized [19], which has made them preferable over the most common alternative method of motion capture, optical motion capture [44], which is limited in practicality and affordability [39]. These developments have increased use of MEMS IMUs in motion capture for varying fields [6, 25]. Some examples are bio-mechanics, video games, sports and navigation.

A MEMS IMU consists of two different tri-axial sensors. One sensor is the gyroscope, which measures the angular velocity. The gyroscope measurement is integrated to find the orientation. When the gyroscope measurement is corrupted by a bias, the error will accumulate, also known as drift from the true orientation [25]. The other sensor is the accelerometer, which measures the gravity vector and the linear acceleration, also known as the specific force. The accelerometer reveals information on the direction of the gravity vector, which gives information on the roll (around the x-axis) and pitch (around the y-axis) angle. This information can be used to compensate for the drift around the roll and pitch angle. However, the accelerometer provides no information on the rotation of the gravity vector, leaving the yaw (around the z-axis) angle uncompensated for. Frequently, a magnetometer is added

Figure 1-1: Xsens [1] MTw Awinda wearable MEMS IMU containing a gyroscope and an accelerometer. Arm for scale.

to compensate for the drift around the yaw angle. However, magnetometers are prone to disturbances in indoor environments due to ferromagnetic materials [39]. This makes the use of magnetometers unfavourable.

MEMS IMUs[1] come in different grades[2], where the higher-grade IMUs have higher measurement accuracy and provide more stable readings [32]. The drift is generally worse for lower-grade IMUs as higher-grade IMUs rely on higher angular velocity accuracy and better bias stability [1]. Bias stability, how much the bias changes over time, is one of the main sources of error [52] of drift. The higher bias stability results in precise calibration in a factory being more effective for higher-grade IMUs than lower-grade IMUs [31]. As a result, calibration for cheaper IMUs is usually left out. However, despite the higher accuracy of higher-grade MEMS IMUs, their high costs limit their use in many applications.

Over the last years, the prices of IMUs have decreased and the availability has increased [38]. This has provided a potential solution to improve the orientation estimate of lower-grade IMUs: combine multiple lower-grade IMUs to increase motion accuracy at low cost. Combining multiple sensors together, also known as Sensor Fusion, forms a synthetic IMU which shows high performance to price ratio [8][3]. The general idea of a synthetic IMU is to mimic higher-grade IMU in terms of accuracy, but there is more to be gained than just fusing multiple sensors [29]. Two important additional benefits of using multiple IMUs are: direct estimation of the angular acceleration and Fault Detection and Isolation (FDI) based on component redundancy. FDI based on component redundancy allows for validating IMU measurements against each other. This gives the overall goal of this thesis:

How can multiple IMUs be used to improve orientation estimation?

Before more detailed research questions can be formulated, we first will look into existing work on multiple IMUs. The fusion of multiple IMUs is discussed first followed by FDI in Sections 1-2 and 1-3 respectively.

[1] Henceforth, the term MEMS shall be dropped. All IMUs referred to in this thesis are MEMS IMUs

[2] In this thesis the term lower-grade is used for the IMUs on the MIMU4844 array from Inertial Elements [23] The term higher-grade is used to refer to the MTi 300 from Xsens [1]

[3] The MIMU4844 array containing 32 IMUs is priced around €500 where the MTi 300 is priced around €3000

1-2 Existing Work on Fusion of Multiple IMUs

The theory of using multiple IMUs is not new. Publications began appearing around the 1960s and has found renewed interest in the last two decades due to the new MEMS technology (see [29] for a detailed literature review on IMU arrays). The main focus of using multiple IMUs for orientation estimation is the increased accuracy and reliability in estimating linear and rotational motion [29]. As previously mentioned, the gyroscopes measure the angular velocity and the accelerometers measure the specific force. By averaging out these measurements, more accurate estimates of the angular velocity and specific force can be attained. While averaging is computationally efficient, more accurate estimates are attainable [40]. This is because accelerometers are also affected by rotational motion. Both the angular velocity and the angular acceleration are measured by the accelerometer. The angular velocity and the angular acceleration induce a Centrifugal and Euler force respectively. The forces measured by the accelerometer are shown in Figure 1-2. By recognizing the rotational forces the accelerometers can provide information on the angular velocity, allowing more accurate angular velocity estimates. Additionally, the angular acceleration can now be estimated directly as well. Because the rotational motion is now separated from the accelerometer measurements, a better estimate on the specific force is attainable, as in previous assumptions, the rotational motion could be considered a source of error.

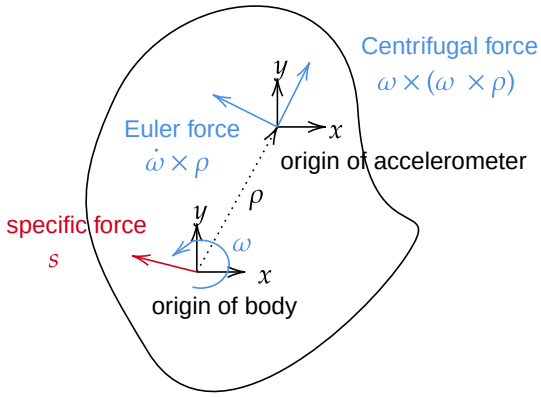

Figure 1-2: Rigid body experiencing linear (red) and rotational (blue) motion, resulting in a specific force s, Centrifugal force $\omega \times (\omega \times \rho)$ and Euler force $\dot{\omega} \times \rho$. The angular velocity is denoted by ρ, angular acceleration is denoted by $\dot{\omega}$ and relative position of the accelerometer by ρ.

Research on fusing multiple IMUs is mostly limited to an array configuration, an example of an IMU array can be seen in Figure 1-3. The relative locations of the accelerometers on the IMU array are small ($\approx 5mm$). The advantage of an IMU array is that the IMU array is practical in use and the locations are fixed with respect to each other. The downside, however, is that the relative locations of the accelerometers are in fact small. As shown in [40], the accelerometer provides more information on the angular velocity and the angular acceleration when the distance between sensors increases. More specifically, the covariance of the angular velocity and the angular acceleration decreases proportionally with the distance between accelerometers squared.

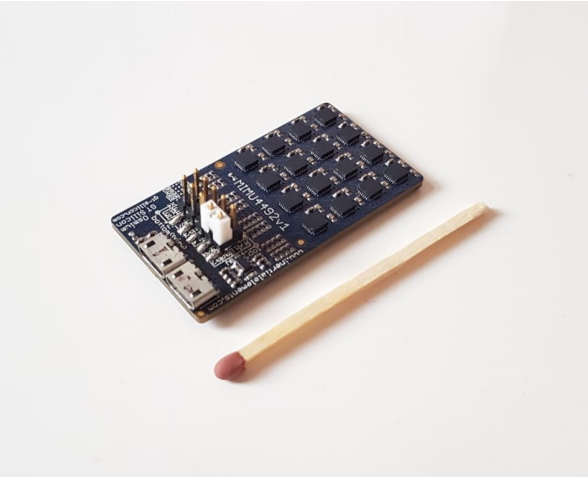

Figure 1-3: MIMU4844 IMU array from Inertial Elements [2] containing 32 accelerometers, 32 gyroscopes and 32 magnetometers. Recreated form [40] Match for scale.

Arrays can be divided into two types: IMU arrays (including gyroscopes and acclerometers) and accelerometer-only arrays. Accelerometer-only arrays are generally used in high dynamic environments [29], as it can be shown that accelerometers provide high accuracy at higher dynamics [40], for example in crash tests for vehicles [29]. The downside of using accelerometer-only arrays is that the angular velocity can be estimated up to a sign ambiguity[4] due to the non-linearities. A lot of different methods have been introduced to resolve the sign ambiguity, see [7] for an overview. IMU arrays are more suitable for lower dynamic environments [29] such as human motion capture [25]. This is because the gyroscopes provide a good estimation on the angular velocity.

A Maximum Likelihood (ML) approach for IMU arrays to estimate the linear and rotational motion the specific force, angular velocity and angular acceleration of a rigid body was presented in [40]. As an extension on this work, [51] showed improvements that can be made using the angular acceleration as a motion model. Resulting in a Maximum A Posteriori (MAP) estimate of the angular velocity. This resulted in higher accuracy of the angular velocity and allows for the extension towards accelerometer-only arrays. The theory presented in [40] was also extended in [39], where the authors used a two-step attitude integration architecture [39] by [36] which allows for computationally cheaper estimations. An IMU array of 8 IMUs was compared to a single IMU and the peak estimation error was reduced by a factor of 2.5. What remains to be investigated is how well multiple lower-grade IMUs compare against a higher-grade IMUs.

1-3 Existing Work on Fault Detection and Isolation

FDI has been a research topic since the early days of Inertial Navigation [9]. The term Inertial Navigation is estimating the position, velocity and orientation of an object using IMUs. The focus was on building fail-safe systems, where fail-safe denotes the ability to identify

[4]The magnitude can be estimated but not the sign.

and isolate faulty sensors [9]. Recently, with the rise of cheaper MEMS IMUs, interest in FDI application for IMU has increased. Commonly, the application of FDI is combined with additional measurements such as GPS. GPS, however, has poor performance in indoor and urban environments [13], this makes the use of GPS for indoor applications undesirable.

The majority of the research on FDI focuses on the parity space. The parity space is a method used for FDI based on either analytical or component redundancy. Analytical redundancy refers to the ability to link past states with current states using a mathematical model. This type of redundancy is applied when using a single IMU. Component redundancy refers to having multiple IMUs, where the measurements can be tested against one another to check for faults. Besides this difference, the method and the application of the parity space method are the same.

Research has been done on optimal configurations using IMUs, see for example [18, 35, 42, 49, 43, 19, 16]. The results show that two types of configurations are optimal in detecting faults: the platonic solids and the half cones which are shown in Figure 1-4. The problem with such configurations, however, is twofold. The first problem is that the costs of such a special configuration outweigh the costs of adding more sensors [15], which would have the same results. The second problem is that placing sensors in such a way is usually impractical. For example, placing sensors on a human body part would not result in the optimal configuration.

Previous work on the parity space using component redundancy, see e.g. [9, 50, 49], assumes that the accelerometer only measures the specific force. This assumption may hold for low angular motion and closely placed accelerometers. For conditions where the angular motion is high and accelerometers are placed further away, this will result in unmodeled errors [3]. The angular acceleration can be found by taking the angular velocity time derivative, however, this is considered to be unreliable as this induces large errors on the angular acceleration estimate. In some accounts these states were added, e.g. [3, 15], however, their works do not make use of the parity space method.

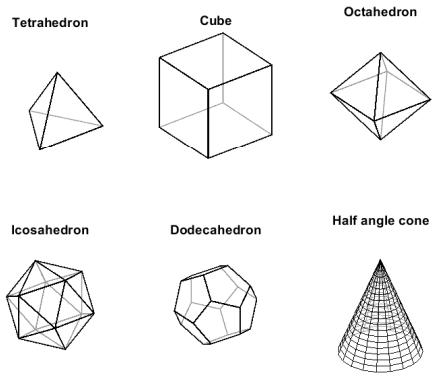

Figure 1-4: All 5 platonic solids and a half angle cone.

1-4 Research Questions

Restating the overall goal of this thesis:

How can multiple IMUs be used to improve orientation estimation?

With the knowledge based on previous work we can now ask more specific questions:

How can the accuracy of multiple lower-grade IMUs compare against a higher-grade IMU? It was shown in [40] that multiple IMUs can be fused into a single synthetic IMU. This synthetic IMU has increased accuracy at higher angular velocities. It was also shown that with more distance between accelerometers, the angular velocity accuracy increases. Additionally, a direct estimation of the angular acceleration becomes attainable. The angular acceleration was used in [51] as a motion model so that the accuracy can be even further increased. The novelty in this thesis is the comparison of not only the array configuration against a higher-grade IMU, but also comparing a larger configuration, the size of human arm against a higher-grade IMU.

Can the parity space method be used for larger configurations? Previous accounts on the parity space method assume that the accelerometer measures the specific force and no angular motion. This can give good results at low angular motion and short distances between the accelerometers. However, if the angular velocity gets higher or the distance between accelerometers increases, the angular motion will create larger differences between the measurements. This may cause the parity space method to wrongfully identify faulty sensors. The authors of [51] and [40] have shown an estimator such that the angular motion and relative distances between accelerometers may be taken into account. The extension of this estimator on the parity space is novel and will be validated using a Monte Carlo simulation.

How can the component redundancy of IMUs be used to reject disturbances? By making use of multiple IMUs, the amount of information available becomes redundant at every time step. The IMUs can be compared to each other to determine if IMUs are corrupted by faults. If a sensor is faulty, it can be isolated until the fault has disappeared. Previous work has focused on specific configurations, however, other configurations might provide useful information. If for example, IMUs were scattered along the length of a human arm and a disturbance is applied that dissipates along the length of this arm, can a scattered configuration reject the disturbance when an IMU array cannot?

1-5 Contributions of this Thesis

This thesis aims to analyze how multiple IMUs can be used to obtain better estimates of the orientation. The contribution of this thesis is threefold. Firstly, based on the works in [40] the covariance analysis is extended towards a real IMU array such that more realistic covariance estimates are shown. Secondly, by making use of several motion priors, different multiple lower-grade IMU configurations are compared to each other and a higher-grade IMU. Thirdly, the parity space method as an extension on the work in [51] and [40].

1-6 Outline of this Thesis

In Chapter 2, it is shown how an orientation estimate can be found using a single IMU. Afterwards, in Chapter 3, it is shown how measurements can be fused into a synthetic IMU. It becomes apparent how much information can be obtained and how it can be combined with Chapter 2 to obtain an orientation estimate. Subsequently, in Chapter 4 results are shown on how much information can be obtained and different setups are compared against a high-grade IMU. In Chapter 5, it is discussed how component redundancy can be used to detect and isolate faulty sensors. It is also shown how this can be integrated with the theory presented in Chapters 2 and 3. It is also shown how FDI can be combined with the orientation estimate found in 3. The novel parity space method is simulated and validated in Chapter 6. Finally, in Chapter 7, the results are discussed, further recommendations are given and conclusions are drawn.

Chapter 2

Single IMU Orientation Estimation

This chapter introduces the fundamental basics for orientation estimation using accelerometers and gyroscopes. In this chapter, a Single IMU is used to find the orientation estimation. It will be shown later in Chapter 3 how a Synthetic IMU can be fused from Multiple IMUs. The Synthetic IMU can take the place of the single IMU to find the orientation estimation. Section 2-1 introduces the coordinate frames necessary to express the orientation. Afterwards, in Section 2-2, different motion models are introduced. Thereafter, in Section 2-3 filtering is introduced, where the measurements used to estimate the orientation. In Section 2-4 calibration is introduced for the gyroscope and accelerometer. Finally, in Section 2-5, the previous sections are brought together into an orientation estimation scheme.

2-1 Describing Orientation

The orientation and position of a body in a 3-D space can be represented in different ways. However, in order for these results to have any meaning, coordinate frames can be introduced to describe the relative change from an initial position. The coordinate frames are shown first after which the relevant orientation parameterizations are discussed.

2-1-1 Coordinate Frames

The orientation estimates of a body in a space can hold many different solutions. It is important to identify which frames play an important role in describing these estimates. The relevant frames here are defined as

- Sensor frame s - frame in which the measurement is done by an individual sensor.
- Body frame b - frame attached to the center of mass of the body of which the position and orientation is to be estimated.
- Navigation frame n - frame in which the body moves.

The measurements take place directly in the sensor frame. These measurements are fused into a single state vector in the body frame. This state vector can then be used to find a navigation solution in the navigation frame.

2-1-2 Orientation Parameterization

The orientation of a rigid body can be parameterized in different ways. The relevant orientation parameterizations are shown below. Additional information on other orientation parameterizations can be found in Section A-1 of Appendix A. These different parameterizations can be used interchangeably as they can be expressed in terms of the SO(3) group, which describes the group of all rotations in $\mathbb{R}^3$. The SO(3) group itself is in the matrix Lie group and the SO(3) group has the following properties [21], where R is a rotation matrix R

$$R^T = R^{-1}, \quad R^T R = \mathcal{I}_3, \quad \det(R) = 1. \tag{2-1}$$

Rotation matrices A rotation matrix describes the relative rotation of a fixed-length vector from one coordinate system to another coordinate system, in this case from s to b. The indices of the matrix in equation (2-2) contain the cosine, shown as C, of the angle between the new axis and the old axis, where the first superscript denotes coordinate system s and the second superscript denotes coordinate system b [4]. The subscripts denote what part of x^s is rotated towards what part of x^b. A rotation matrix that describes the rotation $x^b = R^{bs}x^s$ can be written as

$$\underbrace{\begin{bmatrix} x_1^b \\ x_2^b \\ x_3^b \end{bmatrix}}_{x^b} = \underbrace{\begin{bmatrix} C_{x_1^b,x_1^s} & C_{x_1^b,x_2^s} & C_{x_1^b,x_3^s} \\ C_{x_2^b,x_1^s} & C_{x_2^b,x_2^s} & C_{x_2^b,x_3^s} \\ C_{x_3^b,x_1^s} & C_{x_3^b,x_2^s} & C_{x_3^b,x_3^s} \end{bmatrix}}_{R^{bs}} \underbrace{\begin{bmatrix} x_1^s \\ x_2^s \\ x_3^s \end{bmatrix}}_{x^s}. \tag{2-2}$$

The rotation matrix also has some unique properties that must be upheld to describe a rotation. A rotation matrix is a unique way of describing a rotation, where all nine entries of the rotation matrix describe a different part of the rotation. the properties of the rotation matrix are that given by equation (2-1).

Quaternions Unit quaternions provide another approach on how to parameterize rotations. This 4-dimensional representation has no feasible physical interpretation but is known for its applicability in Gaussian filters such as the Kalman Filter. The main reason for this is that the unit quaternion is not prone to wrapping, unlike the other parameterizations and gives a non-singular orientation [21]. The quaternion, however, is non-unique as $-q$ and q describe the same orientation. Furthermore, the unit quaternion is said to be in the group $\mathbb{S}^3$ under the following conditions

$$q = \begin{bmatrix} q_0 & q_1 & q_2 & q_3 \end{bmatrix}^T = \begin{bmatrix} q_0 \\ q_v \end{bmatrix}, \qquad q \in \mathbb{R}^4, \qquad ||q||_2 = 1. \tag{2-3}$$

To multiply quaternions, the quaternion multiplication $\odot$ must be introduced and is denoted by

$$p \odot q = \begin{bmatrix} p_0 q_0 - p_v \cdot q_v \\ p_0 q_v + q_0 p_v + p_v \times q_v \end{bmatrix}. \tag{2-4}$$

Orientation deviations Orientation deviations are closely related to the previously discussed orientation parameterizations as the rotation can be described as a linearization point around the quaternion or rotation matrix and an orientation deviation around this point [25]. This property can be used in the Extended Kalman Filter (EKF) and will be shown in section 2-3. The rotation quaternion using orientation deviations can be defined as

$$q_k^{nb} = \exp_q(\frac{\bar{\eta}_t^n}{2}) \odot \tilde{q}_k^{nb}. \tag{2-5}$$

Where the mapping from quaternion to SO(3) with the approximations based on the small angle theorem are denoted as

$$\exp_q(\bar{\eta}) = \begin{bmatrix} \cos ||\eta||_2 \\ \frac{\eta}{||\eta||_2} \sin ||\eta||_2 \end{bmatrix}, \qquad \exp_q(\eta) \approx \begin{bmatrix} 1 \\ \eta \end{bmatrix}. \tag{2-6}$$

2-2 Motion Models for Orientation Estimation

The orientation of a rigid body can be determined by making use of motion and measurement models. These can be described into a state-space model which in turn can be used in a Kalman Filter, as will be shown later in Section 2-3. The measurements are sampled at a certain frequency at which the motion models are Euler discretized. The measurements are assumed to take on additive Gaussian white noise and a bias.

2-2-1 Discretized Motion Models

For orientation estimation the gyroscopes provide direct measurements of the angular velocity. Combing these and by making use of quaternions, the attitude can be expressed as (where $\odot$ is the quaternion multiplication operator defined in equation (2-4)) with its Euler discretization [25] gives the following update equations

$$\omega_{k+1}^b = \omega_k^b + e_{\omega,k}, \tag{2-7a}$$

$$q_{k+1}^{nb} = q_k^{nb} \odot \exp_q(\frac{T}{2}\omega_k^b). \tag{2-7b}$$

Where q_k^{nb} denotes the quaternion from the body to the navigation frame, T denotes the sampling time and ω_k^b denotes the angular velocity. Furthermore, $e_{\omega,k} \sim \mathcal{N}(0, \Sigma_\omega)$ is Gaussian white noise.

2-2-2 Measurement Models

MEMS IMUs are readily available and provide angular velocity and linear acceleration measurements. These measurements, however, are corrupted by bias and noise [52]. For better estimations it becomes key to account for these biases as these errors will accumulate when integrating [52]. For this a measurement model is needed. Additionally, for the accelerometers in orientation estimation, one would like to estimate the gravity vector as well.

Gyroscope Measurement Model The gyroscope, which measures the angular velocity of the body, which is independent of its location is given by

$$y_{\omega,k} = \omega_k^s + b_{\omega,k}^s + e_{\omega,k}^s. \tag{2-8}$$

The noise vector $e_{\omega,k}^s$ can be modeled as Gaussian white noise with only diagonal entries and is assumed to be modeled as $e_{\omega,k}^s \sim \mathcal{N}(0, \Sigma_\omega)$. The bias, denoted by $b_{\omega,k}^s$ is slowly time-varying and can be modeled as constant or as a random walk model given by

$$b_{\omega,k+1}^s = b_{\omega,k}^s + e_{b_\omega,k}^s, \tag{2-9}$$

depending on the duration time of the experiment [25]. Furthermore, the assumptions are often made that the earth's rotation and that the navigation frame have negligible effect on the measurements [25].

Accelerometer Measurement Model The accelerometer, which measures the specific force has a measurement model which can be denoted by

$$y_{a,k} = s_k^s + b_{a,k}^s + e_{a,k}^s. \tag{2-10}$$

The noise term $e_{a,k}^s \sim \mathcal{N}(0, \Sigma_a)$ can be modeled as Gaussian white noise. Furthermore, the assumption here is made that the Coriolis acceleration is negligible [25]. The bias, being similar to the bias in the gyroscope, can again be modeled as constant or slowly time-varying as a random walk model given by

$$b_{a,k+1}^s = b_{a,k}^s + e_{b_a,k}^s. \tag{2-11}$$

The specific force can be recast into terms of a gravity vector and linear acceleration which results into the final measurement model of

$$y_{a,k} = \underbrace{(a_k^s - R_k^{sn} g^n)}_{s_k^s} + b_{a,k}^s + e_{a,k}^s, \tag{2-12}$$

however, in some applications a good enough approximation would be to ignore the linear acceleration and only take the gravity vector into account [25]. This results in the simplified measurement model of

$$y_{a,k} = -R_k^{sn} g^n + b_{a,k}^s + e_{a,k}^s. \tag{2-13}$$

2-2-3 Orientation Estimation

Using the previously mentioned motion models and the previously mentioned measurement models one can combine these to estimate orientation. The assumptions can be made that the rotation of the earth and the Coriolis acceleration are negligible [25]. An assumption can be made that the linear acceleration measured by the accelerometer is almost zero [25], which results in only the gravity vector having effect.

$$q_{k+1}^{nb} = q_k^{nb} \odot \exp_q(\frac{T}{2}(y_{\omega,k} - b_\omega - e_{\omega,k})), \tag{2-14a}$$

$$y_{a,k} = -R_k^{bn} g^n + e_{a,k}. \tag{2-14b}$$

2-3 Using the Extended Kalman Filter for Orientation Estimation

Now that the motion models and the measurements models are discussed they can be fused together. This can be done via smoothing or filtering; the former approach gives better state estimates but is computationally more heavy [25]. Filtering is less accurate but faster and can be used on-line. These features make filtering an attractive solution for longer data sets and for many sensors. For this reason, the Extended Kalman Filter (EKF) will be discussed which forms a basis for the Multiplicative Extended Kalman Filter (MEKF). How a prior is obtained is not discussed, the reader is referred to [25].

2-3-1 Extended Kalman Filter

The Kalman Filter (KF) is an unbiased minimum-variance estimator that gives a one-step-ahead prediction of the filtered state estimate [48]. The Extended Kalman Filter (EKF) is an adaptation on the KF that allows for estimation of non-linear functions. The EKF works under the assumption that noise is additive and that both process noise Q and measurement noise R are distributed as Gaussian zero-mean with constant or time-varying co-variance [25]. The EKF is presented in [25] and is as follows. The EKF is a recursive algorithm that performs two updates, one being a time update where the next state $x_{k+1|k}$ and its covariance $P_{k+1|k}$ are estimated. The algorithm is shown in Algorithm 1.

Algorithm 1: Extended Kalman Filter

compute the prior and set initial values;
for $k = 2 : N$ **do**
Time update;

$$\hat{x}_{k+1|k} = f_k(\hat{x}_{k|k}, u_k), \tag{2-15a}$$
$$P_{k+1|k} = F_k P_{k|k} F_k^T + G_k Q G_k^T. \tag{2-15b}$$

Measurement update;

$$\hat{x}_{k+1|k+1} = \hat{x}_{k+1|k} + K_{k+1}\epsilon_{k+1}, \tag{2-16a}$$
$$P_{k+1|k+1} = P_{k+1|k} - K_{k+1} S_{k+1} K_{k+1}^T, \tag{2-16b}$$

end

Where the terms F_k and G_k are linearized as defined as

$$F_k = \frac{\partial f_k(x_k, u_k, e_k)}{x_k}|_{e_k=0, x_k=\hat{x}_{k|k}}, \quad G_k = \frac{\partial f_k(x_k, u_k, e_k)}{e_k}|_{e_k=0, x_k=\hat{x}_{k|k}}, \tag{2-17}$$

and ϵ_{k+1}, S_{k+1} and K_k are defined as

$$\epsilon_{k+1} = y_{k+1} - \hat{y}_{k+1|k}, \quad S_{k+1} = H_{k+1} P_{k+1|k} H_{k+1}^T + R, \quad K_{k+1} = P_{k+1|k} H_{k+1}^T S_{k+1}^{-1}, \tag{2-18}$$

and $\hat{y}_{k+1|k}$ and H_k are defined as

$$\hat{y}_{k+1|k} = h(\hat{x}_{k+1|k}), \qquad H_k = \frac{h_{k+1}(x_{k+1})}{\partial x_{k+1}}|_{x_{k+1}=\hat{x}_{k+1|k}}. \tag{2-19}$$

2-3-2 Multiplicative Extended Kalman Filter for Orientation Estimation

The regular EKF must be adapted to be used for orientation estimation so that it can utilize quaternions. One of the useful adaptations on the EKF is the MEKF, which is designed for orientation estimation. The MEKF uses orientation deviations instead of unit quaternions as described in Section 2-1-2. This is advantageous as orientation deviations are computationally more attractive than quaternions due to their lower-dimensional state [25, 33]. When using unit quaternions in the EKF one has to normalize the resulting solution for it to represent a pure rotation, creating an additional step compared to the regular EKF. This is because the orientation deviation $\hat{\eta}_k^n$ is non-zero and has to be reset [25]. The orientation deviation can be seen as an error between the truth and the estimate [28] in terms of an angle-axis representation. The MEKF described in [25] is shown in the Algorithm 2.

Algorithm 2: Multiplicative Extended Kalman Filter for Orientation Estimation

compute the prior and set initial values;
for $k = 2 : N$ **do**
Time update;

$$\hat{q}_{k|k-1}^{nb} = \hat{q}_{k-1|k-1}^{nb} \odot \exp_q(\frac{T}{2} y_{\omega,k-1}), \tag{2-20a}$$
$$P_{k|k-1} = P_{k-1|k-1} + G_{k-1} Q G_{k-1}^T. \tag{2-20b}$$

Measurement update;

$$\hat{\eta}_k^n = K_k \epsilon_k, \tag{2-21a}$$
$$\tilde{P}_{k|k} = P_{k|k-1} - K_k S_k K_k^T. \tag{2-21b}$$

Relinearize;

$$\hat{q}_{k|k}^{nb} = \exp_q(\frac{\hat{\eta}_k^n}{2}) \odot \tilde{q}_{k|k-1}^{nb}. \tag{2-22}$$

end

Where ϵ_k, K_k and S_k are defined in equation (2-18). G_k, y_k, $\hat{y}_{k|k-1}$ and H_k are defined as

$$G_{k-1} = T\tilde{R}_{k|k-1}^{nb}, \tag{2-23a}$$
$$y_k = y_{a,k}, \tag{2-23b}$$
$$\hat{y}_{k|k-1} = -\tilde{R}_{k|k-1}^{bn} g^n, \tag{2-23c}$$
$$H_k = -\tilde{R}_{k|k-1}^{bn} [g^n \times]. \tag{2-23d}$$

2-4 IMU Calibration

MEMS IMUs as mentioned are prone to systematic errors and biases. Even though these bias errors can be precisely calibrated by the manufacturers, this is usually rather expensive [31] and these errors may change over time [52]. For these lower quality sensors, cross-axis errors can take place which has to be compensated for. Compensating for the errors leads to three additional parameters [14] that have to be estimated. Additionally, when multiple sensors are placed on a rigid body, the relative orientation between sensors becomes important, again leading to three additional parameters to be estimated [31]. The accuracy of these sensors is therefore highly dependent on the calibration techniques to remove these biases and systematic errors [30].

2-4-1 Gyroscope Calibration

Gyroscopes are straightforward to calibrate, this is because their measurements are independent of any geometric configuration and when laid on a flat surface, the measured rotational velocity should be zero. This is illustrated by the measurement model as shown in (2-24). Furthermore, the co-variance can be determined as the deviation of the measurements. The measurement model of the gyroscope

$$y_{\omega,k} = \omega_k + b_{\omega,k} + e_{\omega,k}, \tag{2-24}$$

is assumed to be corrupted by a slowly time-varying $b_{\omega,k}$ bias and zero-mean white noise denoted by $e_{\omega,k} \sim \mathcal{N}(0, \Sigma_\omega)$ [25]. Under the assumption that the gyroscope is properly calibrated, the axes are assumed to have independent measurements and as such, the co-variance Σ_ω can be modeled as a diagonal matrix which takes the form of

$$\Sigma_\omega = \begin{bmatrix} \sigma^2_{\omega,x} & 0 & 0 \\ 0 & \sigma^2_{\omega,y} & 0 \\ 0 & 0 & \sigma^2_{\omega,z} \end{bmatrix}. \tag{2-25}$$

By placing the gyroscope on a flat surface ensures that the only rotation that the gyroscope undergoes is due to the earth's rotation. However, the magnitude of this rotation is insignificant compared to sensor noise and can thus be assumed to be approximately zero [34].

When the gyroscope lies flat on the table for a period of time, the biases can be found as the mean of the data and the co-variance can be found as the standard deviation from this mean [30], which can be modeled as in (2-24).

2-4-2 Accelerometer Calibration

Contrary to gyroscope calibration, accelerometer calibration is not as straightforward. The calibration technique described, taken from [30, 31, 26, 34] does not require the accelerometer to be placed in specific orientations. This calibration method requires the accelerometer to be rotated, slowly, in as many directions as possible. The basic principle of the calibration

technique relies on the assumption that the local gravity field is constant. It is also important to note that accelerometer calibration is often neglected for orientation estimation. This is because the approximation is made that the gravity vector is equal to the specific force. It is, however, useful for FDI. Not accounting for these errors leads to larger residuals which compromises FDI capabilities [3].

The accelerometer measurement model,

$$y_{a,k} = s_k^s + b_a^s + e_{a,k}^s, \tag{2-26}$$

where $b_{a,k}$ is the bias and $e_{a,k}$ is Gaussian white noise denoted by $e_{a,k} \sim \mathcal{N}(0, \Sigma_a)$ similar to the gyroscope noise. The measurement model can be adapted to include cross-axis interference and non-orthogonal axes, which makes for 9 parameters to be estimated [14]. This results in the following uncalibrated measurement model, under the assumption that linear acceleration is approximately zero

$$y_{a,k} = -DR_k^{sn} g^n + b_a^s + e_{a,k}^s. \tag{2-27}$$

The rotation matrix R_k^{sn} denotes the orientation of the sensor. The matrix D models the errors in cross-axis interference, inter-sensor misalignment and non-orthogonal axes, which can be decomposed into a scaling matrix K and a misalignment matrix T. These take the form of

$$D = \underbrace{\begin{bmatrix} k_x & 0 & 0 \\ 0 & k_y & 0 \\ 0 & 0 & k_z \end{bmatrix}}_{K} \underbrace{\begin{bmatrix} 1 & -\alpha_{yz} & \alpha_{zy} \\ \alpha_{zy} & 1 & -\alpha_{zx} \\ -\alpha_{xy} & \alpha_{yz} & 1 \end{bmatrix}}_{T}. \tag{2-28}$$

The assumption is made that the x-axis of the sensor coincides with the x-axis of the body and that y-axis of the body is in the plane of x-y-plane of the sensor, then the lower triangular part under the diagonal of the matrix is reduced to zero [31, 34]. Then, matrix D is constructed as

$$D = \underbrace{\begin{bmatrix} k_x & 0 & 0 \\ 0 & k_y & 0 \\ 0 & 0 & k_z \end{bmatrix}}_{K} \underbrace{\begin{bmatrix} 1 & -\alpha_{yz} & \alpha_{zy} \\ 0 & 1 & -\alpha_{zx} \\ 0 & 0 & 1 \end{bmatrix}}_{T}. \tag{2-29}$$

The 9 parameters to be estimated are then the three scaling factors, three misalignment angles and three biases, given by $\theta_a = \{D, b_a\}$ When these parameters are known, the calibrated accelerometer measurements can be found by

$$y_{a,k}^{cal} = D^{-1}(y_{a,k} - b_a). \tag{2-30}$$

The calibration problem can be formulated as a ML estimation problem, as formulated by [30] is given by

$$\hat{\theta}_a = \arg\max_{\theta_a} p_{\theta_a}(y_{a,1:N}) = \arg\max_{\theta_a} \prod_{k=1}^{N} p_{\theta_a}(y_{a,k}|y_{a,1:k-1}). \tag{2-31}$$

Under the assumption that the measurement model is Gaussian, this formulation can be rewritten into a minimization problem. The negative log-likehood provides the expression of

the following cost function

$$\hat{\theta}_a = \arg\min_{\theta_a} \frac{1}{2} \sum_{k=1}^{N} \left(||y_{a,k} - \hat{y}_{a,k|k-1}(\theta_a)||_{S_k^{-1}(\theta_a)} + \log\det S_k(\theta_a) \right). \tag{2-32}$$

This problem is non-linear as the innovation co-variance matrix S_k and $\hat{y}_{a,k|k-1}(\theta_a)$ have to be solved recursively in the MEKF to avoid a degenerative solution.

2-5 Orientation Estimation Solution

In the previously discussed sections of this chapter, the foundations have been introduced upon which an orientation estimation using a single IMU can be made. The accelerometer and gyroscope can be calibrated beforehand so that the error sources can be compensated for. After removal, the angular velocity provides the time update in the MEKF and the specific force provides the measurement update in the MEKF. The result after the MEKF updates is the rotation. This scheme is summarized in Figure 2-1.

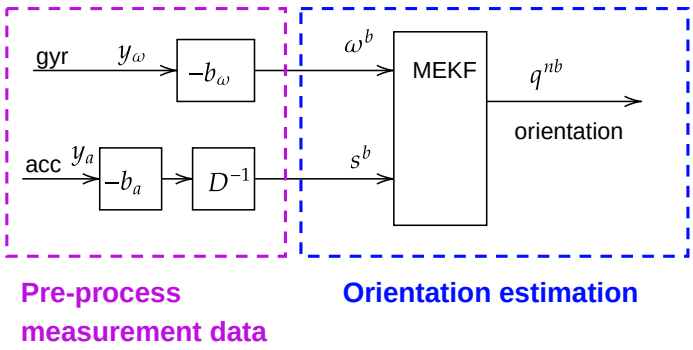

Figure 2-1: An Orientation estimation scheme using a single gyroscope and accelerometer. In the purple rectangle, the measurement data is pre-processed and in the blue rectangle, the orientation is estimated.

The orientation estimation scheme shown in Figure 2-1 forms the basis orientation estimation using an IMU. It will become apparent in Chapter 3 how Multiple IMUs can be fused to form a Synthetic IMU. This Synthetic IMU has an estimate of the angular velocity and the specific force, similar to a single IMU. The Synthetic IMU can then take the place of the single IMU for orientation estimation.

Chapter 3

Fusion of Multiple IMUs into a Synthetic IMU

This chapter expands on Chapter 2, where an orientation estimation method was introduced based on a single gyroscope and accelerometer pair. This chapter introduces how multiple IMUs can be fused into a synthetic IMU with more accurate estimations on the angular velocity and specific force. In Section 3-2, the identifiability conditions are shown, where it becomes apparent that it now also possible to estimate the angular acceleration. Afterwards, the Cramér-Rao Lower Bound (CRLB) is introduced and it is shown that more accurate estimates are possible than simply averaging out sensor readings in Section 3-3. In Section 3-4, a method is shown on how the measurements can be fused to estimate the angular velocity, angular acceleration and the specific force. Subsequently, in Section 3-5 it is shown how the angular acceleration can aid as a motion model. Once the angular velocity and specific force are estimated, together they can classified as a synthetic IMU. To provide good estimations on the angular velocity and angular acceleration, the position needs to be known accurately. Therefore, in Section 3-6 an accelerometer position calibration method is explained. Finally, in Section 3-7 it is shown how the information presented in this chapter can be used for orientation estimation.

3-1 Introduction

When using a single IMU, it is often assumed that the accelerometer only measures the specific force. This results in the measurement equation as shown in (2-27) and is restated here as

$$y_{a,k} = \underbrace{(a^s_k - R^{sn}_k g^n)}_{s^s_k} + b^s_{a,k} + e^s_{a,k}.$$

This can be a good approximation, however, if there is high rotational motion and the accelerometer is far away from the origin, the accelerometer will experience additional forces. This motivates the user to estimate the rotational motion in the accelerometer measurement.

The benefit here is twofold. Firstly, the rotational motion does not influence the estimate of the specific force. Secondly, more information on the rotational motion becomes available.

Two additional forces act on the accelerometer, the Euler force and the Centrifugal force. The specific force is the same magnitude for all accelerometers. However, the rotational acceleration increases linearly with ρ, the displacement of an accelerometer from the origin. The angular acceleration $\dot{\omega}_k$ induces the Euler force and the angular velocity ω_k induces the Centrifugal force. The effects the forces on the IMU array are shown in Figure 3-1.

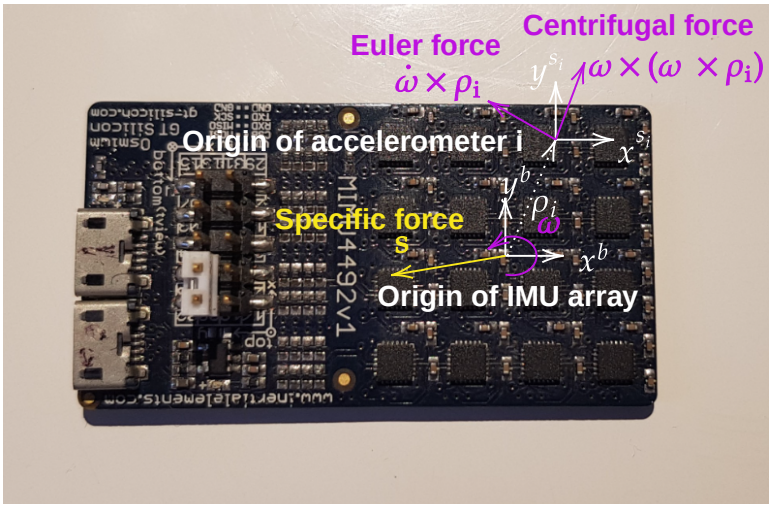

Figure 3-1: IMU array experiencing linear (yellow) and rotational (purple) motion, resulting in a specific force s, Centrifugal force $\omega \times (\omega \times \rho)$ and Euler force $\dot{\omega} \times \rho$. The angular velocity ω, angular acceleration $\dot{\omega}$ and relative position of an accelerometer ρ_i are shown. Recreated and adapted from [29].

Adding these forces to the accelerometer measurement equation gives

$$y_{a,k} = s_k^s + R^{sb}\left(\omega_k^b \times (\omega_k^b \times \rho) + \dot{\omega}_k^b \times \rho\right) + b_{a,k}^s + e_{a,k}^s, \tag{3-1}$$

where $\dot{\omega}_k^b$ denotes the angular velocity and ρ denotes the relative distance between origin of the accelerometer and the origin of the IMU array.

Measurement equation (3-1) is non-linear and has added extra terms that need to be estimated. This provides several challenges that need to be overcome as a single tri-axial accelerometer will not be able to estimate all the parameters as the problem is underdetermined. To understand how many accelerometers and gyroscopes are needed to get a unique solution, the identifiability conditions needed are identified, which will be shown in Section 3-2.

In literature, a distinction is made between accelerometer-only arrays and multiple IMU arrays. The identifiability conditions are shown for multiple IMU arrays is largely the same for accelerometer-only arrays. The main difference that separates the two is the ability to identify the sign of the angular velocity ω_k. This can be resolved by using a motion model prior for the angular velocity. The angular acceleration as a motion model will be shown in

Section 3-5. Additional methods have been developed to resolve the sign ambiguity and the reader is referred to [7] for an overview.

3-2 Identifiability Conditions of the Specific Force, Angular Velocity and Angular Acceleration

This section largely follow the work presented in [40]. For multiple IMUs the identifiability conditions of the angular acceleration is depends on the number of accelerometers [40] and their relative position ρ. The measurement models of the accelerometer and the gyroscopes are defined as [40]

$$y_{a,k} = s_k + \dot{\omega}_k \times \rho + \omega_k \times (\omega_k \times \rho) + e_{a,k}, \tag{3-2a}$$

$$y_{\omega,k} = \omega_k + e_{\omega,k}. \tag{3-2b}$$

The accelerometer measures both the specific force s and the angular acceleration $\dot{\omega}_k$ linearly. The accelerometer measures the angular velocity ω_k non-linearly, while the gyroscopes measure the angular velocity ω_k linearly. The angular velocity and angular acceleration cross products can be rewritten intro skew-symmetric matrices by introducing the cross-product matrix $[x\times]$. A real skew-symmetric matrix of size $n \times n$ has eigenvalues that are purely imaginary or zero, for a 3×3 skew-symmetric matrix this implies a rank of either 2 or 0 [11]. The skew-symmetric matrix is defined as

$$[x\times] \triangleq \begin{bmatrix} 0 & -x_3 & x_2 \\ x_3 & 0 & -x_1 \\ -x_2 & x_1 & 0 \end{bmatrix}. \tag{3-3}$$

The measurement equations can now be separated into a non-linear part $h(\omega_k)$ and a linear part H as seen below. The non-linear part holds information on the angular velocity while the linear part holds information on the angular acceleration and specific force.

In the formulation below, the relative positions of the accelerometers, denoted by ρ are time-invariant while the other variables θ_k and ω_k are time-varying and are estimated each time-step. The resulting expression is found below, (note that the timestep $*_k$ is left out for brevity)

$$\underbrace{\begin{bmatrix} y_{a_1} \\ \vdots \\ y_{a_N} \end{bmatrix}}_{Y_a} = \underbrace{\begin{bmatrix} [\omega\times]^2\rho_1 \\ \vdots \\ [\omega\times]^2\rho_N \end{bmatrix}}_{h_a(\omega)} + \underbrace{\begin{bmatrix} -[\rho_1\times] & \mathcal{I}_3 \\ \vdots & \vdots \\ -[\rho_N\times] & \mathcal{I}_3 \end{bmatrix}}_{H_a} \underbrace{\begin{bmatrix} \dot{\omega} \\ a \end{bmatrix}}_{\theta} + \underbrace{\begin{bmatrix} e_{a_1} \\ \vdots \\ e_{a_N} \end{bmatrix}}_{E_a}, \tag{3-4a}$$

$$\underbrace{\begin{bmatrix} y_{\omega_1}^T & \dots & y_{\omega_N}^T \end{bmatrix}^T}_{Y_\omega} = \underbrace{\begin{bmatrix} \omega^T & \dots & \omega^T \end{bmatrix}^T}_{h_\omega(\omega)} + \underbrace{\begin{bmatrix} e_{a_1}^T & \dots & e_{a_N}^T \end{bmatrix}^T}_{E_\omega}, \tag{3-4b}$$

$$\underbrace{\begin{bmatrix} Y_a \\ Y_\omega \end{bmatrix}}_{Y} = \underbrace{\begin{bmatrix} h_a(\omega) \\ h_\omega(\omega) \end{bmatrix}}_{h(\omega)} + \underbrace{\begin{bmatrix} H_a \\ H_\omega \end{bmatrix}}_{H} \theta + \underbrace{\begin{bmatrix} E_a \\ E_\omega \end{bmatrix}}_{E}. \tag{3-4c}$$

From this formulation several important conditions can be established. The first one, is that when no gyroscopes are available, the angular velocity can only be estimated up to a sign ambiguity. Secondly, through analysis of H, it can be determined how many accelerometers are needed in order for H to have full rank. By assuming that the position of the first accelerometer is at the origin the rank of H becomes

$$\text{rank}(H) = \text{rank}\left(\begin{bmatrix} 0_3 & \mathcal{I}_3 \\ -[\rho_2\times] & \mathcal{I}_3 \\ \vdots & \vdots \\ -[\rho_3\times] & \mathcal{I}_3 \\ \mathbf{0}_{3N,3} & \mathbf{0}_{3N,3} \end{bmatrix}\right) = 3 + \text{rank}\left(\begin{bmatrix} -[\rho_2\times] \\ \vdots \\ -[\rho_N\times] \end{bmatrix}\right). \tag{3-5}$$

It is evident that the rank of H is independent on the number of gyroscopes and solely depends on the location and the number of accelerometers placed. The locations of the accelerometers are in skew-symmetric form which has rank 2. Thus in order for matrix H to have full rank, a third accelerometer must be present that is not collinear with the other two placed sensors. The sufficient and necessary conditions for the identifiability conditions for the specific force, angular velocity and the angular acceleration are then

1. At least one gyroscope must be present.
2. At least three accelerometers are present whose locations are non-collinear.

Not only does the amount of sensors have influence on the identifiability conditions of the system, more accurate estimates can be found as well by using more sensors. For the gyroscopes this is rather simple as every gyroscope provides the same information (assuming they are of the same type). The co-variance of the estimation then scales linearly by $\frac{1}{N}$ by simply taking the average of the all the gyroscopes [29]. This simple solution is not the ideal method for the accelerometers, more sophisticated methods are presented in the next section.

3-3 Cramér-Rao Lower Bound of the Specific Force, Angular Velocity and the Angular Acceleration

The Cramér-Rao Lower Bound (CRLB) allows one to evaluate the performance of the estimator [27]. The CRLB is useful statistical tool that allows the user to identify how covariance of the measurement space translates to the covariance of the state space. More specifically, the CRLB gives the lower bound for all statistical unbiased estimators for a given parameter [10]. The CRLB is based on the Fisher Information Matrix (FIM) I. The FIM, as the name suggests tells us how much information can be obtained for the state space. The inverse of the Fisher information matrices show the information obtained. That is, the eigenvalues of the co-variance matrix are larger than the eigenvalues of the CRLB

$$\text{Cov}(x) \succeq I(x)^{-1}. \tag{3-6}$$

This section largely follows the work presented in [40].

Fisher Information Matrix The FIM is defined as follows [20], consider the linear measurement model with Gaussian noise $e \sim \mathcal{N}(0, \Sigma)$

$$y = Hx + e. \tag{3-7}$$

From this measurement equation the FIM is defined as

$$I(x) = E\left[\left(\frac{d\log(p(y|x))}{dx}\right)\left(\frac{d\log(p(y|x))}{dx}\right)^T\right], \tag{3-8a}$$

$$I(x) = H^T\Sigma^{-1}H. \tag{3-8b}$$

Information from multiple sensors is additive in the FIM [20]. This allows us to rewrite the stacked matrix into summation as

$$I(x) = \sum_{n=1}^{N} H_n^T\Sigma^{-1}H_n. \tag{3-9}$$

This can be applied to the measurement equation (3-4c) as well. First, the equation needs to be linearized around ω. The Jacobion for a single accelerometer and gyroscope pair of $h(\omega)$ is found to be [51]

$$J_{h,n} = \begin{bmatrix}([[\rho_n\times]\omega\times] - [\omega\times][\rho_n\times]) \\ \mathcal{I}_3\end{bmatrix}, \tag{3-10a}$$

$$H_n = \begin{bmatrix}-[\rho_n\times] & \mathcal{I}_3 \\ 0_3 & 0_3\end{bmatrix}. \tag{3-10b}$$

The information matrices of the accelerometer and gyroscopes can then be found as

$$I_n = \begin{bmatrix}J_{h,n}^T \\ H_n^T\end{bmatrix} Q_n^{-1}\begin{bmatrix}J_{h,n} & H_n\end{bmatrix}, \tag{3-11}$$

where Q_n is Gaussian white measurement noise of the specified accelerometer and gyroscope pair. $N_{max} = \max(N_{acc}, N_{gyr})$, such that every $J_{h,n}$ and H_n contains information of an accelerometer or gyroscope. If the $N_{acc} < N_{gyr}$ or $N_{acc} > N_{gyr}$ then the corresponding $J_{h,n}$ and H_n would be left empty as there is no sensor information. The formulation above can again be summed as information is additive, resulting in

$$I = \sum_{n=1}^{N_{max}} I_n = \sum_{n}^{N_{max}} \begin{bmatrix}J_{h,n}^T \\ H_n^T\end{bmatrix} Q_n^{-1}\begin{bmatrix}J_{h,n} & H_n\end{bmatrix}. \tag{3-12}$$

This formulation can then be expressed in its stacked equivalent as

$$I(\theta) = \begin{bmatrix}J_h^T \\ H^T\end{bmatrix} Q^{-1}\begin{bmatrix}J_h & H\end{bmatrix}, \tag{3-13}$$

The derivation of which was done by [40] under a few assumptions. The assumptions are as follows

1. The sensors are placed in a 2D plane with the origin at the center.

2. The measurements errors are uncorrelated
3. The accelerometers have the same error variance and gyroscopes have the same error variance.
4. The sensors are placed equidistant with distance α between two neighbouring sensors.
5. The number of sensors placed is a quadratic number

By finally taking the inverse of the FIM, the CRLB for the ML estimate of the angular velocity ω, angular acceleration $\dot{\omega}$ and the specific force s can be found. The results here are based solely on the measurement equation as given in (3-4c). The results of these are shown below in equations (3-14), (3-15) and (3-16).

CRLB of the specific force

$$I_s^{-1} = \frac{\sigma_a^2}{N_a}\mathcal{I}_3. \tag{3-14}$$

The previous equation shows that the accuracy of the information on linear acceleration obtained is inversely proportional to the number of accelerometers placed on the body. It also shows that the accuracy of the information obtained is independent on the geometry of the system.

CRLB of the angular velocity

$$I_\omega^{-1} \overset{(|\omega|>\gamma)}{=} \frac{6\sigma_a^2}{\alpha_a^2(N_a^2-N_a)} \begin{bmatrix} \frac{1}{\omega_x^2+\omega_y^2} & 0 & * \\ 0 & \frac{1}{\omega_x^2+\omega_y^2} & * \\ * & * & \frac{1}{2\omega_z^2} \end{bmatrix}. \tag{3-15}$$

It is of interest to know how the accelerometers estimate angular velocity, therefore, it is assumed that the gyroscopes are saturated: $|\omega| > \gamma$, where $\gamma \in \mathbb{R}$ is a positive constant and defined as the saturation threshold of the gyroscopes. Some off-diagonal entries have been left out for conciseness. It can then be seen that the information obtained about the rotational velocity is inversely proportional on the number of accelerometers placed squared N_a^2 and the distance between the accelerometers squared α_a^2. Finally, it can be observed that the information obtained is independent of the geometry of the system.

CRLB of the angular acceleration

$$I_{\dot{\omega}}^{-1} \overset{(\omega=0)}{=} \frac{12\sigma_a^2}{\alpha_a^2(N_a^2-N_a)} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \frac{1}{2} \end{bmatrix}. \tag{3-16}$$

The information achieved from the previous equation is most accurate when the angular velocity is zero. Furthermore, the information obtained for the angular acceleration is proportional to the distance between the accelerometers squared α_a^2 and the number of accelerometers placed squared N_a^2.

3-4 Fusing Multiple IMUs into a Synthetic IMU using Seperable Least Squares

The previous section discusses how much information can be obtained from the angular velocity ω, angular acceleration $\dot{\omega}$ and specific force s. The problem formulation of the stacked measurements is, however, partly non-linear. One approach to solving the non-linear problem is the separable least squares method. In this method, the non-linear part is separated from the linear part. The non-linear terms can be found using Gauss-Newton optimization, resulting in a weighted least squares solution for the remaining linear part. How the estimates can be calculated is summarized in Algorithm 5 in Appendix C, from the works [40].

Separable Least Squares (SLS) can be applied to a niche selection of non-linear equations, where part of the state vector enters the measurement vector [20]. This is the case for (3-4c) as is shown by [40], where the measurement equation is party linear. The non-linear part of the equation can enter the measurement vector and then WLS can be used to find the linear part. The methodology of SLS applied to (3-4c) is shown below.

First, the non-linear part can be fixed at ω^* such that a WLS expression can be found for the linear part θ. This is denoted by

$$Y = h(\omega^*) + H\theta(\omega^*) + E, \tag{3-17a}$$

$$\hat{\theta}(\omega^*) = (H^T Q^{-1} H)^{-1} H^T Q^{-1}(y - h(\omega^*)). \tag{3-17b}$$

Here, $\hat{\theta}$ is assumed to be stochastic and distributed as

$$\theta \sim \mathcal{N}(\hat{\theta}, (H^T Q^{-1} H)^{-1}), \tag{3-18}$$

and it is assumed that the stochastic error of $\hat{\theta}$ is independent in time [24]. The result of $\hat{\theta}$ can then be substituted back into Equation (3-4c), which results in

$$Y - h(\omega) = H\hat{\theta} + E. \tag{3-19}$$

The formulation above is frequently used in sensor array processing as the new covariance gives a better estimate of the covariance as it includes the structure of the array [5]. The covariance of the two stochastic sources are independent and additive, the expression can be expressed as

$$Y - h(\omega) \sim \mathcal{N}(H\hat{\theta}, H(H^T Q^{-1} H)^{-1} H^T + Q), \tag{3-20}$$

where the proof can be found in [41]. The corresponding Information matrix is given by [40] and is expressed as

$$P = Q^{-1} - Q^{-1} H (H^T Q^{-1} H)^{-1} H^T Q^{-1}. \tag{3-21}$$

The solution for $\hat{\omega}$ can then be expressed as [40]

$$\hat{\omega} = \arg\max_{\omega}(\mathcal{L}(\omega)), \tag{3-22}$$

where $\mathcal{L}(\omega)$ is defined as

$$\mathcal{L}(\omega, \hat{\theta}(\omega)) = -\frac{1}{2}||Y - h(\omega)||_P^2 + c, \tag{3-23}$$

where c is a constant remainder and negligible in the minimization operation. The expression of $\hat{\omega}$ can now take the form of

$$\hat{\omega} = \arg\min_{\omega} ||Y - h(\omega)||_P^2 \tag{3-24}$$

This problem formulation can be solved using a Gauss-Newton optimization. For the interested reader, the theory on Gauss-Newton optimization can be found in Section A-2 of Appendix A.

3-5 Angular Acceleration as a Motion Model

Previously, in Chapter 2, an orientation estimate for a single gyroscope and accelerometer pair was established. In this chapter, however, it has been shown that when multiple IMUs are used, they can be fused to not only obtain more accurate estimates of the angular velocity and specific force but also to estimate the angular acceleration. The angular acceleration can be used as a motion model to add additional information on the angular velocity.

The work presented in [40], which has been has been described in Section 3-4, describes a ML estimator which fuses the measurements in each individual time step. There is no relationship that couples the dynamics from one time step to the next time step. Based on their work in [40], the ML estimation of ω has been extended into a MAP estimation in [51] using motion models. The formulation presented in [51] takes a smoothing approach. This approach, however, would require all the data to be known and thus cannot be used in real time. The same motion models can also be applied as a one step ahead predictor.

The angular acceleration at time step k can now be used as a prior on the angular velocity at time step $k+1$. The motion model gives a prior of the angular velocity $\tilde{\omega}_{k+1}$ which takes the form of

$$\tilde{\omega}_{k+1} = \hat{\omega}_k + T\dot{\hat{\omega}}_k, \tag{3-25}$$

where

$$\hat{\omega}_k = \omega_k + e_{\omega,k}, \tag{3-26a}$$

$$\dot{\hat{\omega}}_k = \dot{\omega}_k + e_{\dot{\omega},k}, \tag{3-26b}$$

where $e_{\omega,k} \sim \mathcal{N}(0, \Sigma_\omega)$ and $e_{\dot{\omega},k} \sim \mathcal{N}(0, \Sigma_{\dot{\omega}})$ are Gaussian white noise. These two error sources can be assumed to be independent [51], this leads to the summation of their error giving the motion model as

$$\tilde{\omega}_{k+1} = \omega_k + \dot{\omega}_k + e_{\tilde{\omega},k}, \tag{3-27}$$

where $e_{\tilde{\omega},k} \sim \mathcal{N}(0, \Sigma_{\tilde{\omega}})$. In the current form, no motion models (random walk) models are used for the specific force and the angular acceleration, however, they can be incorporated in the same manner. Only motion model (3-27) is used such that the potential benefits of this motion model can be highlighted. Additionally to the increase in accuracy, [51] also has shown that the motion model allows identifiability conditions described in Section 3-2 to be altered to be free of any gyroscope. The sign ambiguity now disappears as prior information is available. This allows the method described in Section 3-4 for the extension towards accelerometer-only arrays.

3-6 Position Calibration of Accelerometers

The fusion method described in 3-4 does not assume any error in the relative locations. However, the relative spatial geometry between accelerometers must be known accurately for the optimal fusion of the measurements [40]. A geometric calibration for the relative positions, based on their work in [40], is shown in [8]. The calibration allows the user to estimate the relative positions with enough precision to make the estimation in Section 3-4 reliable [8]. The calibration algorithm is summarized below by making use of the measurement models as in (3-2a) and its stacked form (3-4c)

As initial position, the relative distances can be measured with e.g. a ruler and the first accelerometer is set at the origin. Afterwards, set an arbitrary accelerometer as the reference position ρ_1. The other positions are updated relative to this reference position. Furthermore, one can define the state x_k as

$$x_k = \begin{bmatrix} \omega_k^T & \dot{\omega}_k^T & {}_k^T \end{bmatrix}^T \tag{3-28}$$

for the unknown motion dynamics at every time step, this allows one to define the ML estimation problem as

$$\{\hat{x}_{1:K}, \hat{\rho}_{1:N}\} = \arg \max_{x_{1:K}, \rho_{1:N}} p(Y_{1:K} | x_{1:K}, \rho_{1:N}). \tag{3-29}$$

Assuming the motion dynamics to be independent in time the problem can be reformulated into

$$\{\hat{x}_{1:K}, \hat{\rho}_{1:N}\} = \arg \max_{\theta_\rho} \prod_{k=1}^{K} \max_{x_k} p(y_k | x_k, \rho_{1:N}), \tag{3-30}$$

such that x_k is found to maximize the probability of observing measurement y_k and θ_ρ to maximize the probability of observing all measurements $y_{1:K}$. A solution to this problem can be found by alternatively calculating and fixing $x_{1:k}$ and θ_ρ until a stopping criteria is satisfied. To find the solution to $x_{1:k}$ the solution from Section 3-4 can be used. Additionally the solution for θ_ρ can be found as follows

$$\rho_i = \left(\sum_{k=1}^{K} \Omega_k^T \Omega_k \right)^{-1} \left(\sum_{k=1}^{K} \Omega_k^T (y_{a_i,k} - s_k^b) \right). \tag{3-31}$$

The position calibration method is summarize in Algorithm 6 (from [8]) in Appendix C.

3-7 Multiple IMUs Orientation Estimation

An orientation estimation has been shown previously in Section 2-5, where use was made of a single gyroscope and accelerometer pair. In this chapter the orientation estimation is done in the same manner and data is pre-process in the same manner. This chapter introduced how multiple sensors can be fused together into a synthetic IMU, where the synthetic IMU is classified as the final estimates of the angular velocity and the specific force. Furthermore, it also allows for the estimation of the angular acceleration, which in turn can be used as a motion model prior to further improve the estimation of the angular velocity. An updated

orientation estimation scheme is shown in Figure 3-2 where the additional benefit the fusion of IMUs and the motion model are enclosed by the red rectangle.

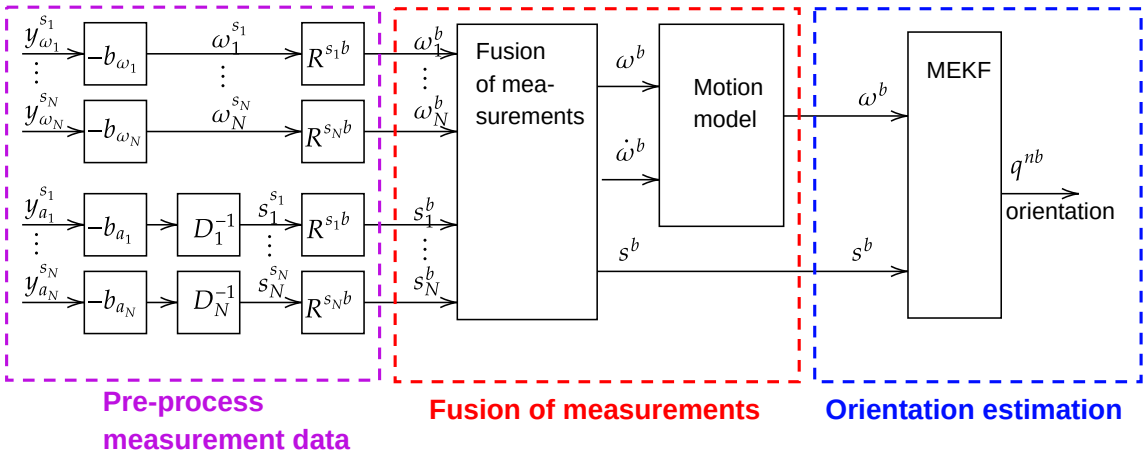

Figure 3-2: An orientation estimation scheme using multiple IMUs. In the purple rectangle, the measurement data is pre-processed after which in the red rectangle all the measurement are fused into a single synthetic IMU. The synthetic IMU can then estimate the orientation as seen in the blue rectangle

The theory presented in this chapter largely follows the works of [40], [51] and [8]. To the extent of the knowledge of the writer, this is the first time the theory is combined in such a way. Note that a similar scheme was used by [39], where they combinied the fusion of measurements and an EKF, but excluded the motion model. This chapter serves as a basis for the simulation setup in Chapter 4 and will be used to propose a novel parity space method for FDI as will be shown in Chapter 5.

Chapter 4

Multiple IMUs Simulation setup and results

In Chapter 3 it was introduced how multiple IMUs can be fused to form a synthetic IMU. It was also shown that additional information became available in the form of angular acceleration which was used as a motion model. The relationship between certain parameters and how much information can be obtained was also introduced. This chapter starts with Section 4-1 an introduction is given for the motivation of this chapter. This is followed by Section 4-2 where the first simulation setup is discussed. In the setup the Cramér-Rao Lower Bound (CRLB) is analyzed for the MIMU4844 array. Afterwards, in Section 4-3 a general procedure on how to generate measurement data is explained. This will be used as a basis for measurement data generation in the simulations. Afterwards, the results are shown and discussed in Section 4-4. An additional simulation setup and results are shown in Sections 4-5 and 4-6 respectively. Where different setups are tested on how accurate the angular velocity and specific force estimate is against an higher-grade IMU.

4-1 Introduction

In Chapter 3 it was shown how multiple IMUs can be fused. After fusion we had an estimated covariance of the specific force s, angular velocity ω, but also of the angular acceleration $\dot{\omega}$. How these covariances change according to different parameters was also shown. These are briefly summarized (with the assumption that $\omega = \omega_x = \omega_y = \omega_z$ for brevity) here as:

$$\underbrace{I_{\omega}^{-1} \overset{(|\omega|>\gamma)}{\propto} \frac{\sigma_a^2}{\alpha_a^2(N_a^2 - N_a)\omega^2}}_{\text{Angular velocity}}, \qquad \underbrace{I_{\dot{\omega}}^{-1} \overset{(\omega=0)}{\propto} \frac{\sigma_a^2}{\alpha_a^2(N_a^2 - N_a)}}_{\text{Angular acceleration}}, \qquad \underbrace{I_s^{-1} \propto \frac{\sigma_a^2}{N_a}}_{\text{Specific force}}, \tag{4-1}$$

where I denotes the FIM, σ_a denotes the variance of the accelerometer, γ denotes the saturation value for the gyroscopes, N_a denotes the number of accelerometers and α denotes the distance between two adjacent sensors. These analytical relationships were found by [40]

for an array configuration, where several assumptions were made. These assumptions are restated here:

1. The sensors are placed in a 2D plane with the origin at the center.
2. The measurement errors are uncorrelated
3. The accelerometers have the same error variance and gyroscopes have the same error variance.
4. The sensors are placed equidistant with distance α between two neighbouring sensors.
5. The number of sensors placed is a quadratic number

One of the limitations is the assumption that the sensors are placed in a 2D plane. Another limitation of the assumptions is that the sensors are placed equidistantly. These assumptions are violated for the MIMU4844 IMU array from Inertial Elements shown in Figure 1-3. A schematic representation of this IMU array is shown in Figure 4-1. In the IMU array, two layers of sensors are placed on top of each other. This makes the IMU array more compact and easier to use. However, as some assumptions are violated, the question is: does this have an influence on the covariance of the specific force, angular velocity and the angular acceleration? This will be shown in the results in Section 4-4.

Firstly, it was shown in 3 how multiple IMUs can be fused to form a synthetic IMU. It was also shown that an additional motion parameter can be estimated, the angular acceleration. The angular acceleration can be used as a motion model prior for the angular velocity, as shown in [51]. They also noted in their work that the angular acceleration and specific force can be modeled as a random walk model as well. In this chapter, different motion models are used in a simulation to compare a synthetic IMU against a higher-grade IMU.

There are two configurations for a synthetic IMU that will be tested against a higher-grade IMU. The first is the MIMU4844 IMU array, as shown in Figure 1-3. The distance between the sensors is very small for this configuration $\approx 5mm$. If the same sensors were instead scattered across the surface of a human arm (resulting in a distance of $\approx 100mm$ between sensors in the x-axis). This would result in more accurate estimations of the angular velocity. The second configuration represents a human arm and will be shown in Figure 4-6. The IMUs are placed equidistantly on the surface of a cylinder of length $0.3m$ and radius $0.05m$. The results will be shown in Section 4-6.

4-2 Simulation Setup to Estimate the CRLB of the MIMU4844 array

The simulation is based on the MIMU4844 IMU array, with 32 accelerometers and 32 gyroscopes from Inertial Elements. This is done to analyze the limitations of the assumptions made in [40] with the MIMU4844 IMU array.

Violation of previous assumptions The assumptions made in [40] that do not hold for the MIMU4844 IMU array are:

- The sensors are placed in a 2D plane with the origin at the center.
- The accelerometers have the same error variance and gyroscopes have the same error variance.
- The sensors are placed equidistant with distance α between two neighbouring sensors.
- The number of sensors placed is a quadratic number.

While the first and the last assumption can be met by only using one of the two layers as shown in Figure 1-3, in reality, one would want to use all the sensors, not part of the sensors.

Assumptions of this simulation For simulation purposes, the estimated positions for the accelerometers and estimated covariances for the gyroscopes and accelerometers were used.

- Accelerometer covariances are as in Figure B-4 and they have a mean of $\begin{bmatrix} 3.10 & 0 & 0 \\ 0 & 3.10 & 0 \\ 0 & 0 & 3.41 \end{bmatrix} 10^{-2}$ $\mathrm{m/s^2}$,
- Gyroscope covariances are as in Figure B-3 and they have a mean of $\begin{bmatrix} 8.40 & 0 & 0 \\ 0 & 8.70 & 0 \\ 0 & 0 & 6.20 \end{bmatrix} 10^{-3}$ rad/s,
- The positions of the IMUs with their respective numbers are shown in Figure 4-1.
- The biases and cross-axis inference can be compensated for properly and are, therefore, left out.
- The rotational velocity (unless varied) will be set at $\omega = \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T$ rad/s.

For the chosen angular velocity magnitude, this can be considered a fast movement for the arm as after 1 second the arm has rotated almost 60° around each axis.

Calibration The IMU array has 32 gyroscopes and 32 accelerometers each not yet calibrated. First, the gyroscopes are calibrated for the covariances and biases as discussed in Section 2-4-1. Afterwards, the accelerometers can be calibrated for as discussed in Section 2-4-2 for its bias, covariance and cross-axis interference. After all sensors have been calibrated, the calibration for the relative positions can be done as discussed in Section 3-6. The positions and coordinate frames of the IMUs are plotted in Figure 4-1. For the simulation, it is assumed that the calibrated positions are the true positions.

Configuration The real-life array, shown in Figure 1-3, consists of two layers of sensors on top of each other. In each layer, the axes of the sensors point in the same direction. However, the axes of the two layers do not. The locations, orientations of each sensor after calibration is illustrated in Figure 4-1.

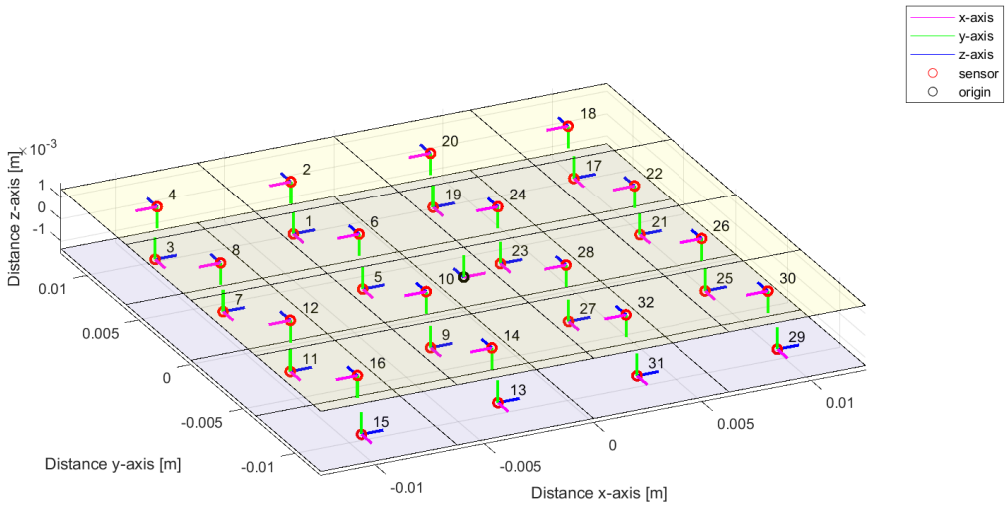

Figure 4-1: Locations and orientation of 32 accelerometers and gyroscopes of the MIMU4844 IMU array from Inertial Elements. The IMUS are located with a certain error on two layers on top of each other. The IMUs in each layer have their axes pointing in the same direction.

Each IMU had a slight offset in the orientation and a misalignment error. It is assumed that it is known a priori how the IMUs are rotated and what the misalignment error is. More specifically, the rotation matrices $R^{s_i b}$ and $R^{s_j b}$ for $i = 1 : N_a$ and $j = 1 : N_\omega$ are known a priori with sufficient accuracy and are, therefore, neglected in the simulation.

Varying parameters Several different varying parameters were used for the simulation. From four parameters, the results of four are shown in Section 4-4. These are

- Varying number of accelerometers,
- Varying relative distance between sensors,
- Varying displacement of the origin.
- Varying magnitude of the angular velocity.

For all simulations, a Monte Carlo simulation was run with 10^3 realizations to approximate the ML CRLB. The results of the Monte Carlo simulation are also shown in every plot.

4-3 Generating Simulation Measurement Data for Multiple IMUs

As is previously shown, the measurement equation of the gyroscope depends on the angular velocity and Gaussian white noise. The accelerometer, in turn depends on the angular velocity, the angular acceleration, the specific force and also Gaussian white noise. The measurement equations (3-2a) are repeated here for clarity

$$y_{a,k} = s_k + \dot{\omega}_k \times \rho + \omega_k \times (\omega_k \times \rho) + e_{a,k},$$
$$y_{\omega,k} = \omega_k + e_{\omega,k}.$$

The angular velocity and angular acceleration is the same for every accelerometer, and has to be calculated only once. By introducing the term Ω_k as

$$\Omega_k^b = [\dot{\omega}_k^b \times] + [\omega_k^b \times]^2, \tag{4-3}$$

where $i = 1 : N_a$ accelerometers $j = 1 : N_\omega$ gyroscopes. .The measurement equations of the accelerometers and gyroscopes can be expressed as

$$y_{a_i,k} = R^{s_i b}\left(\left(a_k^b - R_k^{bn} g^n\right) + \Omega_k^b \rho_i\right) + b_{s_i k}^{s_i} + e_{a_i,k}^{s_i}, \tag{4-4a}$$
$$y_{\omega_j,k} = R^{s_j b}\omega_k^b + b_{\omega_j}^{s_j} + e_{\omega_j,k}^{s_j}. \tag{4-4b}$$

The sensors can be placed at any location but also in any orientation. The sensor frames and body frame do not necessarily align. For the gyroscope measurement equation, the angular velocity is expressed in the body frame and is measured in the sensor frame. Afterwards, Gaussian white noise and a bias can be added. For the accelerometers, the angular motion can be grouped together into a single term Ω_k^b. The linear motion is the combination of the linear acceleration and the gravity vector. The linear acceleration can be set in the body frame, however, the gravity vector is constant in the navigation frame and must be rotated towards the body frame using rotation matrix R_k^{bn}.

The rotation matrix can be found from the quaternion q^{nb}, which keeps track of the rotation from body to navigation frame as

$$R_k^{bn} = \left[q_v q_v^T + q_0^2 \mathcal{I}_3 + 2q_0[q_v \times] + [q_v \times]^2.\right] \tag{4-5}$$

The quaternion can be updated (note that the quaternion after the update will be at unit length) for every time step as

$$q_k^{nb} = \exp_q(\frac{\omega_k^b}{2}) \odot q_{k-1}^{nb}. \tag{4-6}$$

The linear acceleration and the rotational velocity in the body frame can now be set as desired. The angular acceleration can be found by taking the discrete time derivative of the angular velocity as

$$\dot{\omega}_k^b = \frac{\omega_{k+1}^b - \omega_k^b}{T}. \tag{4-7}$$

The assumptions are made that the bias is constant and that noise is additive. It is further assumed that the relative accelerometer locations $\rho_{1:N_a}$, biases and all rotation matrices $R^{s_i b}$ and $R^{\omega_j b}$ are known a priori. The generation of simulation data is summarized in Algorithm 4 in Appendix C.

4-4 Results on the Covariance of the Specific Force, Angular Velocity and Angular Acceleration of the MIMU4844 Array

In this section the results are shown of the covariance of the angular velocity, the angular acceleration and the specific force for a varying number of accelerometers, a varying relative distance between sensors and a varying displacement of the origin.

4-4-1 Varying Number of Accelerometers

The plots for the number of accelerometers below are shown from 3 to 32 accelerometers as that is the minimum amount of accelerometers necessary to comply to the identifiability conditions as discussed in 3-2. There were 32 gyroscopes used and the angular velocity was set to $\omega = \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T$ rad/s. The results are plotted in Figure 4-2. The relationship in (4-1) is repeated for the number of accelerometers as

$$\underbrace{I_{\omega}^{-1} \propto \frac{1}{N_a^2}}_{\text{Angular velocity}} , \qquad \underbrace{I_{\dot{\omega}}^{-1} \propto \frac{1}{N_a^2}}_{\text{Angular acceleration}} , \qquad \underbrace{I_{s}^{-1} \propto \frac{1}{N_a}}_{\text{Specific force}} . \tag{4-8}$$

Angular velocity The trace of the covariance of the angular velocity is plotted on the left of Figure 4-2. The analytical expression (blue) is verified with a MC approximation (yellow). They are both compared to trace of the covariance of the mean of all the gyroscopes (purple). From the relationship shown in (4-8), the covariance trace of the angular velocity is expected to decrease quadratically. The trace of the covariance in Figure 4-2 does not follow this expectation. Instead, the covariance remains almost constant and decreases slightly when more sensors are placed. This can be explained by that the information provided by the gyroscopes outweighs the information provided by the accelerometers. As more accelerometers are added, add a small amount of information. This shows that at this angular velocity and this configuration the mean of the gyroscopes provides roughly the same accuracy on the estimate of the angular velocity.

Angular acceleration In the middle of Figure 4-1, the trace of the angular acceleration covariance (blue) is plotted and verified with a MC approximation (yellow). As can be seen in (4-8), the angular angular acceleration covariance is expected scale proportionally with N_a^2. However, the simulation shows that the angular acceleration scales with N_a rather than N_a^2. This can be attributed to the configuration of the array. The sensors are placed equidistantly in the xy-plane, but they are placed in two planes with small separation in the z-direction. The assumption that the relative distance is the same between sensors does not hold for this IMU array. If the angular acceleration were to be used as a motion model prior, it would contain less information.

Specific force As shown in the right of Figure 4-2, the trace of the covariance of the specific force is plotted. The analytical covariance (blue) is verified with a MC approximation (yellow) and compared to averaging out the accelerometer measurements. By looking at (4-8), we

would expect the analytical covariance to coincide with averaged covariance. It can be seen that the three lines overlap and thus (4-8) shows the expected relationship. A small difference between the analytical and the averaged covariance can be observed at 3 accelerometers. This can be explained by the displacement of the origin. As only 3 accelerometers are placed, the origin is not in the same place as the centre of these 3 accelerometers. This effect will be elaborated on in more detail in Section 4-4-3.

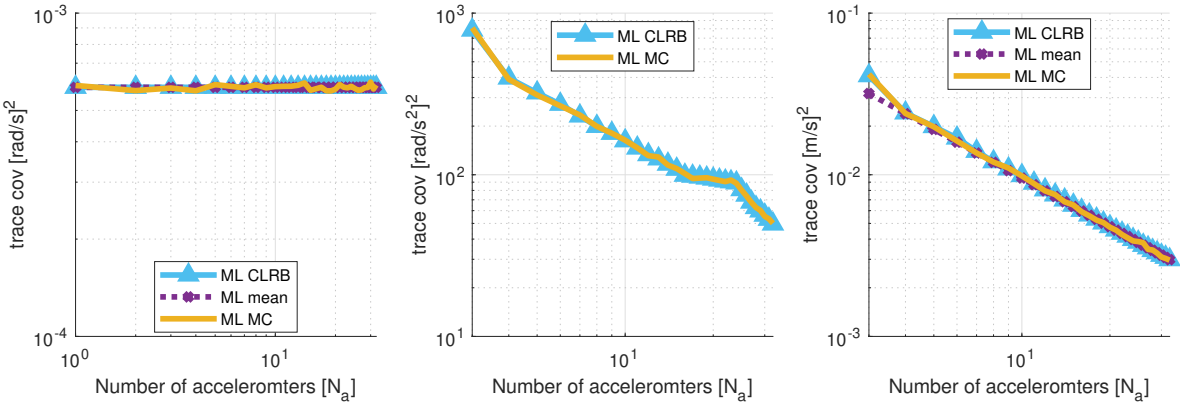

Figure 4-2: Angular velocity covariance (left) Angular acceleration covariance (mid) and Specific force covariance (right) plotted against number of accelerometers N_a. All 32 gyroscopes were used and the angular velocity was set at $\omega = \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T$ rad/s and the IMU locations are based on the MIMU4844 array also shown in Figure 4-1.

4-4-2 Varying Relative Distance Between Sensors

In this simulation, 32 accelerometers and 32 gyroscopes were used at $\omega = \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T$ rad/s. The relative distance was increased with a scale factor from 1 to 10^3 times ρ. This was done to preserve the structure of the array and only the relative distance between the sensors is increased. The results are plotted in Figure 4-3. The relationship in (4-1) is repeated for the distance between sensors as

$$\underbrace{I_{\omega}^{-1} \propto \frac{1}{\alpha_a^2}}_{\text{Angular velocity}} \quad , \qquad \underbrace{I_{\dot{\omega}}^{-1} \propto \frac{1}{\alpha_a^2}}_{\text{Angular acceleration}} \quad . \tag{4-9}$$

Angular velocity In Figure 4-3, the analytical covariance trace (blue) of the angular velocity is plotted and verified with a MC approximation (yellow). As a baseline, the averaged covariance of the gyroscopes (purple) is also plotted. From (4-9) we would expect that the angular velocity decreases quadratically as the scale factor increases. Initially, the analytical covariance is almost the same as the averaged covariance. This is because, at this magnitude of the scale factor, the information provided by the gyroscopes outweighs the information provided by the accelerometers. As the magnitude of the scale factor increases, the accelerometers can add more information to the angular velocity estimate. Eventually, at higher scale factors, the covariance decreases quadratically. At lower angular velocities, the gyroscopes are dominant in providing information.

Angular acceleration The analytical covariance trace (blue) of the angular acceleration and a MC approximation (yellow) are plotted in the middle of Figure 4-3. By looking at (4-9) we expect that the covariance decreases quadratically with the scale factor. The plotted results show the same trend. at higher magnitudes As was stated in the results of the angular velocity, it also decreases quadratically with the scale factor. This affects the information that the angular acceleration as a motion model prior can have. More specifically, increasing the distance between the accelerometers increases the information gained from the angular acceleration as a motion model prior.

Specific force On the right in Figure 4-3 the trace of the analytical covariance trace of the specific force covariance (blue) is plotted with a MC approximation (yellow). Additionally, the averaged covariance trace of all accelerometers is plotted (purple) as a baseline. The expectation from (4-9) is that the specific force remains the same as the baseline. The results coincide with the expectation from (4-9).

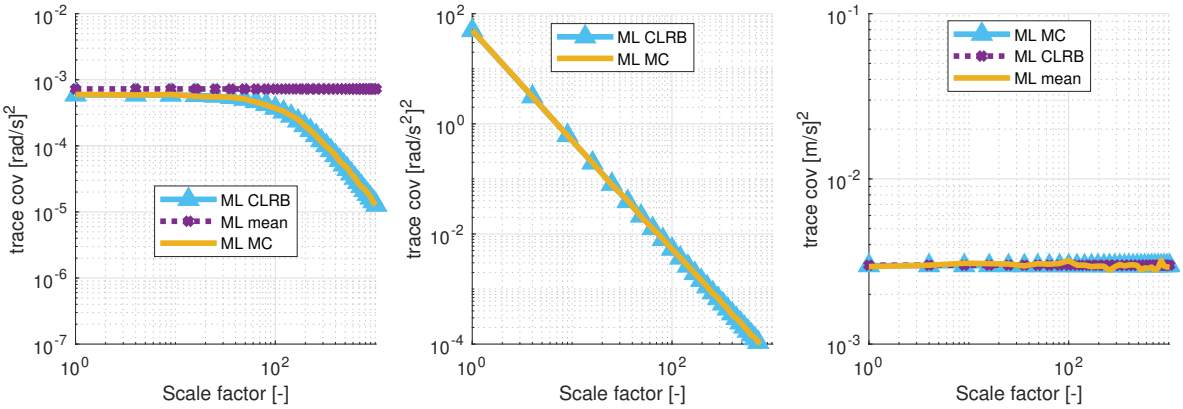

Figure 4-3: Angular velocity covariance (left) Angular acceleration covariance (mid) and Specific force covariance (right) plotted against distance between sensors. The base MIMU4844 array locations with 32 accelerometers and 32 gyroscopes is as in Figure 4-1. The array is scaled up with a scale factor from 1 to 10^3 the original configuration size. The angular velocity was set at $\omega = \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T$ rad/s.

4-4-3 Varying Displacement of the Origin

This simulation again has 32 accelerometers, 32 gyroscopes, an angular velocity of $\omega = \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T$ rad/s and the array structure. In this case, however, the origin was displaced, i.e. where synthetic IMU measurements are estimated. The results are shown in Figure 4-4. The relationship in (4-1) does no include the varying displacement. We would expect that changing this origin will not affect the covariance of the specific force, angular velocity or angular acceleration.

Angular velocity The trace of the angular velocity covariance is plotted in Figure 4-4. The analytical covariance (purple) is verified numerically with a MC approximation (yellow). A baseline is added from the averaged covariance (purple) of the gyroscopes. The covariance did not change when displacing the origin of the synthetic IMU as was expected.

Angular acceleration In Figure 4-4 the trace of the angular acceleration is plotted in the middle. Both the analytical covariance (blue) and a MC approximation (yellow) are shown. As expected, the covariance of the angular acceleration remains constant under changing the location of the synthetic IMU.

Specific force On the right of Figure 4-4, the specific force covariance trace is shown. The analytical covariance (blue) is verified with a MC approximation (yellow) and compared to a baseline. The baseline is the averaged covariance (purple) of the accelerometers. We expected the covariance to remain unchanged under shifting the origin of the synthetic IMU. As can be observed, however, the covariance starts to grow quadratically with the displacement of the origin. This does not seem problematic at first glance as the origin of the synthetic IMU can be chosen arbitrarily, so the most accurate would be in the centre. However, if these locations were to change, or sensors would be momentarily removed (as will be the case in Chapter 6). The uncertainty increase may result in a performance loss of orientation estimation. This effect was observed in future results (see Figure B-8) and are discussed in Section 6-2.

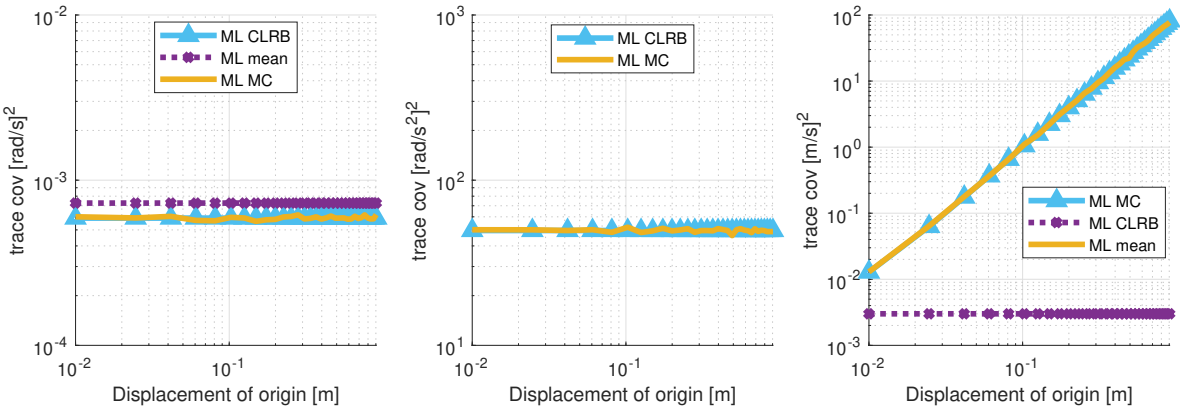

Figure 4-4: Angular velocity covariance (left) Angular acceleration covariance (mid) and Specific force covariance (right) plotted the displacement of the origin. The base angular velocity was set at $\omega = \begin{bmatrix}1 & 1 & 1\end{bmatrix}^T$ rad/s. The IMU locations are based on the MIMU4844 array also shown in Figure 4-1 with 32 accelerometers and 32 gyroscopes.

4-4-4 Varying Magnitude of the Angular Velocity

For the final simulation, 32 accelerometers and 32 gyroscopes were used with the standard array configuration. The angular velocity was changed from $\omega = \begin{bmatrix}1 & 1 & 1\end{bmatrix}^T$ rad/s to $\omega = 10^4\begin{bmatrix}1 & 1 & 1\end{bmatrix}^T$ rad/s. The results are shown in Figure 4-5. The relationship in (4-1) is repeated for the distance between sensors as

$$\underbrace{I_\omega^{-1} \propto \frac{1}{\omega^2}}_{\text{Angular velocity}} . \tag{4-10}$$

Angular velocity In the left of Figure 4-5, the trace of the covariance is plotted. The analytical covariance (blue) is verified with a MC approximation (yellow) and compared to the averaged gyroscope covariance (purple) as a baseline. At low angular velocities and using the base configuration of the MIMU4844 array, the analytical covariance is slightly below the baseline. However, as the magnitude of the angular velocity increases (from 10^2 rad/s and higher), the covariance of the angular velocity decreases quadratically. This result is in expectation with (4-10). At these higher angular velocities, the accelerometers provide more information than the gyroscopes. This verifies the use of accelerometer-only arrays in highly dynamic environments. Note that the gyroscope still provides useful information. Adding at least one gyroscope in an accelerometer-only array helps resolve the sign ambiguity.

Angular acceleration The covariance trace of the angular acceleration is plotted in the middle of Figure 4-5. The analytical covariance (blue) is plotted and verified with a MC approximation (yellow). In (4-10) we expected that the angular acceleration covariance would remain constant, however, the assumption was made that the angular velocity was zero. At higher angular velocities, the covariance of the angular acceleration increases slightly, the impact is rather insignificant. More importantly, the angular velocity covariance decreases quadratically while the angular acceleration covariance slightly increases. This impacts the angular acceleration as a motion model prior for the angular velocity. At lower angular velocities, the angular acceleration motion model is more potent than at higher angular velocities.

Specific force In Figure 4-5 the trace of the covariance of the specific force is shown on the right. The analytical covariance (blue) is verified with a MC approximation (yellow) and compared to the averaged covariance (purple) of the accelerometers. We expect that the specific force covariance remains invariant under changing the angular velocity and this is verified in the result.

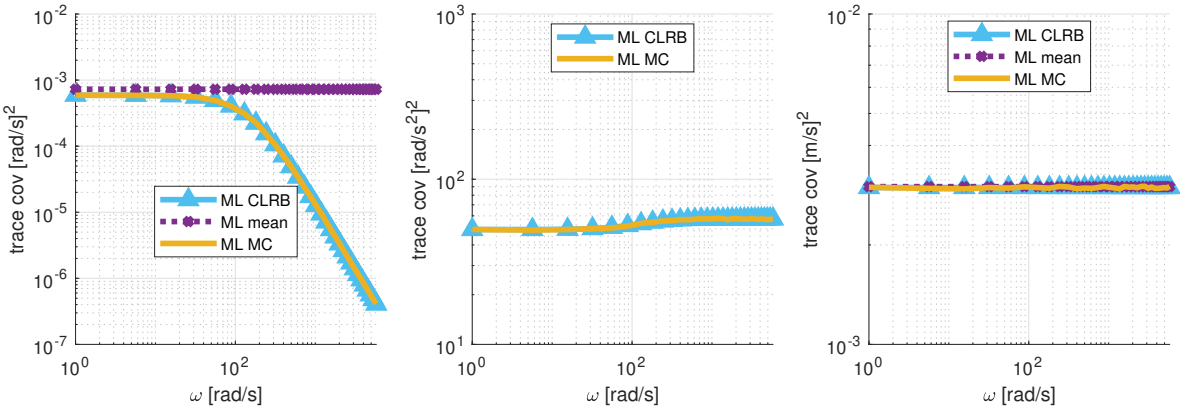

Figure 4-5: Angular velocity covariance (left) Angular acceleration covariance (mid) and Specific force covariance (right) plotted against angular velocity ω. The base angular velocity was set at $\omega = 10^0 \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T$ rad/s and increased to $\omega = 10^4 \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T$ rad/s . The IMU locations are based on MIMU4844 array also shown in Figure 4-1 with 32 accelerometers and 32 gyroscopes.

4-5 Simulation Setup to Compare Different Setups Against a Higher-grade IMU

One of the questions of this thesis is how do multiple lower-grade IMUs compare against a higher-grade IMU. To provide an answer to this question, first it must be addressed that there are many different factors that influence the answer. For this reason some assumptions have to be made. The choices and assumption made are shown below. Note that measurement data was simulated as described in Algorithm 4 in Appendix C.

Different Accelerometer Configurations The first configuration that will be analyzed is the MIMU4844 IMU array from Inertial Elements. On this configuration all 32 IMUs will be used with their locations taken from their position calibration, see Figure 4-1 a reference. Another configuration to be analyzed is a cylinder the size of a human arm segment where the same IMUs are used. The radius of the cylinder is 0.05 m and the length of the cylinder is 0.3m, to resemble the size of a human arm. The accelerometers placed in four rings with 8 sensors. The rings are placed equidistantly in the x-axis, so that the entire length of the cylinder is used. The locations of the accelerometers can be seen in Figure 4-6.

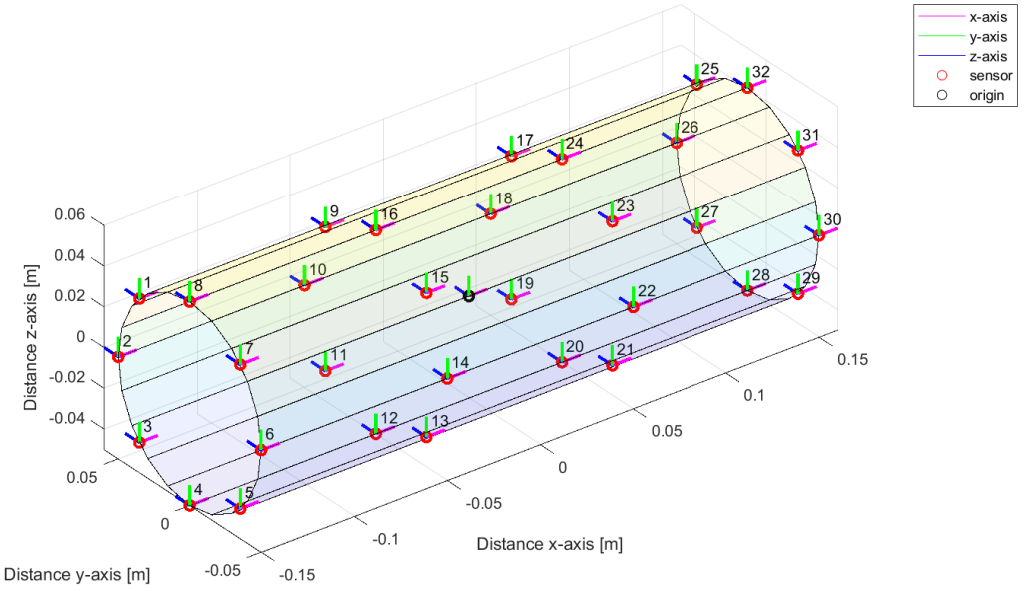

Figure 4-6: Locations and orientation of 32 accelerometers and gyroscopes. The IMUs are placed on the surface of the cylinder and have their sensitive axes point in different directions.

Different Angular Velocity Magnitudes As was shown in the results in Section 4-4-4, the covariance of the angular velocity depends on on the magnitude of the angular velocity, that is, the angular velocity covariance decreases as the angular velocity magnitude increases. The least amount of information that can be obtained from the accelerometers on the angular velocity is when the angular velocity is zero. The angular velocity will be set from

$10^0 \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T$ rad/s to $10^4 \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T$ rad/s.

Different Choices in Motion Models There are different choices to make when it comes to using motion models as a prior or not. These already have been discussed in Sections 2-2 and 3-5. In this simulation two models are assumed for the specific force and the angular acceleration; no motion model and a random walk model. The angular velocity can take no motion model, a random walk model or make use of the angular acceleration. In summary, these take the form of

$$s_{k+1} = s_k + e_{s,k}, \tag{4-11a}$$
$$\omega_{k+1} = \omega_k + e_{\omega,k}, \tag{4-11b}$$
$$\omega_{k+1} = \omega_k + \dot{\omega}_k + e_{\dot{\omega},k}, \tag{4-11c}$$
$$\dot{\omega}_{k+1} = \dot{\omega}_k + e_{\dot{\omega},k}. \tag{4-11d}$$

Constants Throughout the Simulation The covariance matrices of the accelerometers and gyroscopes are all set equal to mean of the covariances found experimentally of the MIMU4844 array. The gyroscope and accelerometer covariance matrices are

$$\Sigma_{\omega_j} = \begin{bmatrix} 8.40 & 0 & 0 \\ 0 & 8.70 & 0 \\ 0 & 0 & 6.20 \end{bmatrix} 10^{-3}, \qquad \Sigma_{a_i} = \begin{bmatrix} 3.10 & 0 & 0 \\ 0 & 3.10 & 0 \\ 0 & 0 & 3.41 \end{bmatrix} 10^{-2}. \tag{4-12}$$

For $i = 1:32$ accelerometers and $j = 1:32$ gyroscopes. The sampling rate is set at 100 Hz. No biases or misalignment errors were added as the assumption is made these can be compensated for. It is also assumed that the position of the accelerometers can be calibrated properly and hence is known with high enough accuracy. The motion model covariances were chosen to be equal to the estimated covariances. In reality, the motion model covaraince are often used as tuning parameters and are likely taken to be higher. However, for this simulation we assume that we know them with sufficient accuracy.

Higher-grade IMU As comparison, the MTi 300-series from Xsens [1] was chosen as a higher-grade IMU. The noise covariance for the gyroscope and accelerometer were found experimentally at 100 Hz as

$$\Sigma_{\omega_X} = \begin{bmatrix} 2.096 & 0 & 0 \\ 0 & 2.076 & 0 \\ 0 & 0 & 1.474 \end{bmatrix} 10^{-6}, \qquad \Sigma_{a_X} = \begin{bmatrix} 3.335 & 0 & 0 \\ 0 & 3.323 & 0 \\ 0 & 0 & 2.470 \end{bmatrix} 10^{-5}. \tag{4-13}$$

The higher-grade IMU was simulated without motion model for both the gyroscope and the accelerometer. Note that in reality, a motion model for the angular velocity and specific force can be chosen for the higher-grade IMU. This would further increase the accuracy of the higher-grade IMU.

4-6 Results on Comparing Different Configurations Against a Higher-grade IMU

This section compares two different configurations against a higher-grade IMU. First, the MIMU4844 array is shown after which a cylinder the size of a human arm is shown. The angular velocity RMSE and specific force RMSE are compared for different motion models as mentioned in Section 4-5. Two figures are shown in which the x-, y- and z-axis angular velocity RMSE are compared against the higher-grade IMU. Two tables are presented in which a performance comparison for the angular velocity and the specific has been done on the z-axis. The corresponding Tables B-1, B-2, B-3 and B-4 for the x- and y-axis are almost equal and are shown in Appendix B.

4-6-1 MIMU4844 Array configuration

In Figure 4-7 the angular velocity RMSE of the MIMU4844 configuration is plotted. The x- (left), y- (middle) and z-axis (right) are shown and the gyroscope covariance of the Xsens MTi-300 IMU is added as a baseline. From the results in Section 4-4, we would expect that the angular acceleration motion models perform the best. In this configuration, however, the angular velocity motion models perform the best. This is because the process noise covariances were taken to be equal to the estimated covariances and due to the gyroscopes present, the angular velocity covariance is low. At this distance between the sensors, the angular acceleration covariance is too high to provide any information as a motion model. Note that in reality, however, the process noise is usually treated as a tuning parameter and the angular acceleration will likely give a better estimate. The angular velocity RMSE at lower angular velocities does not outperform the higher-grade IMU from Xsens. If the angular velocity gets higher than a threshold, roughly $\omega \approx 2000 \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T$ rad/s, the IMU array, with and without motion models start to outperform the Xsens gyroscope sensor.

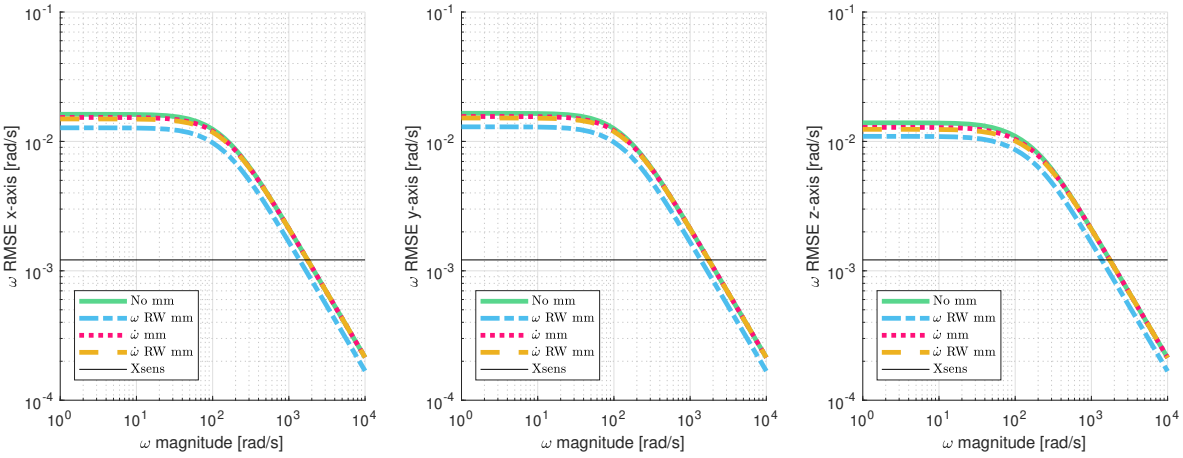

Figure 4-7: Synthetic IMU angular velocity covariance after fusion of the MIMU4844 IMU array of Inertial Elements. The covariances of the x-axis (left) y-axis (mid) and z-axis (right) using different motion models are shown, which are compared to MTi 300 IMU from Xsens.

Table 4-1 contains the RMSE (angular velocity and specific force) ratios of the synthetic IMU

with respect to the higher-grade IMU. There were 32 accelerometers used, which is the only influence on the specific force RMSE. the ratio of the RMSE with respect to the high-grade IMU can be seen in Table 4-1. It is interesting to note that the angular velocity random walk is the lowest at $\omega = 10^4 \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T$ rad/s. The reason for this has been discussed in the paragraph above. From the results in 4-4 we would have expected that the angular acceleration motion models perform better at lower velocities, which is verified by the results. This is shown by that the angular acceleration motion models have the same ratio as no motion model at high velocities. The specific force RMSE ratio is constant as only adding more accelerometers would decrease the ratio.

motion model	Angular velocity magnitude in x- y- and z- axis [rad/s]				
Angular velocity	$\omega = 10^0$	$\omega = 10^1$	$\omega = 10^2$	$\omega = 10^3$	$\omega = 10^4$
ω no motion model	11.465	11.434	9.035	1.689	0.175
ω random walk	9.013	8.989	7.103	1.328	0.138
$\dot{\omega}$ motion model	10.58	10.555	8.539	1.684	0.175
$\dot{\omega}$ random walk	10.213	10.19	8.313	1.682	0.175
Specific force					
s no motion model	6.568				
s random walk	5.164				

Table 4-1: Performance comparison different motion models of the MIMU4844 IMU array of Inertial Elements. The performance is indicated by the z-axis ratio of the RMSE with respect to the MTi 300 IMU from Xsens.

4-6-2 Multiple IMU Cylinder Configuration

The x- (left), y-(mid) and z-axis (right) RMSE error velocities are plotted in Figure 4-8. It can be observed that at lower angular velocities, the motion models using the angular acceleration outperform the angular velocity random walk model. This is to be expected and the reason is twofold as was shown in 4-4. The accelerometers provide the least amount of information on the angular velocity at lower angular velocities while providing more information on the angular acceleration. It is then interesting to note that as the angular velocity magnitude gets higher, the angular velocity random walk model starts to outperform the angular acceleration motion models. This is because the process noise covariances were taken to be equal to the estimated covariances. As the angular velocity increases, the accelerometers start providing more information on the angular velocity estimate. In reality, however, we would expect that the angular acceleration motion models still perform better even at higher angular velocities. Finally, it can be seen that even before $\omega \approx 300 \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T$ rad/s, the synthetic IMU starts to outperform the Xsens gyroscope accuracy.

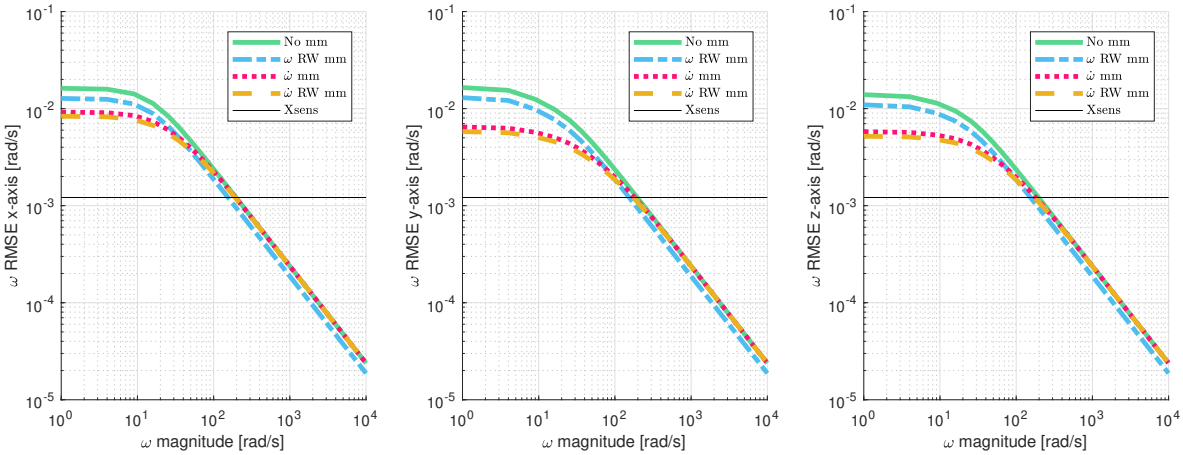

Figure 4-8: Synthetic IMU angular velocity covariance after fusion of 16 IMUs placed on a cylinder. The covariances of the x-axis (left) y-axis (mid) and z-axis (right) using different motion models are shown, which are compared to MTi 300 IMU from Xsens.

Table 4-1 contains the RMSE (angular velocity and specific force) ratios of the synthetic IMU with respect to the higher-grade IMU. At this configuration, it is apparent that the angular acceleration motion models provide more information as a motion model prior on the angular velocity than at a smaller configuration. To compare, at $\omega = \begin{bmatrix}1 & 1 & 1\end{bmatrix}^T$ rad/s the angular acceleration motion model here has a ratio of 4.763 where the MIMU4844 IMU configuration the ratio was 10.58. This coincides with the expectation from the results in 4-4. Similar to the MIMU4844 IMU configuration, the angular velocity random walk model performs the best at higher angular velocities. For the same reason as mentioned there, we would expect that in real life the angular acceleration motion models still perform better.

motion model	Angular velocity in x- y- and z- axis [rad/s]				
Angular velocity	$\omega = 10^0$	$\omega = 10^1$	$\omega = 10^2$	$\omega = 10^3$	$\omega = 10^4$
ω no motion model	11.465	9.399	1.939	0.192	0.02
ω random walk	9.013	7.39	1.524	0.151	0.015
$\dot{\omega}$ motion model	4.763	4.403	1.614	0.191	0.02
$\dot{\omega}$ random walk	4.276	3.966	1.523	0.191	0.02
Specific force					
s no motion model	6.568				
s random walk	5.164				

Table 4-2: Performance comparison different motion models of the IMU cylinder. The performance is indicated by the ratio of the z-axis RSME with respect to the MTi 300 IMU from Xsens

4-7 Concluding Remarks

In Chapter 3 it was shown how multiple IMUs can be fused and how the angular acceleration can be used as a motion model to obtained estimates of the angular velocity and the specific

force. In this chapter, we simulated two different configurations, the MIMU4844 array and a cylinder the size of a human arm. The CRLB was analyzed for the cylinder. Afterwards, the angular velocity and the specific force of the two configurations were compared against a higher-grade IMU. The relevant results are summarized below.

Influence of shifting the origin When the synthetic IMU shifts away from the origin, the covariance of the specific force starts to increase quadratically. For applications where all the sensors are used this probably will not pose a problem. This is because the origin of the synthetic IMU can be chosen arbitrarily. It might pose a problem, however, when FDI is applied. When all the sensors on one side are assumed false and thus isolated, the origin is not in the same place as the centre of the remaining accelerometers. This will lead to a loss in the accuracy of the specific force estimate. This was experienced during the simulation in Chapter 6 and the implications are elaborated on in that chapter.

Influence of the configuration The second interesting effect is that due to the configuration of the array, the covariance of the angular acceleration when adding more sensors scales linearly rather than quadratically. The same effect holds for the angular velocity, although this effect is not visible in Figure 4-2 as the accelerometers add negligible information. This shows a 2D array is not the best configuration for estimating rotational motion. If the accelerometers were placed in a 3D array, where the accelerometers are equidistant in x- -y and z-direction, the covariance would start to decrease quadratically. This suggests that is is more important to optimize the distance between sensors in all directions rather than in a 2D plane. For example, instead of placing two layers of 16 IMUs on top of each other, more information can be obtained by placing 4 layers of 8 IMUs on top of each other.

Angular acceleration as a motion model A third interesting result regards the angular acceleration that can be used as a motion model prior for the angular velocity. The angular acceleration and the angular velocity both decrease quadratically with the distance between sensors squared. As both covariances decrease at the same rate, the ratio between the two remains constant. This suggests that the amount of information that the angular acceleration provides is independent of the configuration. While this is true for accelerometer-only arrays, the gyroscopes also provide information on the angular velocity. Starting around 10 times the scale factor, the accelerometers add enough information so that the angular velocity covariance starts to decrease. Before this point, the angular acceleration as a motion model will also provide less information. This shows that the benefit of bigger configurations is twofold. It is also apparent that as the angular velocity magnitude increases, the angular velocity covariance decreases while the angular acceleration covariance slightly increases. This shows that at lower angular velocities, the angular acceleration provides more information on the angular velocity as a motion model prior. This result has also been verified in [51].

Specific force As the RMSE of the specific force of the IMU array and the cylinder when using a random walk model are about 6.568 times and 5.164 times larger than the accelerometer from Xsens. More accelerometers ($\approx$ 43 and $\approx$ 27 times more) are needed to get the

same level of accuracy. Placing that many accelerometers in an array or on a human arm segment does not seem feasible. Alternatively, more accurate accelerometers can to be used.

Angular velocity The RMSE of the angular velocity estimate, however, can compete with the higher-grade IMU. As the angular velocity increases, the accelerometers are also able to provide information on the estimate. The results show that using more IMUs allow for a good estimate of the angular velocity. The motion models add a lot more information at lower angular velocities, in particular when use is made of the angular acceleration. Furthermore, at higher angular velocities a high accuracy estimate of the angular velocity can be obtained. This automatically extends to accelerometer-only arrays as the information of the accelerometers on the angular velocity outweigh the gyroscope information.

Motion model covariance It must be said, however, that although the angular velocity random walk model outperforms the other motion models, in reality, this does not have to be the case. It was assumed here that the process noise was equal to the noise of the estimates. It would be reasonable to assume that the motion models using the angular acceleration as well have lower process noise than the angular velocity random walk model. These parameters would have to be tuned in reality depending on the application.

Effects on orientation estimation Going back to Figure 3-2 shown at the end of Chapter 3, where a schematic was shown on how the angular velocity and specific force can provide an estimation of the orientation. The results shown above directly influence the orientation estimation. A more accurate estimation of the angular velocity provides less drift around the roll, pitch and yaw angle. The specific force can compensate for the drift around the pitch and the roll angle. A bigger configuration, such as the size of a human arm, already allows for more accurate estimations of the angular velocity. Not only do the accelerometers provide more information on the angular velocity, they also provide more information on the angular acceleration motion model. The angular acceleration as a motion model provides the most information as lower angular velocities. This can be observed at $\omega = \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T$ rad/s, where the angular acceleration motion model RMSE ratios are 10.58 and 4.763 for the MIMU4844 IMU configuration and the cylinder configuration respectively. In short for better orientation estimation at lower angular velocities. The angular acceleration can provide good information. This effect is increased if the distance between sensors is increased. For higher angular velocities the accelerometers provide an accurate estimate.

Chapter 5

Fault Detection and Isolation for Multiple Inertial Measurement Units

In the previous two chapters assumed that the only errors in sensor measurements were either the bias or the white noise. It is more realistic, however, to assume that other faults may also potentially happen. Such a fault can be described as a disturbance. Using 3 or more sensors, which can be categorized as component redundancy, allows for Fault Detection and Isolation (FDI). This chapter is focused on how a fault can be recognized and how it can be rejected. The assumption is made that all measurement noise is Gaussian, from this assumption a hypothesis test can be derived as shown in 5-1. Afterwards, in Section 5-2, the parity space method is introduced. Thereafter, in Section 5-3 a novel parity space method is introduced which builds on the work shown in Chapter 3. This method allows for the accelerometer positions and angular motion to be used as well. Finally, an orientation estimation scheme is shown in Section 5-4.

5-1 Hypothesis Testing

If one wants to know whether a change occurs in a signal or whether a fault is present it can be tested against a number of hypotheses. One of the more simplistic tests is the binary hypothesis test. The test takes on two hypotheses, the null hypothesis and the alternate hypothesis, which are expected to be true and false respectively. In case of fault detection the null hypothesis $\mathcal{H}_0$ would state that no fault is present and the alternate hypothesis $\mathcal{H}_1$ states that there is a fault present [45]. This can be expressed as

$$\mathcal{H}_0 : \text{Fault is not present,} \tag{5-1a}$$
$$\mathcal{H}_1 : \text{Fault is present.} \tag{5-1b}$$

In order to reject the $\mathcal{H}_0$ or not reject $\mathcal{H}_0$ the following definitions are needed

- The detection probability : $P_D = P(\mathcal{H}_1|\mathcal{H}_1)$

- The correct rejection probability : $P_C = P(\mathcal{H}_0|\mathcal{H}_0)$
- The false alarm probability, or type I error : $P_{FA} = P(\mathcal{H}_1|\mathcal{H}_0) = 1 - P_C$
- The miss probability, or type II error : $P_M = P(\mathcal{H}_0|\mathcal{H}_1) = 1 - P_D$

To decide between the two hypotheses a likelihood ratio $\mathcal{L}(y)$ can be established as

$$\mathcal{L}(y) = \frac{p(y|\mathcal{H}_1)}{p(y|\mathcal{H}_0)} > \gamma, \tag{5-2}$$

where γ is the threshold and is chosen arbitrarily depending on the application. This rather simple hypothesis test is according to the Neymans-Pearson's lemma optimal [12] as it maximizes the detection probability P_D for a given false alarm rate P_{FA} [20], where P_{FA} is a function of γ. More specifically, one can set the false alarm probability to a fixed significance level such that γ can be found [37]

$$P_{FA} = \int_{y:\mathcal{L}(y)>\gamma} p(y|\mathcal{H}_0)dy = \alpha. \tag{5-3}$$

Two types of binary hypothesis are discussed in theory, the first is the Simple Hypothesis test, where a state can be either of two a priori known values (This require the user to know these values, which is not true in our case). The second, known as composite hypothesis testing is where the state can take up an expected value or not that expected value. The second test, is more useful for realistic applications and is explained below.

5-1-1 Composite Hypothesis Testing

The composite hypothesis test considers two hypothesis. The null hypothesis $\mathcal{H}_0$, where the variable to be estimated takes up a known value and an alternative hypothesis $\mathcal{H}_1$, where the variable does not take up said value. This binary test is useful when one wants to test whether a fault is present or not. There are two common applications of this test. One is the Generalised Likelihood Ratio Test (GLRT) and the other is the Marginalised Likelihood Ratio Test (MLRT). Generally, one uses the MLRT when one wants to eliminate an unknown variable and GLRT when one can estimate the variable. Due to numerous sensors, a ML estimate can be found. The theory of composite hypothesis testing is shown below.

Consider the measurement given measurement model

$$y = Hx + e, \qquad e \sim \mathcal{N}(0, \Sigma), \tag{5-4}$$

where y is the measurements of a variable, H is the projection vector form the state space x onto measurements space y and e is Gaussian white noise. The likelihood function is then given by

$$p(y|x) = \frac{1}{(2\pi)^{\frac{N}{2}}\sqrt{\det(\Sigma)}} \exp\left(-\frac{1}{2}(y - Hx)^T\Sigma^{-1}(y - Hx)\right). \tag{5-5}$$

The parameter x can take up an unknown value. The binary hypothesis test now takes the shape of

$$\mathcal{H}_0 : x = x_0, \tag{5-6a}$$
$$\mathcal{H}_1 : x \neq x_0. \tag{5-6b}$$

There are two major approaches to constructing a detector for such a hypothesis test [37], the former being the GLRT and the latter being Marginalization.

Generalised Likelihood Ratio Testing The Generalised Likelihood Ratio Test takes the ML estimate of x under $\mathcal{H}_1$, which from linear least squares is given by

$$\hat{x} = (H^T H)^{-1} H^T y, \tag{5-7}$$

and plugs this in as the value for x. This gives the following hypothesis test

$$\mathcal{L}(y) = \frac{p(y|x = \hat{x})}{p(y|x = x_0)}. \tag{5-8}$$

The resulting distributions are

$$2 \ln \mathcal{L}(y) \sim \begin{cases} \chi^2_{n_x}, & \mathcal{H}_0, \\ \chi^2_{n_x}((x_1 - x_0)^T I_{x_0}(x_1 - x_0)), & \mathcal{H}_1, \end{cases} \tag{5-9}$$

where I_{x_0} is the Fisher Information matrix under x_0.

5-2 Parity Space Method for Multiple IMUs Assuming Linear Accelerometer Measurements

One of the most used FDI method on IMUs is the parity space. The parity space method exists on two accounts, one is analytical redundancy and the other is component redundancy. The analytical redundancy refers to expressing past states towards the current state so that they can be compared. Component redundancy refers to having a multiplicity of the same measurement. Other than this difference, the method is the same for both applications.

The parity space is based upon GLRT as it creates a projection matrix that is orthogonal to the ML estimate state vector $\hat{x}$. The resulting vector is now free of the ML estimate state vector and is called the residual vector r. If no fault is present in the measurements, the residual vector contains only information on erroneous measurements due to noise. If a fault is present then the parity vector contains information on noise and the fault. The parity space method is largely followed by [46] and is explained below.

Consider the measurement equation and assume that the fault vector $F = 0$.

$$y = Hx + F + e, \qquad e \sim \mathcal{N}(0, \Sigma), \tag{5-10}$$

The ML estimate of is

$$\hat{x} = (H^T H)^{-1} H^T y. \tag{5-11}$$

Using the ML estimate a projection matrix can be made as

$$\hat{y} = H\hat{x} = \underbrace{H(H^T H)^{-1} H^T}_{\mathcal{P}_{\mathcal{H}}} y, \tag{5-12}$$

where the projection matrix $\mathcal{P}_{\mathcal{H}}$ maps the measurement space onto the expected measurement space based on the ML estimate. Furthermore, $\mathcal{P}_{\mathcal{H}}$ must adhere to ${\mathcal{P}_{\mathcal{H}}}^2 = \mathcal{P}_{\mathcal{H}}$ for it to be a projection matrix [45]. The prediction error can now be found as

$$\epsilon = y - \hat{y}, \tag{5-13}$$

which leads to the orthogonal projection matrix

$$\epsilon = \underbrace{(I - \mathcal{P}_H)}_{\mathcal{P}_{H_\perp}} y, \tag{5-14}$$

where $\mathcal{B}_{H_\perp}$ is of $\text{rank}(\mathcal{P}_{H_\perp}) = n - \text{rank}(H)$. It was shown, however, by [47] that due to the rank deficiency of the projection matrix $\mathcal{P}_{H_\perp}$, the prediction error vector is of lower rank and thus its covariance matrix can become non-singular. This can be circumvented by pre-multiplying the prediction error vector by the orthonormal basis $\mathcal{B}_{H_\perp}$. As was shown in [**?**].

Furthermore, the projection matrix can also be constructed via the orthonormal basis $\mathcal{B}_{H_\perp}^T$ of the range space of H, denoted by $\mathcal{R}(H_\perp)$ [45]. The projection matrix can be constructed as

$$\mathcal{P}_{H_\perp} = \mathcal{B}_{H_\perp} \mathcal{B}_{H_\perp}^T. \tag{5-15}$$

It can be written in this form and to the fact that $\mathcal{B}_{H_\perp}^T \mathcal{B}_{H_\perp} = \mathcal{I}$, prediction error can be rewritten into the residual vector r [45]

$$r = \mathcal{B}_{H_\perp}^T r = \mathcal{B}_{H_\perp}^T \mathcal{B}_{H_\perp} \mathcal{B}_{H_\perp}^T y = \mathcal{B}_{H_\perp}^T y, \tag{5-16}$$

of which the co-variance is

$$\text{cov}(r) = \mathcal{B}_{H_\perp}^T \Sigma \mathcal{B}_{H_\perp}. \tag{5-17}$$

By multiplying the residual by the square root of its covariance, the normalized residual has covariance identity. This is also known as whitening [**?**]. The normalized residual then takes the shape of

$$z = \underbrace{(\mathcal{B}_{H_\perp}^T \Sigma \mathcal{B}_{H_\perp})^{-\frac{1}{2}} \mathcal{B}_{H_\perp}^T}_{W} y. \tag{5-18}$$

From this another hypothesis can be made based on whether the fault in (5-10) is zero or not, such that the null hypothesis states there is no fault and the alternative hypothesis states there is a fault.

$$\mathcal{H}_0 : z \sim \mathcal{N}(0, \mathcal{I}), \tag{5-19a}$$

$$\mathcal{H}_1 : z \sim \mathcal{N}(WF, \mathcal{I}). \tag{5-19b}$$

The Chi-square test statistic was derived similar to [45] as

$$\mathcal{L} \sim \begin{cases} \chi^2_\nu, & \mathcal{H}_0, \\ \chi^2_\nu(F^T\mathcal{B}_{H_\perp}(\mathcal{B}^T_{H_\perp}\Sigma^{-1}\mathcal{B}_{H_\perp})^{-1}\mathcal{B}^T_{H_\perp}F)., & \mathcal{H}_1, \end{cases} \tag{5-20}$$

where $\nu = \text{rank}(W) = \text{rank}(\mathcal{P}_{H_\perp})$. From this distribution a threshold can be chosen such that the null hypothesis can be rejected with a certain probability. In the chi-square test, when no fault is present the distribution takes the shape of the central chi-square distribution. If a fault is present, the distribution takes the shape of the non-central chi-square distribution. Both distributions can be seen in Section A-3 in Appendix C.

5-3 Parity Space Method for Multiple IMUs Assuming Non-linear Accelerometer Measurements

The parity space method has been used in IMU applications, see for example [9, 50, 49, 45], however, to the knowledge of the writer, this is the first time it is used as an extension on the formulation in Section 3-4 by [40]. This extension is motivated to remove the effects due to rotational motion and larger spatial separation between accelerometers. Previous accounts have assumed that the accelerometer only measure the specific force. This can be a good approximation for low-dynamic applications and where accelerometers are placed closely together. However, if the sensors are placed further apart, the rotational motion starts having a bigger effect on the accelerometer measurements. Take for example the angular velocity $\omega = \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}^T$ rad/s and angular acceleration $\dot{\omega} \begin{bmatrix} 40 & 40 & 40 \end{bmatrix}^T$ rad/s^2 (these values are taken from Figure B-5 Figure B-6 at their peaks). Two different configurations, the MIMU4844 array and the cylinder have their accelerometers at a maximum distance between sensors of $\begin{bmatrix} 0.0185 & 0.0185 & 0.03 \end{bmatrix}^T$ m and $\begin{bmatrix} 0.30 & 0.1 & 0.1 \end{bmatrix}^T$ m respectively. This would result in an unmodeled error between the two accelerometers of $\begin{bmatrix} -0.64 & 0.64 & 0.03 \end{bmatrix}^T$ m/s^2 and $\begin{bmatrix} -0.4 & 8.2 & -7.8 \end{bmatrix}^T$ m/s^2 respectively. This would mean that sensors are detected to be false due to unmodeled errors in bigger configurations (but also at higher angular motion). This is the motivation to separate the angular motion in larger configurations.

The application of FDI on MIMUs will be applied in two steps. In the first step, the linear parity space method as discussed in Section 5-2 is applied to the gyroscopes measurements only. It is of importance to eliminate faulty measurements first as these faulty measurements will influence the initial estimate of the angular velocity. The angular velocity estimate will then removed as this is the non-linear part from the stacked measurement equation (3-4c). The closer the angular velocity estimate is to the true angular velocity, the fewer errors are induced into the measurement equation. The stacked measurement equation is restated here

as

$$\underbrace{\begin{bmatrix} y_{a_1} \\ \vdots \\ y_{a_N} \end{bmatrix}}_{Y_a} = \underbrace{\begin{bmatrix} [\omega\times]^2\rho_1 \\ \vdots \\ [\omega\times]^2\rho_N \end{bmatrix}}_{h_a(\omega)} + \underbrace{\begin{bmatrix} -[\rho_1\times] & \mathcal{I}_3 \\ \vdots & \vdots \\ -[\rho_N\times] & \mathcal{I}_3 \end{bmatrix}}_{H_a} \underbrace{\begin{bmatrix} \dot{\omega} \\ a \end{bmatrix}}_{\theta} + \underbrace{\begin{bmatrix} e_{a_1} \\ \vdots \\ e_{a_N} \end{bmatrix}}_{E_a},$$

$$\underbrace{\begin{bmatrix} y_{\omega_1}^T & \dots & y_{\omega_N}^T \end{bmatrix}^T}_{Y_\omega} = \underbrace{\begin{bmatrix} \omega^T & \dots & \omega^T \end{bmatrix}^T}_{h_\omega(\omega)} + \underbrace{\begin{bmatrix} e_{a_1}^T & \dots & e_{a_N}^T \end{bmatrix}^T}_{E_\omega},$$

$$\underbrace{\begin{bmatrix} Y_a \\ Y_\omega \end{bmatrix}}_{Y} = \underbrace{\begin{bmatrix} h_a(\omega) \\ h_\omega(\omega) \end{bmatrix}}_{h(\omega)} + \underbrace{\begin{bmatrix} H_a \\ H_\omega \end{bmatrix}}_{H} \theta + \underbrace{\begin{bmatrix} E_a \\ E_\omega \end{bmatrix}}_{E}.$$

The stacked observations (3-4c) are non-linear. The non-linear part, however, can be removed similar to the method described Section 3-4, however, the assumption in this case is made that some sensors contain a fault. Under this assumption, the fusion of all measurements would lead to a biased ML estimate. Instead of using the Gauss-Newton optimization to find the angular velocity, use of a prior from the motion model as described in Section 3-5 can be made.

Simultaneously, the parity space method can also be applied to gyroscope measurements alone. It is assumed that the prior $\hat{\tilde{\omega}}$ and the healthy gyroscope measurements are now only corrupted by white noise. It is also assumed that all the estimates of the angular velocity are independent of each other. The estimates of the angular velocity estimates are distributed as

$$\begin{aligned} \hat{\tilde{\omega}} &\sim \mathcal{N}(\tilde{\omega}, \Sigma_{\tilde{\omega}}), \\ y_{\omega,i} &\sim \mathcal{N}(\omega, Q_{\omega,i}). \end{aligned}$$

Where $i = 1 : N_{\omega,\text{healthy}}$ denotes the healthy gyroscopes. A fused estimate of the angular velocity at a time instance can be found by taking the WLS as [20]

$$\hat{\omega} = \left(\Sigma_{\tilde{\omega}}^{-1} + \sum_{i=1}^{N_{\omega,\text{healthy}}} Q_{\omega,i}^{-1} \right)^{-1} \left(\Sigma_{\tilde{\omega}}^{-1} \hat{\tilde{\omega}} + \sum_{i=1}^{N_{\omega,\text{healthy}}} Q_{\omega,i}^{-1} y_{\omega,i} \right). \tag{5-23}$$

After the faulty gyroscopes have been identified, the weighted average of the gyroscopes can be taken for an estimate of the angular velocity. This leads to two independent estimates of the angular velocity, which can be fused to remove the non-linear part $h(\omega)$ from the measurement space.

In is assumed that the estimate of the angular velocity close enough to the true angular velocity. The adapted measurement equation has covariance $R = H(H^T Q^{-1} H)^{-1} H^T + Q$, as was shown in Section 3-4 and takes the form of

$$Y - h(\hat{\omega}_0) = H\hat{\theta} + E. \tag{5-24}$$

The parity space method as discussed in Section 5-2 can now be applied to the previous adapted measurement equation. The method is summarized in Algorithm 3.

Algorithm 3: A novel parity space method for multiples IMUs to identify faulty measurements

Run FDI scheme for gyroscope measurements, set $T_\omega = 1$;
while $T_\omega = 1$ **do**

$$\begin{aligned}
\mathcal{B}_{H_{\omega,\perp}} &= \text{null}(H_\omega^T), \\
\nu_\omega &= \text{rank}\left(\mathcal{B}_{H_{\omega,\perp}}\right), \\
r_\omega &= \left(\mathcal{B}_{H_{\omega,\perp}}^T \mathcal{B}_{H_{\omega,\perp}}\right) Y_\omega, \\
z_\omega^T z_\omega &= Y_\omega^T \mathcal{B}_{H_{\omega,\perp}} \left(\mathcal{B}_{H_{\omega,\perp}}^T Q_\omega \mathcal{B}_{H_{\omega,\perp}}\right)^{-1} \mathcal{B}_{H_{\omega,\perp}}^T Y_\omega.
\end{aligned}$$

Run hypothesis test on $z_\omega^T z_\omega \sim \chi^2_{\nu_\omega}$ to update T_ω ;
if $T_\omega = 1$ **then**
find sensor with highest residual in vector r_ω ;
remove said sensor from measurements ;
end
end
Run FDI scheme for accelerometer + gyroscope measurements ;

$$\hat{\omega} = \left(\Sigma_{\tilde{\omega}}^{-1} + \sum_{i=1}^{N_{\omega,\text{healthy}}} Q_{\omega,i}^{-1}\right)^{-1} \left(\Sigma_{\tilde{\omega}}^{-1}\hat{\omega} + \sum_{i=1}^{N_{\omega,\text{healthy}}} Q_{\omega,i}^{-1} y_{\omega,i}\right).$$

while $T = 1$ **do**

$$\begin{aligned}
\mathcal{B}_{H_\perp} &= \text{null}(H^T), \\
\nu &= \text{rank}\left(\mathcal{B}_{H_\perp}\right), \\
r &= \left(\mathcal{B}_{H_\perp}^T \mathcal{B}_{H_\perp}\right)\left(Y - h(\hat{\omega})\right), \\
z^T z &= \left(Y - h(\hat{\omega})\right)^T \mathcal{B}_{H_\perp} \left(\mathcal{B}_{H_\perp}^T R \mathcal{B}_{H_\perp}\right)^{-1} \mathcal{B}_{H_\perp}^T \left(Y - h(\hat{\omega})\right).
\end{aligned}$$

Run hypothesis test on $z^T z \sim \chi^2_\nu$ to update T ;
if $T = 1$ **then**
find sensor with highest residual in vector r ;
remove said sensor from measurements ;
end
end

5-4 Fault Detection and Isolation with Multiple IMUs Orientation Estimation

Two orientation estimation schemes have been shown earlier in Section 2-5 and Section 3-7 where an orientation estimation scheme was shown for a single IMU and multiple IMUs respectively. In this chapter, a novel method for FDI was introduced which allows for the detection and isolation of faulty sensors. This takes place after the sensors are pre-processed and before they are fused into a synthetic IMU. Use of the angular acceleration motion model prior not only improves the estimate of the angular velocity, but also allows for this novel parity space method to be extended towards accelerometer-only arrays. The new orientation estimation scheme is shown in Figure 5-1 where the green rectangle shows how FDI can be implemented into the previous orientation estimation.

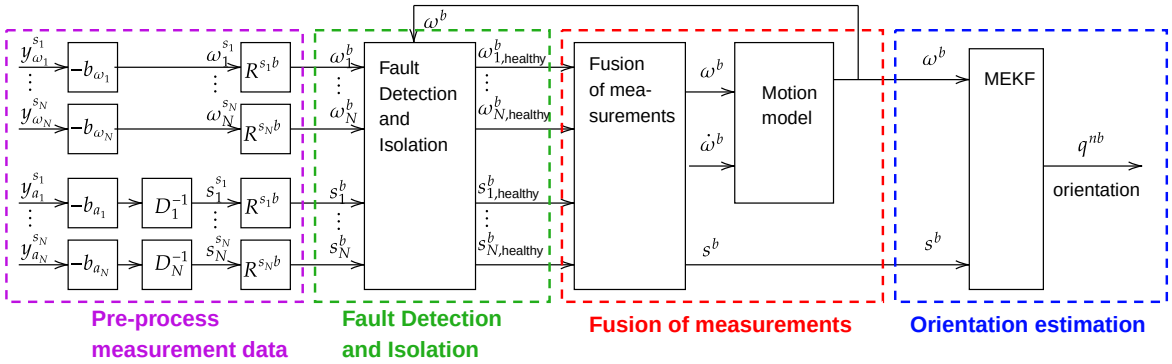

Figure 5-1: An orientation estimation scheme using multiple gyroscopes and accelerometers and FDI. The measurement data is first pre-processed in the purple rectangle after which in the green rectangle faulty sensors are removed. The remaining healthy sensors are fused into a synthetic IMU in the red rectangle. The synthetic IMU then estimates the orientation in the blue rectangle.

In this chapter, the parity space method for multiple IMUs has been discussed. Previous accounts assumed a linear accelerometer measurement model. A novel parity space method to separate the rotational motion from the accelerometer measurements has been presented. The new parity space method is included in the orientation estimation scheme from Chapter 3 and is shown in Figure 5-1. This orientation estimation scheme will be used in Chapter 6 to validate the new parity space method.

Chapter 6

Fault Detection and Isolation Simulation Setup and Results

This chapter will validate the novel parity space method as presented in 5-3. In Section 6-1 the simulation setup for the cylinder, how the simulation data is generated and how the disturbance is simulated is discussed. Afterwards, in 6-2 the results of the simulation are shown. In the results, first, the states are shown and to what extent the disturbance can be rejected, afterwards how the covariance of the specific force is affected is shown. Finally, it is shown how the orientation estimation is affected by the disturbance.

6-1 Simulations to test FDI

In Chapter 2 it was shown how an orientation estimate can be found using angular velocity and specific force estimates. Afterwards in Chapter 3 it was shown how multiple IMUs can be fused into a single synthetic IMU. In Chapter 5 a novel parity space method to detect and isolate faulty sensors was proposed. How everything combines is summarized in Figure 5-1. In this simulation, the novel parity space method will be tested with a MC simulation using 100 data realizations. Additionally, a comparison using the newly developed parity space method on the cylinder configuration, the MIMU4844 configuration and the single Xsens MTi 300 IMU (without FDI as only one sensor is used).

A short summary on orientation estimation To analyze the parity space method the orientation estimation scheme in Section 5-4 is used. The first step here is to pre-process the data, however, as it is assumed that the measurements do not have a bias and are not prone to misalignment errors, this step can be disregarded. The measurement data at this point is in the individual sensor frames and first must be rotated into the body frame. Afterwards, the data in the body frame can be tested against each other using the novel parity space method described in Section 5-3. The remaining healthy sensors can then be fused into a Synthetic IMU and use of the angular acceleration can be made to obtain a new angular velocity prior. Finally, the synthetic IMU can be used to obtain the orientation estimation in the MEKF.

Simulation setup for a cylinder Previously, the capabilities of placing IMUs in an array and placing IMUs on a cylinder were analyzed. Here, the simulation continues on the cylinder in the shape and size of a human arm segment. The cylinder is $0.3m$ long and has a radius of $0.05m$, the IMUs are placed equidistantly and can be seen in Figure 6-1. It is assumed that biases and misalignment errors can be calibrated for properly. It is also assumed that the positions can be calibrated for properly. The covariances of the accelerometers and gyroscopes are normally distributed with covariances matrices

$$\Sigma_{a_i} = \begin{bmatrix} 3.10 & 0 & 0 \\ 0 & 3.10 & 0 \\ 0 & 0 & 3.41 \end{bmatrix} 10^{-2}, \qquad \Sigma_{\omega_j} = \begin{bmatrix} 8.40 & 0 & 0 \\ 0 & 8.70 & 0 \\ 0 & 0 & 6.20 \end{bmatrix} 10^{-3}, \tag{6-1}$$

for $i = 1:32$ accelerometers and $j = 1:32$ gyroscopes.

Simulation setup for the MIMU4844 array The assumptions made for the IMU are the same as for the cylinder. The difference here is the placement of the IMUs. Instead of spreading the IMUs across the arm, they are placed in the middle of x- and y-axis and on top of the cylinder in the z-axis as can be seen in Figure 6-1. They maintain their configuration of the MIMU4844 as their positions were calibrated for.

Simulation setup for the Xsens MTi 300 The other two configurations used multiple IMUs where this sensor uses a single higher-grade IMU. This results of this sensors are used as a comparison and no FDI is applied. The sensor is placed at the origin of the MIMU4844 array in Figure 6-1. The covariance matrices for the accelerometer and gyroscope are

$$\Sigma_{a_{\mathrm{X}}} = \begin{bmatrix} 3.335 & 0 & 0 \\ 0 & 3.323 & 0 \\ 0 & 0 & 2.470 \end{bmatrix} 10^{-5}, \qquad \Sigma_{\omega_{\mathrm{X}}} = \begin{bmatrix} 2.096 & 0 & 0 \\ 0 & 2.076 & 0 \\ 0 & 0 & 1.474 \end{bmatrix} 10^{-6}. \tag{6-2}$$

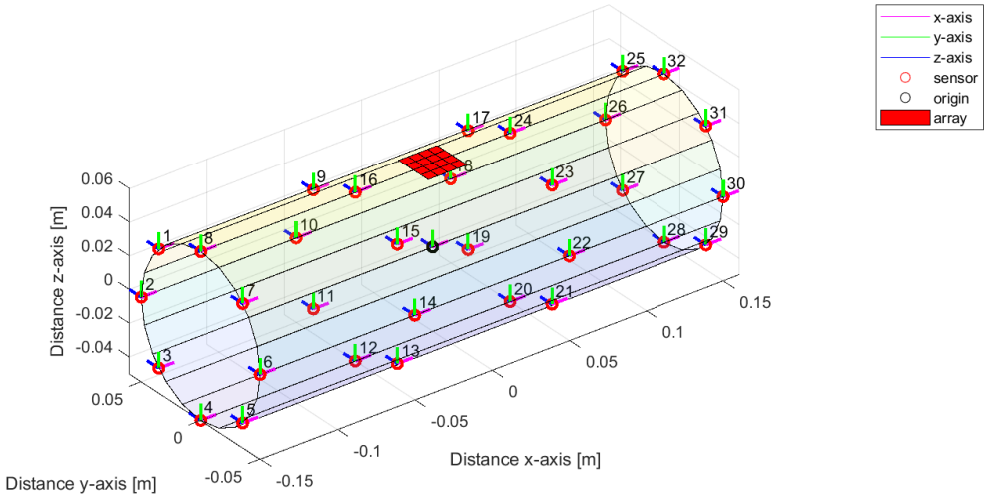

Figure 6-1: Locations and orientation of 32 accelerometers and gyroscopes placed on the surface of the cylinder and have their sensitive axes point in the same directions. The location of the MIMU4844 array is also shown. At the centre of the array, the higher grade IMU is placed.

Simulation data The simulation of measurement data was done as described in Section 4-3 and summarized in Algorithm 4 in Appendix C. The linear acceleration was set at $\begin{bmatrix} 0 & 0 & 0 \end{bmatrix}^T$ m/s. The magnitude of the linear acceleration should have no effect on the novel parity space method and could degrade the orientation estimate. For this reason it is set at zero. Instead of setting the angular velocity to arbitrary values, a random walk model for an angle ϕ^b was generated. Afterwards, the angle was smoothed using a moving average filter in to represent an arm smoothly rotating through space. The angular velocity can be calculated from this angle as the first order discrete time derivative. The simulation of this data is as

$$\phi^b_{k+1} = \phi^b_k + e_{\phi,k}, \tag{6-3a}$$

$$\omega^b_{k+1} = \frac{\phi^b_{k+1} - \phi^b_k}{T}. \tag{6-3b}$$

Simulating a disturbance One of the problems in motion capture for humans is disturbances caused by hitting or banging against an object. As the surface of the arm is soft tissue, the disturbance should be the largest at the point of contact of where the arm would hit an object. The magnitude of the disturbance should dissipate along the length of the arm. The disturbance is simulated for 1 second as a step input and the magnitude varies per set of accelerometers. The magnitude of the disturbance are shown in Table. The disturbance will be applied in the x-, y- and z-axis of the body frame. The magnitude (for the x-, -y and z-axis) of the disturbance per accelerometer are shown in Figure 6-2.

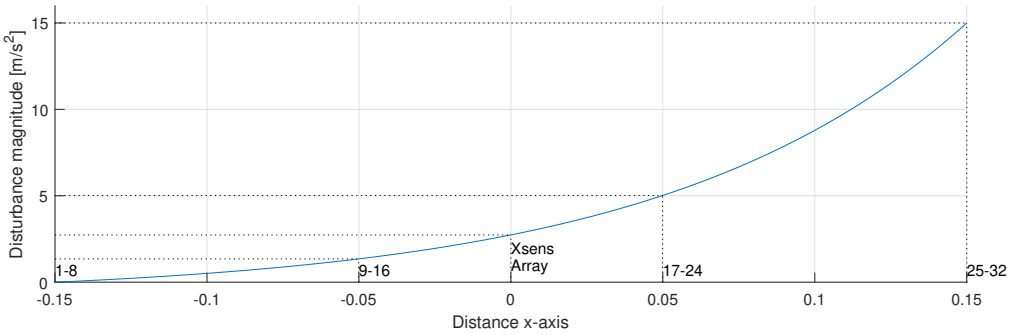

Figure 6-2: Disturbance applied on the x-, y- and z-axis of the accelerometers in the body frame. The disturbance diminishes over the length of the x-axis of the cylinder. The numbers, array and Xsens correspond to the accelerometers as in Figure 6-1.

Accelerometer disturbance on orientation estimation The disturbance will cause a spike in the accelerometer readings resulting in erroneous measurements. When using a single IMU, the fault will influence the approximation of the specific force. The specific force is used in the MEKF and now updates the orientation with a fault. Using multiple IMUs, a difference between two accelerometer measurements will result in rotational motion. However, this difference is now due to the fault. As a results, there will be faults in the angular velocity and the angular acceleration. Is the gyroscopes measure the angular velocity, it is expected that the angular acceleration will absorb most of the fault. The MEKF will have a process and a measurement update containing faults. This will also result in a faulty estimate of the orientation. In this simulation, we will try to reject the disturbance using the novel parity space method. Note that it in this simulation it is assumed that the gyroscopes are unaffected by the disturbance which is a limitation.

6-2 Results on Testing the Novel Parity Space Method

The results are shown in the subsections below, first the states (s, ω, $\dot{\omega}$) are shown and to what extent the disturbance applied influences the estimate of the states. Afterwards, the covariance of the specific force is shown and how it is impacted by the location of where the synthetic IMU is fused. Thereafter, the orientation estimation is shown to analyze the impact on the orientation.

6-2-1 Detection Probabilities of the Accelerometers

Accelerometers 1-8 experience no disturbance and accelerometers 25-32 experience the most disturbance. From the applied disturbance we know that accelerometers 1-8 never contain a fault whereas accelerometer 9-32 always contain a fault. We would expect that the accelerometers containing a fault are rejected and the sensors not containing a fault are deemed healthy. This is not entirely what the results show as can be seen in Figure 6-8. In this figure, the detection probabilities of the accelerometers are shown. As can be observed, accelerometers 17-32 are detected correctly 100% of the time as expected. Accelerometers 9-16, however,

have a detection probability of about 70% (with accelerometers 9 and 10 above 80%). Simultaneously, accelerometers 1-8 have a false alarm probability of about 65%. The parity space method looks at the ML estimate of the measurements and looks at which measurements is the farthest off from the ML estimate. After sensors 19-32 have been isolated, a ML estimate is found on the accelerometer measurements 1-16. This ML estimate contains the a fault from accelerometers measurements 9-16. As the parity space method does not know which sensors are faulty and not it labels sensors 1-16 as faulty with equal probability. The remaining sensors afterwards fused into a synthetic IMU which may contain part of the disturbance. This will have an effect on the states (s, ω and $\dot{\omega}$), which will be looked at in the next subsection.

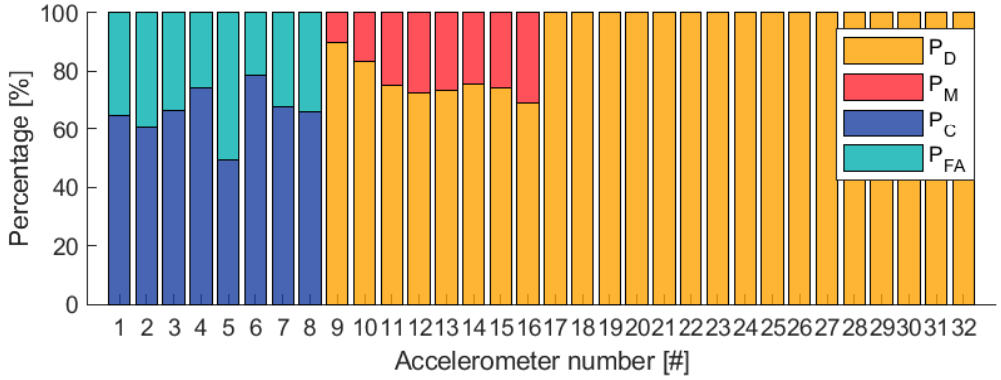

Figure 6-3: Detection probabilities with 1000 MC realization for all 32 accelerometers. Accelerometers 1-8 contain no fault and accelerometers 9-32 contain a fault. This fault increases in magnitude per set of 8 accelerometers.

6-2-2 Effect of the disturbance on the angular velocity, angular acceleration and specific force

As was previously observed, not all of the disturbance was rejected. Some of the measurements that are fused into a synthetic IMU contain a fault. This will result in the fault appearing in the angular velocity, angular acceleration and specific force. There are two types of motion that the cylinder (a rigid body) will undergo: linear and rotational. If an accelerometer at different location experiences the same measurement, it will be interpreted as linear motion. If two sensors experience a difference in the measurement, this will be interpreted as rotational motion. At the time, the gyroscopes and the rotational acceleration (in the form of a motion model prior) provide information on the angular velocity. As the gyroscope measurements are not influenced by the disturbance, they give a good estimate of the true angular velocity. For this reason, we would expect that the angular velocity estimate is only slightly disturbed whereas the rotational acceleration experiences a larger disturbance.

Angular velocity The x- (left), y- (middle) and z-axis (right) mean angular velocity errors are plotted in Figure 6-4 (the actual mean angular velocities are shown in Figure B-5 in Appendix B). In this figure, the novel parity space method (pink) is compared to no FDI method (blue). A disturbance is applied at from 1 to 2 seconds. We can see that the novel parity space method exhibits only minor errors. When no FDI is applied, we can observe

minor errors in the x-axis and bigger errors in the y- and z-axis. This is to be expected as the accelerometers are more spread out in the x-axis (0.3m total) of the cylinder than in the x- and y-axes (both 0.1m total). At the time the disturbance magnitude increases (for all 3 axes) in the x-direction. The accelerometers measure the same difference (between different accelerometers) over a smaller distance (with respect to the x-axis). This difference between measurements results in the fusion algorithm thinking this is due to rotational motion. The gyroscopes also provide information on the angular velocity, which only contains noise. This has the most impact on rejecting the disturbance. However, as can be seen, some error remains. From the results, it can be concluded that the disturbance has no significant effect on the angular velocity estimate when using the novel parity space method.

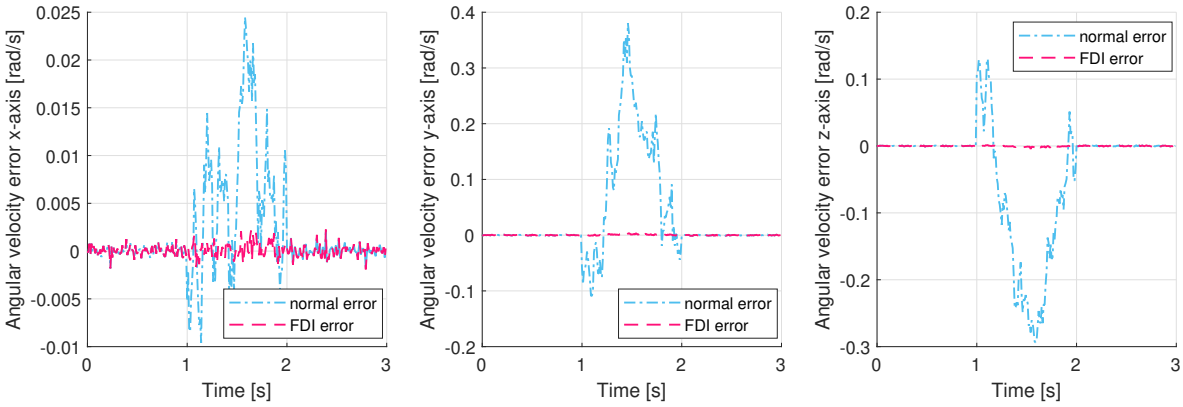

Figure 6-4: Mean angular velocity errors of the synthetic IMU with (blue) and without (pink) FDI using 1000 MC realizations for the x-axis (left), y-axis (mid) and z-axis (right) of the body frame.

Angular acceleration In Figure B-6 and Figure 6-5 the mean angular acceleration errors are plotted (the actual mean angular accelerations are shown in Figure B-6 in Appendix B). The angular acceleration error using no FDI (blue) is compared to the novel parity space method (pink). As expected, when not using FDI, the angular acceleration is disturbed significantly (almost 45 rad/s^2 for the y- and z-axes). Interestingly, the no FDI method does not exhibit a disturbance in the x-axis. This can be explained by that the magnitude of the disturbance increases along the x-axis of the cylinder. For this reason, the accelerometers measurements are the same around the x-axis. The novel parity space method can effectively reject most of the distance on the y- and z-axes. This, however, does come at a small cost of inducing errors on the x-axis. Overall, the error has been reduced more than 4.5 times for the y- and z-axes. Some error remains, which is to be expected not all accelerometers are rejected. Overall the novel parity space method rejects the majority of the influence of the disturbance on the angular acceleration.

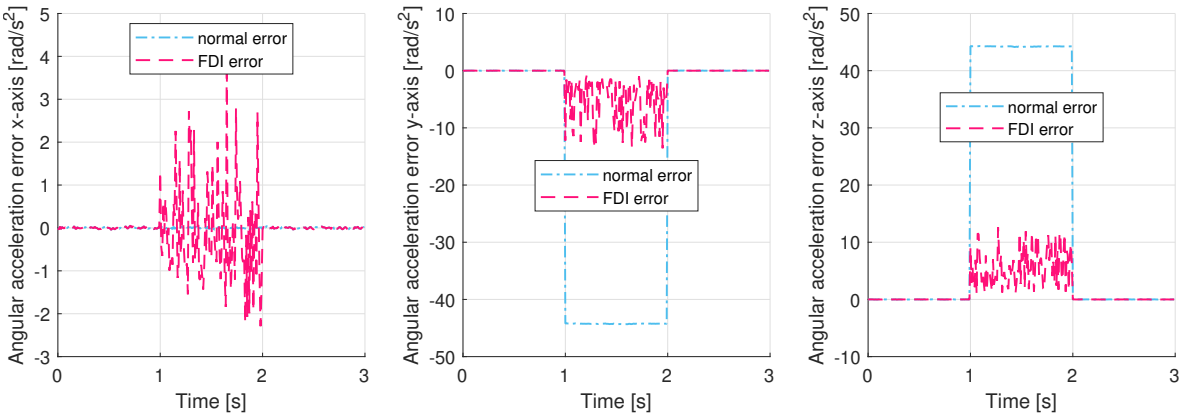

Figure 6-5: Mean angular acceleration errors of the synthetic IMU with (blue) and without (pink) FDI using 1000 MC realizations for the x-axis (left), y-axis (mid) and z-axis (right) of the body frame.

Specific force The mean specific force errors errors are shown in Figure 6-6 respectively (the actual mean angular accelerations are shown in Figure B-7 in Appendix B. The x-(left), y-(middle) and z-(axis) mean errors are plotted for using the novel parity space method (pink) and not using FDI (blue). When not using FDI, every axis undergoes a disturbance of almost 5.5 m/s^2. This is to be expected as the disturbance applied grows exponentially (see Figure 6-2 for a reference). This means that part of the disturbance causes an error in the angular acceleration and part causes an error in the specific force. The novel parity space method is able to reject most of the influence of the disturbance on the specific force. Interestingly, less is rejected along the y- and z-axis compared to the x-axis. In the end, there is still a disturbance around 1 m/s^2 in the y- and z-axis. This will still have an influence on the orientation estimation.

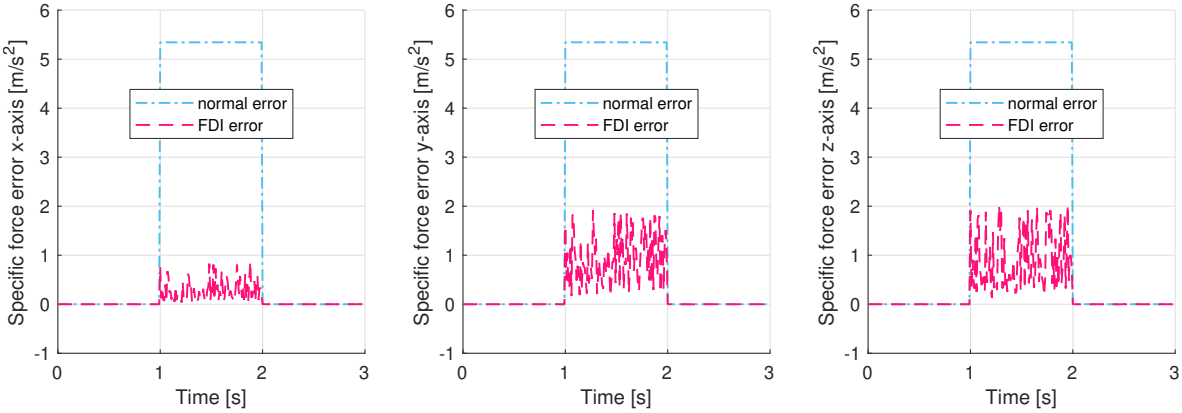

Figure 6-6: Mean specific force errors of the synthetic IMU with (blue) and without (pink) FDI using 1000 MC realizations for the x-axis (left), y-axis (mid) and z-axis (right) of the body frame.

Specific force covariance During the disturbance, most of the accelerometers are rejected, therefore, they are not taken into account when fusing the measurements. Due to the nature of the disturbance, however, the remaining accelerometers are all on one side of the cylinder.

This causes the origin to be displaced away from the centre of the accelerometers. This effect was shown earlier in Section 4-4-3. There, an analysis was done the CRLB and the displacement of the origin. When the origin displaces itself from the centre of the accelerometers, the covariance of the specific force increases. Both covariances are plotted in Figure 6-7. The uncertainty of the specific force has now increased such that the measurement update in the MEKF has negligible effect. This can be resolved by changing the origin from the centre of the cylinder towards the centre of the healthy accelerometers. The assumption is then made that a rigid body is used and thus the linear and rotational motion should be the same for every point in the rigid body. Thus, if the origin of the synthetic IMU is placed on any part of the rigid body the covariance estimate should be valid.

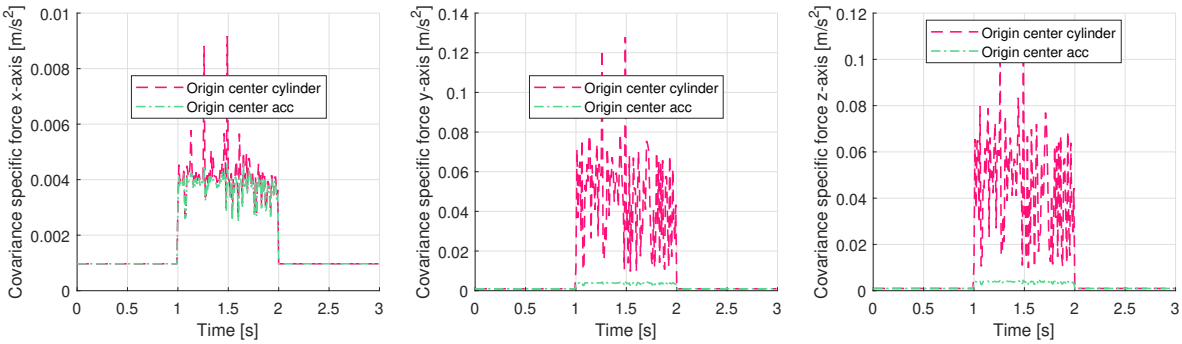

Figure 6-7: MC simulation of 1000 realizations with an applied disturbance. The x-axis (left), y-axis(mid) and z-axis (right) covariances of the specific force of the synthetic IMU is shown when the origin is in the centre of the cylinder (pink) and when the origin is in the centre of the healthy accelerometers (yellow).

6-2-3 Orientation Estimation

It is important to highlight two important update steps of the MEKF. These will help explain the results that follow. The two updates in the MEKf are the time update and the measurement update. In the time update, the angular velocity estimate is used to update our belief of the orientation. In the measurement update, the gravity vector (for orientation estimation it is assumed this is equal to the specific force) corrects for drift around the pitch and the roll angle. The MEKF is discussed in more detail in Section 2-3. The simulation results of the true data, no FDI method and the novel parity space method are plotted in Figure B-8. The respective errors are plotted in Figure 6-8. For all the synthetic IMUs, the roll and pitch angle converge back towards the true angles after the disturbance. This effect is due to the accelerometer compensates for the drift. The gravity vector, however, does not provide any information on the yaw angle. Therefore, in the aftermath of the disturbance, the yaw angles may experience a constant offset from the true angle.

Synthetic IMU without FDI In Figure 6-8 the orientation error estimate (actual estimates are shown in Figure B-8 in Appendix B) using no FDI are plotted in blue. It can be observed that at the time of the disturbance, when not using the parity space method (blue), the roll angle undergoes a large error which peaks at about $-180°$. The pitch angle error peaks at about $130°$ while the yaw angle error peaks at about $-80°$. These errors are to be expected as

all accelerometers are used. The accelerometers contain a disturbance that induces an error on the states (s ω and $\dot{\omega}$). Most notably, the angular velocity (y- and z- axes), the angular acceleration (y- and z-axes) and the specific force (all axes) have a large error. The error in the angular motion causes faulty time update in the MEKF. The specific force then causes a measurement update with wrong measurements. The main source of error in the orientation is due to the specific force containing a large fault. This effect can be observed from the roll and pitch angle. The error on the yaw angle, however, is caused by the faulty angular motion. Using no FDI for such a disturbance results in a wrong orientation estimation.

Synthetic IMU with FDI (pink) The novel parity space method where the origin of the synthetic IMU remains in the centre of the cylinder is plotted in pink. The error of the orientation can be seen in. At the time of the disturbance (from 1 to 2 seconds), the y- and z-axes of the specific force start to increase. This causes the measurement update in the MEKF to provide less information on our belief of the orientation, resulting in negligible effect. The error in the orientation is caused by the faults in the angular velocity estimate and no correction by the gravity vector. If we want to have a measurement update with more effect, the origin of the synthetic IMU must be shifted towards the centre of the healthy accelerometers. The result of this is shown in yellow.

Synthetic IMU with FDI (yellow) Finally, the novel parity space with the synthetic IMU origin at the centre of the healthy accelerometers (yellow) is shown. The orientation error estimate is plotted in Figure B-8. As can be observed, the angel has little drift away from the true angle. This improvement is due to the decrease in the covariance of the specific force. This results in more compensation using the gravity vector. From this we can conclude that the novel parity space effectively rejects the disturbance.

Different configurations The same simulation was run using three different configurations. The MIMU4844 array (purple) and the cylinder configuration (yellow) using multiple lower-grade IMUs and the third using a single higher-grade IMU: the MTi-300 from Xsens (orange). For the two lower-grade configurations, the novel parity space method was used. The higher-grade sensor had no FDI applied as there is only one sensor. The orientation error estimate results are compared against each other in Figure 6-9. As can be seen, the cylinder configuration can reject most of the disturbance, where the array configuration and the Xsens sensor both cannot reject the disturbance. This is to be expected as the Xsens sensor cannot make use of component redundancy. The array configuration has its sensors closely placed together (see Figure 6-1). The disturbance has almost the same magnitude for these sensors.From these results we can conclude that using a larger configuration can help in rejecting disturbances.

6-3 Concluding Remarks

In Chapter 5 the parity space method was introduced. The parity space method can be used when a multiplicity of sensors is available. The method checks at each time step whether sensors (against each other) are faulty or not. In previous accounts, the parity space method applied assumed that the accelerometer measurement model was linear. This assumption may

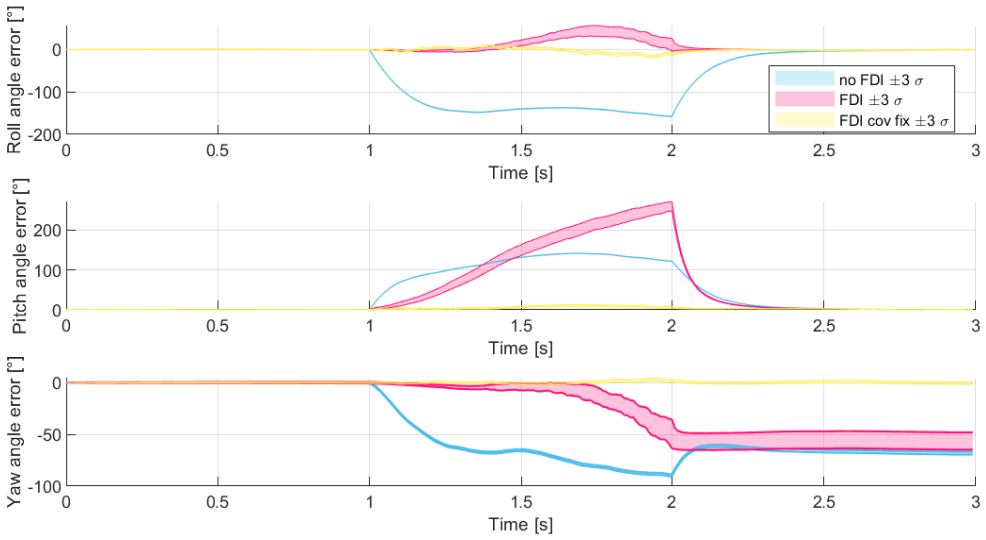

Figure 6-8: MC simulation of 1000 realizations with an applied disturbance. The synthetic IMUs undergo no FDI (blue), and undergo FDI (pink and yellow). Pink has the origin at the centre of the cylinder and yellow at the centre of the accelerometers. The estimated errors are shown of the roll (top), pitch (mid) and yaw (bottom).

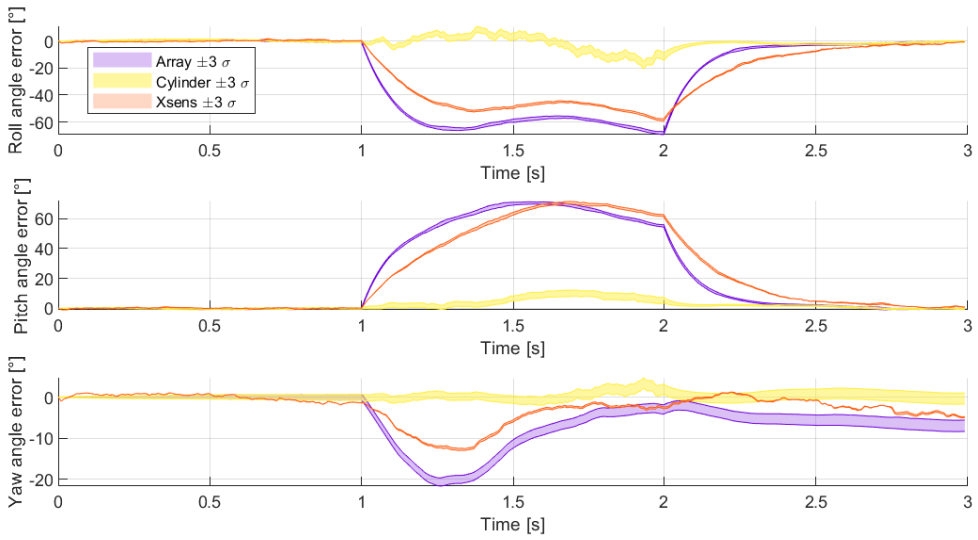

Figure 6-9: Roll (top), pitch (middle) and yaw (bottom) error estimates with an step disturbance from 1 to 2 seconds. Three configurations are compared, the cylinder (yellow), the MIMU4844 (purple) and a single higher-grade sensor (orange). Data is from a MC simulation of 1000 data realizations.

hold for closely placed accelerometers and low rotational motion. At higher angular motion and more distantly placed accelerometers, this assumption will be violated. For this reason, we proposed a novel parity space method that uses the ML estimator from [40, 51]. The novel

parity space method takes the angular velocity and angular acceleration into account. We tested the novel parity space method in a simulation. In this simulation, the accelerometers were scattered across the length of a cylinder. This configuration was compared against an array and a single higher-grade sensor (not using FDI).

In the results, we showed the effectiveness of the novel parity space method when the accelerometers are scattered all over a cylinder the size of a human arm. We also see that the array and the single Xsens sensor are not able to reject the disturbance. There is almost no advantage of using the parity space method compared to the Xsens sensor. From this, we can draw an important conclusion. The novel parity space method is successful because it allows the user to make full advantage of the size of an object. The novel parity space method can be used for even bigger configurations.

Chapter 7

Discussion and Conclusion

In this chapter a conclusion is drawn and a discussion. Firstly, in 7-1 the answers to the research questions are presented. Secondly, in Section 7-2 this thesis is compared to related work. Thirdly, in Section 7-3 how future work can build on this thesis is discussed. Thereafter, in Section 7-4 the limitations of this thesis are discussed. Finally, in Section 7-5 a conclusion is drawn on the main goal of this thesis.

7-1 Answers to the Research Questions

In the introduction, we highlighted the problem with magnetometer-free orientation estimation using IMUs: The orientation estimate drifts away from the true orientation and cannot be compensated for in the yaw angle. This problem is less severe for higher-grade IMUs, however, the higher costs relative to lower-grade IMUs limits their usage. Due to the decreasing costs of lower-grade IMUs and increasing availability, a potential solution to improve the orientation estimate using lower-grade IMUs was found: combine multiple lower-grade IMUs to increase motion accuracy at low cost. This led to the main goal of this thesis: How can multiple IMUs be used to improve orientation estimation? This question can be answered by answering the research questions:

How can the accuracy of multiple lower-grade IMUs compare against a higher-grade IMU?
The results in Chapter 4 show that using more IMUs increases the accuracy of the angular velocity and the specific force estimate. The angular velocity accuracy can additionally be increased by larger spatial separation of the accelerometer and at higher angular velocities. It was found that the array configuration has its geometric limitations as the IMUs are placed in a 2D plane. Higher accuracy of the angular velocity can be obtained if they are placed in a 3D plane. A higher-grade IMU was compared against 32 lower-grade IMUs scattered along the surface of a bigger configuration, a cylinder the size a of human arm. This showed that the angular velocity RMSE starts exceeding the higher-grade gyroscope at $\omega \approx \begin{bmatrix}300 & 300 & 300\end{bmatrix}^T$

rad/s (300 rad/s $\approx$ 48 rotations/s). This angular velocity is considered too fast for human motion capture. At a more practical angular velocity of $\omega = \begin{bmatrix}1 & 1 & 1\end{bmatrix}^T$ rad/s the angular velocity RMSE of the lower-grade IMUs using an angular acceleration motion model and the specific force RMSE using a random walk model were 4.27 and 5.16 times higher. From this we can conclude that these lower-grade IMUs, when used for human motion capture, cannot practically outcompete a higher-grade IMU.

Can the parity space method be used for larger configurations? Most previous accounts on the parity space method assumed that the accelerometer only measures the specific force. This assumption may hold for small separation between accelerometers (such as an array with $\approx$ 5mm separation) and low angular velocities. However, at larger configurations, this will lead to unmodeled errors in the accelerometer measurement model. The authors of [51] and [40] have shown an estimator such that the angular motion and relative distances between accelerometers may be taken into account. The extension of this estimator on the parity space has not been explored and was introduced in Chapter 5. This novel parity space approach was tested in a MC simulation in Chapter 6. The results show that the novel parity space approach can reject most of the disturbance when using a larger configuration. As a comparison, the maximum yaw angle deviation during a step disturbance were 4° and 21° for the cylinder (larger) and array (smaller) configuration respectively.

How can the component redundancy of IMUs be used to reject disturbances? As as proof of concept, a cylinder the size of a human arm undergoing a collision was simulated. A step disturbance was added that dissipates over the length of the cylinder. Using this simulation a higher-grade IMU was compared against IMUs distributed equally all over the cylinder. It was found that the cylinder configuration could effectively reject most of the disturbance whereas the higher-grade IMU could not. The peak yaw angle errors of the cylinder configuration and the higher-grade IMU were 4° and 13° respectively. This benefit was obtained by the larger spatial separation of the accelerometers.

7-2 Related Work

The work explained in [39] is similar to the work presented here, however, no use is made of FDI. Instead, the fusion of measurement was incorporated into an Iterative Extended Kalman Filter. This allowed for the addition of gyro-bias tracking. This is useful to compensate for the bias instability for lower-grade IMUs. For longer-term experiments, this would help improve the orientation estimation and the FDI presented in this thesis. Additionally, a bandpass filter as a model for the linear acceleration was included. This helped to separate the specific force into the gravity vector and the linear acceleration for improved orientation estimation. They also added a two-speed attitude integration architecture, which was used to increase computational efficiency. This, however, was used for gyroscope measurement rate of 800 Hz, which far exceeds the measurement rate in this thesis of a 100 Hz. Incorporating the methods described here would help improve the orientation estimation in real life. The parity space method presented in this thesis can be used in its current form with their work as well.

Another paper, [3], is closely related to the novel parity space method proposed in this thesis. In their work, they propose that FDI is not used for a synthetic IMU as scale factors and biases are not compensated for. Furthermore, their application regards pedestrian tracking, which uses closely placed sensors. The residuals obtained are largely due to high dynamic spikes during a gait cycle (when walking, stepping on the ground produces large accelerometer and gyroscope outputs). In this thesis we show that using bigger configurations, this can be used as an advantage instead of a disadvantage. The IMUs scattered across a human arm may experience large accelerometer and gyroscope measurements at one end (due to a collision) but not at the other end. The faulty sensors can be detected and isolated from fusion.

7-3 Future Work

The novel parity space method shown in the thesis has only been shown to work in simulation, it remains to be seen, however, if it will hold up in a real-life scenario. We propose that this novel parity space method should be tested on a robotic arm in which a disturbance is applied to part of the IMUs. Additionally, it is interesting to see how well this parity space method will work on a real human arm. As the skin on a human arm is soft tissue artifact, which means that relative movement between sensors is possible. By placing multiple IMUs on the human arm and moving it around the false positive rate can be analyzed. A specific interesting research would be to test how well this method works on a pitchers arm in baseball. When batting a ball, one part of the arm will undergo a large disturbance. A field test can be setup to compare the novel parity space method against no FDI.

The current parity space method is based solely on the ML estimate of the current time step. A MAP estimate can be found by including motion models of the specific force and the angular acceleration. It would be valuable to see how these motion models affect the ability to accurately identify faulty measurements.

The current parity space method based on component redundancy only uses information of the current time step. There is information to be gained from remembering which sensors were faulty in the past. If a sensor is continuously faulty, it would be logical to assume that the IMU has been either displaced or has been rotated. This sensor can then be permanently removed. Additionally, it could be possible to re-calibrate for the position and orientation error so that in the future the sensor is deemed healthy again. Another, easier approach, could be to increase the covariance matrix of the faulty sensor such that the fault can be modeled as well.

Previous work has assumed that the IMUs are closely placed together. In this thesis we showed that using a larger configuration improves the accuracy of the orientation estimate and improve FDI. The exclusion of magnetometers is due to ferromagnetic materials. When the sensors are placed closely together, they undergo roughly the same disturbance in magnetic field. However, by using a larger configuration, some sensors may be disrupted and other may not. This could allow for the use of magnetometers indoors. By using the parity space approach the magnetometers measurements can be checked against each other in order

to eliminate the disturbance from ferromagnetic materials. This approach has to be validated in an experiment.

7-4 Limitations

Throughout this thesis we have only considered one type of lower-grade IMU and one type of higher-grade IMU. This was done as a proof of concept. Moreover, the theory is not limited to lower-grade IMUs and can also be applied when using any grade of multiple IMUs.

In the current simulation it is assumed that when the accelerometers hit an object, they only undergo a spike in their measurements. In reality, however, it would be more realistic to assume that the accelerometers experience some displacement and/or rotation. These changes in orientation and position on the IMUs would influence future measurement readings that are not taken into account in the current model. This can potentially be resolved by introducing a prior on the measurement covariance matrix of the IMU which has been disturbed for future time steps. The covariance matrix would not only include white noise but would also have to model the other sensor errors.

It is also realistic to assume that when an IMU on one part of the body is determined to be faulty, IMUs that are close to the faulty IMU are more likely to be faulty as well. This knowledge can be used in the form of a prior as well. The IMUs that are close to the faulty IMU can be given a prior on their covariance matrix, for the current and future states, such that it includes other errors than white noise.

Accelerometers and gyroscopes are usually combined into a single IMU, hence if an accelerometer experiences a disturbance, it would be likely that the gyroscopes are disturbed as well. In the current simulation, this has not been taken into account as it was assumed that the gyroscopes can be placed anywhere. If the accelerometer and gyroscope are combined into a single IMU and one of the two is deemed to be a faulty sensor, it could be reasonable to assume that the entire IMU is faulty.

In the current form, where all measurements are fused to form a synthetic IMU, it is not possible to use external measurements such as GPS for FDI. To test the individual IMUs for faults using GPS, the individual navigation solution has to be found, after which it can be fused into a Synthetic IMU. This is a considerable downside of the parity space method where only use is made of the ML estimate.

Throughout the simulations, it has been assumed that the biases were constant and could be compensated for. It might be more realistic to assume that the biases are time-varying and in the current form cannot be included in the state space. The INS scheme would have to be altered slightly such that the fusion of measurements into a synthetic IMU and the motion models are included into the MEKF. Note that this has been implemented by [39], also based on the works of [40], where use was made of an IMU array.

7-5 Conclusion

To conclude, we show that using multiple IMUs provide several benefits to orientation estimation. Firstly, multiple IMUs can be fused into a synthetic IMU to obtain more accurate estimates of the angular velocity and the angular acceleration. Secondly, by using larger configurations FDI can be used to reject disturbances.

Appendix A

Additional Theory

A-1 Orientation Paramaterizations

A-1-1 Euler angles

Building on rotation matrices, Euler angles are three consecutive rotation matrices around specified axes. Different Euler angle parameterizations are possible, one of them is by looking at three consecutive rotations, yaw, pitch and roll (ψ, θ, ϕ) around the (z, y, x) axes respectively. The rotation R^{bs} is thus $R^{bs} = R^{bs}(\phi)R^{bs}(\theta)R^{bs}(\psi)$ which can be expressed as

$$R^{bs} = \underbrace{\begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\phi) & \sin(\phi) \\ 0 & -\sin(\phi) & \cos(\phi) \end{bmatrix}}_{R^{bs}(\phi)} \underbrace{\begin{bmatrix} \cos(\theta) & 0 & -\sin(\theta) \\ 0 & 1 & 0 \\ \sin(\theta) & 0 & \cos(\theta) \end{bmatrix}}_{R^{bs}(\theta)} \underbrace{\begin{bmatrix} \cos(\psi) & \sin(\psi) & 0 \\ -\sin(\psi) & \cos(\psi) & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{R^{bs}(\psi)}, \tag{A-1}$$

or when written as a single matrix it is given as

$$R^{bs} = \begin{bmatrix} \cos(\theta)\cos(\psi) & \cos(\theta)\sin(\psi) & -\sin(\theta) \\ \sin(\phi)\sin(\theta)\cos(\psi) - \cos(\phi)\sin(\psi) & \sin(\phi)\sin(\theta)\sin(\psi) + \cos(\phi)\cos(\psi) & \sin(\phi)\cos(\theta) \\ \cos(\phi)\sin(\theta)\cos(\psi) + \sin(\phi)\sin(\psi) & \cos(\phi)\sin(\theta)\sin(\psi) - \sin(\phi)\cos(\psi) & \cos(\phi)\cos(\theta) \end{bmatrix}. \tag{A-2}$$

As the Euler angles are based on the Rotation matrices the same properties hold and thus the Euler angles are also in SO(3). Contrary to rotation matrices, however, Euler angles are prone to singularities and are not unique in describing rotations. A phenomenon called gimbal lock can take place which is when two rotations describe the same rotations around an axis [4, 25]. Another problem that is frequent with Euler angles is that wrapping also induces non-uniqueness.

A-1-2 Rotation vectors

Every rotation can be described by a single rotation around a single axis. This rotation can be expressed as a normalized unit vector n which lies in the group $\mathbb{S}^2$ with a rotation angle

α in $\mathbb{R}$. A mapping can be made that maps the axis-handle from representation into $SO(3)$ [21]. This mapping can be denoted by $\exp(-\alpha[n\times])$.

The rotated vector x_* can be expressed in terms of the original vector x, a rotation α and a normalized vector n as

$$x_* = x\cos\alpha + n(x\cdot n)(1-\cos\alpha) - (n\times x)\sin\alpha, \tag{A-3}$$

where a detailed derivation can be found in [25]. By introducing the notation $[x\times]$ allows one to write the skew-asymmetric matrix as

$$[x\times] \triangleq \begin{bmatrix} 0 & -x_3 & x_2 \\ x_3 & 0 & -x_1 \\ -x_2 & x_1 & 0 \end{bmatrix}. \tag{A-4}$$

Given 3 arbitrary vectors u, v and w, multiple cross products can be expressed as terms of dot products as

$$u\times(v\times w) = v(w\cdot u) - w(u\cdot v). \tag{A-5}$$

Combining equations (A-4) and (A-5) with (A-3) allows one to rewrite the rotation vector as

$$x_* = (\mathcal{I}_3 - \sin\alpha[n\times]) + (1-\cos\alpha)[n\times]^2)x, \tag{A-6}$$

leading to the compact expression of

$$x_* = \exp(-\alpha[n\times])x. \tag{A-7}$$

Which is an equivalent to a rotation matrix if it were parameterized in terms of α and n [25]. There is, however, a problem that arises when using rotation matrices. The problem is that some rotation cannot be described by the axis-handle paramaterization. By looking at the inverse mapping which is given by (tr is the trace operator)

$$\alpha = \cos^{-1}(\frac{1}{2}(\mathrm{tr}(R)-1)), \tag{A-8a}$$

$$n = \frac{1}{2\sin(\alpha)}\begin{bmatrix} R_{32}-R_{23} \\ R_{13}-R_{31} \\ R_{21}-R_{12} \end{bmatrix}, \tag{A-8b}$$

it can be observed that when $\sin(\alpha)$ goes to 0 the unit vector n goes to infinity. Thus the axis-handle parameterization becomes invalid for angles $\alpha = k\pi$, for $k = 0, 1, ...$ [4].

A-1-3 Quaternions

A unit quaternion is an orientation parameterization in $q \in \mathbb{R}^4$, however, in reality, orientation is in $\mathbb{R}^3$. This means in order to use the quaternions we must exploit a trick such that our coordinate vector can be rewritten as follows

$$\bar{x}^s = \begin{bmatrix} 0 \\ x^s \end{bmatrix}. \tag{A-9}$$

A rotation with unit quaternions is defined as

$$\bar{x}^b = q^{sb} \odot \bar{x}^s \odot (q^{sb})^c, \tag{A-10}$$

where $(q^{sb})^c$ denotes the complex conjugate of the unit quaternion, defined as

$$\begin{bmatrix} q_0 \\ q_v \end{bmatrix}^c = \begin{bmatrix} q_0 \\ -q_v \end{bmatrix}. \tag{A-11}$$

The quaternion multiplication $\odot$ is denoted by

$$p \odot q = \begin{bmatrix} p_0 q_0 - p_v \cdot q_v \\ p_0 q_v + q_0 p_v + p_v \times q_v \end{bmatrix} = p^L q = q^R p, \tag{A-12}$$

where

$$p^L \triangleq \begin{bmatrix} p_0 & -p_v^T \\ p_v & p_0 \mathcal{I}_3 + [p_v \times] \end{bmatrix}, \qquad p^R \triangleq \begin{bmatrix} p_0 & -p_v^T \\ p_v & p_0 \mathcal{I}_3 - [p_v \times] \end{bmatrix}. \tag{A-13}$$

Combining (A-4) and (A-13) into (A-10) results into the following

$$\bar{x}^b = (q^{bs})^L (q^{sb})^R \bar{x}^s = \begin{bmatrix} 1 & 0_{1x3} \\ 0_{3x1} & q_v q_v^T + q_0^2 \mathcal{I}_3 + 2 q_0 [q_v \times] + [q_v \times]^2 \end{bmatrix} \begin{bmatrix} 0 \\ x_v \end{bmatrix}, \tag{A-14}$$

By taking the axis-handle representation and rewriting it in quaternion form, one can choose to express it as follows

$$q^{bs}(n^s, \alpha) = \begin{bmatrix} \cos(\frac{\alpha}{2}) \\ -n^s \sin(\frac{\alpha}{2}) \end{bmatrix}, \tag{A-15}$$

such that (A-14) can be written in the form of

$$\bar{x}^b = \begin{bmatrix} 1 & 0_{1x3} \\ 0_{3x1} & \mathcal{I}_3 - \sin(\alpha)[n^s \times] + (1 - \cos(\alpha))[n^s \times]^2 \end{bmatrix} \begin{bmatrix} 0 \\ x^s \end{bmatrix}. \tag{A-16}$$

Which is an equivalent expression to (A-6).

A-1-4 Orientation Deviations

The mapping of quaternions into an SO(3) group can be found by choosing

$$\exp_q(\bar{\eta}) = \begin{bmatrix} \cos ||\eta||_2 \\ \frac{\eta}{||\eta||_2} \sin ||\eta||_2 \end{bmatrix}, \tag{A-17}$$

which allows one to find the mapping expressed as

$$\exp_q([\eta \times]) = \mathcal{I}_3 + \sin(||\eta||_2) \left[\tfrac{\eta}{||\eta||_2} \times\right] + (1 - \cos(||\eta||_2) \left[\tfrac{\eta}{||\eta||_2} \times\right]^2. \tag{A-18}$$

The rotation quaternion can be defined as

$$q_t^{nb} = \exp_q(\frac{\bar{\eta}_t^n}{2}) \odot \tilde{q}_t^{nb}. \tag{A-19}$$

The reverse mapping can be expressed as

$$\eta = \log_q(q) = \frac{\arccos(q_0)}{||q_v||_2} q_0. \tag{A-20}$$

Finally, use can be made of approximations using the small angle theorem which are as follows

$$\exp_q(\eta) \approx \begin{bmatrix} 1 \\ \eta \end{bmatrix}, \qquad \log_q(q) \approx q_v. \tag{A-21}$$

A-2 Gauss-Newton Optimization

The Gauss-Newton optimization is an iterative optimization method that can be applied to non-linear equations. The Gauss-Newton optimization algorithm is suitable as a solver for a non-linear least squares problem when a good initial estimate is available and the residuals are small [17, 22]. Furthermore, the Gauss-Newton has the benefit that hessian is not computed as it is approximated using the Jacobian which in turn also guarantees positive semi-definiteness [25]. The Gauss-Newton algorithm is explained below

Consider a nonlinear measurement equation of the form

$$y = h(x) + e, \tag{A-22}$$

where y are the observations, x is the state space and $e \sim \mathcal{N}(0, \Sigma)$. The likelihood is then given by

$$p(y|x) = \frac{1}{(2\pi)^{\frac{N}{2}}\sqrt{\det(\Sigma)}} \exp\left(-\frac{1}{2}(y - h(x))^T \Sigma^{-1}(y - h(x))\right). \tag{A-23}$$

Then, under the assumption that Σ is time-invariant, the ML estimation can be found by maximizing the argument of the likelihood function,

$$\hat{x} = \arg\max_x p(y|x), \tag{A-24}$$

which is equivalent to minimizing the cost function $V^{NWLS(x)}$

$$\arg\min_x V^{NWLS}(x) = \arg\min_x \left(y - h(x))^T \Sigma^{-1}(y - h(x))\right), \tag{A-25a}$$

$$\arg\min_x V^{NWLS}(x) = \arg\min_x \left(e^T(x)\Sigma^{-1}e(x)\right), \tag{A-25b}$$

The gradient $\mathcal{G}$ and Hessian $\mathcal{H}$ of the cost function $V^{NWLS}(x)$ can be calculated using the Jacobion $J(x)$. The Jacobian of the cost function is

$$J(x) = \frac{\partial e^T(x)}{\partial x} = -\frac{\partial h^T(x)}{\partial x}. \tag{A-26}$$

The gradient and Hessian are defined as

$$\underbrace{\frac{dV^{NWLS}(x)}{dx}}_{\mathcal{G}} = J(x)\Sigma^{-1}e(x), \tag{A-27a}$$

$$\underbrace{\frac{d^2V^{NWLS}(x)}{dx^2}}_{\mathcal{H}} = J(x)\Sigma^{-1}J^T(x) + \underbrace{e^T(x)\Sigma^{-1}\frac{d^2e(x)}{dx^2}}_{\approx 0}. \tag{A-27b}$$

The second term from the Hessian of the cost function can be approximated as zero under the assumption that the ML estimate of $\hat{x}$ is close enough to the true value of x. Given an initial estimate , a Gauss-Newton optimization algorithm can literately find the minimum using the following iteration

$$x^{k+1} = x^k + \mathcal{H}^{-1}\mathcal{G}, \tag{A-28a}$$

$$x^{k+1} = x^k + (J(x)\Sigma^{-1}J^T(x))^{-1}J(x)\Sigma^{-1}(y - h(x)). \tag{A-28b}$$

A-3 Probability Distributions

The ability to detect a change or a disturbance depends on what the Probability Density Function (PDF) looks like [37]. One of the most commonly used PDF is the normal distribution. It is also assumed in the previous Chapters that all measurements are Gaussian. The sum of multiple normally distributed variables is also normally distributed, however, summing squared normally distributed variables leads to a chi-square distribution. If the variables were zero-mean, then a central chi-square distribution forms, however, if the variables were not zero-mean than a non-central chi-square distribution forms. The different distributions are shown below.

A-3-1 Normal Distribution

The normal distribution also known as the Gaussian distribution is denoted by

$$X \sim \mathcal{N}(\mu, \sigma^2), \tag{A-29}$$

where X is the random stochastic variable, μ is the mean value and σ^2 is the variance. The PDF is shown in Figure A-1 and is defined as

$$f_X(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2\sigma^2}(x-\mu)^2\right), \quad -\infty < x < \infty. \tag{A-30}$$

The probability of X falling in between an interval $[a, b]$ can be found by taking the integral from a to b as

$$P(a \leq X \leq b) = \int_a^b f_X(x)dx. \tag{A-31}$$

The scalar case in (A-30) can be extended into a vector such that X becomes of length n. Giving rise to the more general formulation

$$f_X(x) = \frac{1}{(2\pi)^{\frac{n}{2}}\sqrt{\det(Q)}} \exp\left(-\frac{1}{2}(x-\mu)^T Q^{-1}(x-\mu)\right), \quad -\infty < x < \infty. \tag{A-32}$$

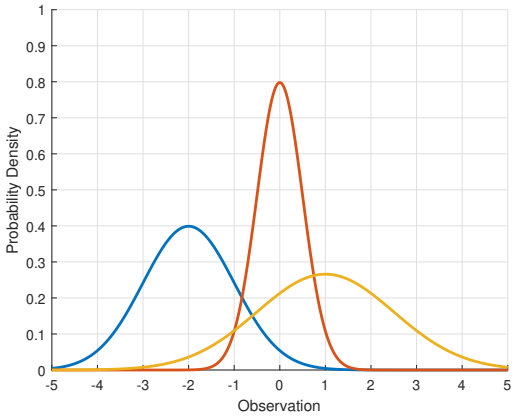

Figure A-1: Normal Probability Density Function for varying μ and σ^2.

A-3-2 Chi-square distributions

The Chi-square distribution arises when the sum is taken of squared normally distributed variables X_i which are also Independent and Identically Distributed (IID) [37], i.e. $Y = \sum_{i=1}^{\nu} X_i^2$. If $X - i$ has zero mean then one gets the central Chi-square distribution, if they have non-zero mean one gets the noncentral Chi-square distribution.

Central Chi-square distribution If the squared variables happen to have zero-mean, i.e. $X_i \sim \mathcal{N}(0, Q)$, then a central Chi-square distribution $Y \sim \chi_\nu^2$ is formed. The degrees of freedom, denoted by ν is assumed to be a positive integer [37] can be seen as how many IID squared variables are summed and as ν goes to infinity the Chi-square distribution goes towards the Normal distribution.

The probability density function is strictly positive and takes the form of

$$f_X(x) = \begin{cases} \frac{1}{2^{\frac{\nu}{2}}\Gamma(\frac{\nu}{2})} x^{\frac{\nu}{2}-1} e^{-\frac{1}{2}x} & x > 0, \\ 0 & x < 0, \end{cases} \tag{A-33}$$

where $\Gamma(u)$, also known as the Gamma function, is defined as

$$\Gamma(u) = \int_0^\infty t^{u-1} e^{-t}. \tag{A-34}$$

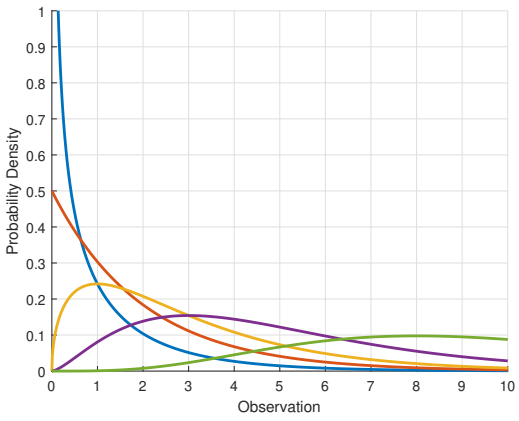

Figure A-2: Central Chi-square Probability Density Function for varying ν.

Non-central Chi-square distribution If, however, the variables has non-zero mean, i.e. $X_i \sim \mathcal{N}(\mu_i, Q)$ then a Non-central Chi-square distribution $Y \sim \chi_\nu^2(\lambda)$ is formed. Where λ is defined as the noncentrality parameter $\lambda = \sum_{i=1}^{\nu} \mu_i^2$. When $\lambda = 0$ the non-central Chi-square distribution takes shape of the central Chi-square distribution [37].

Again, the probability density function is strictly positive but now takes the shape of

$$p(x) = \begin{cases} \frac{1}{2}\left(\frac{x}{\lambda}\right)^{\frac{\nu-2}{4}} I_{\frac{\nu}{2}-1}(\sqrt{\lambda x}) \exp\left(-\frac{1}{2}(x+\lambda)\right) & x > 0, \\ 0 & x < 0, \end{cases} \tag{A-35}$$

where $I_r(u)$ is defined as

$$I_r(u) = \frac{\left(\frac{1}{2}u\right)^r}{\sqrt{\pi}\Gamma\left(r+\frac{1}{2}\right)} \int_0^\pi \sin^{2r}(\theta) e^{u\cos\theta} d\theta. \tag{A-36}$$

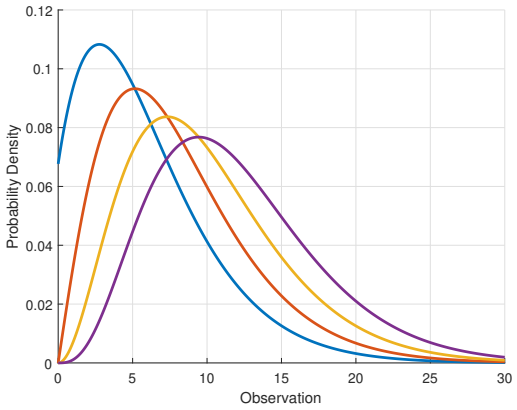

Figure A-3: Non-central Chi-square Probability Density Function for varying ν.

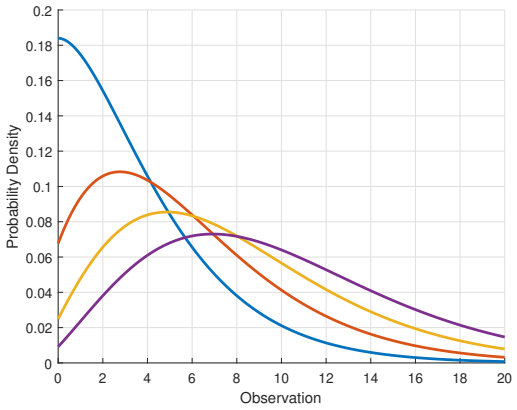

Figure A-4: Non-central Chi-square Probability Density Function for varying λ.

Appendix B

Additional Results

B-1 Biases, covariance of array

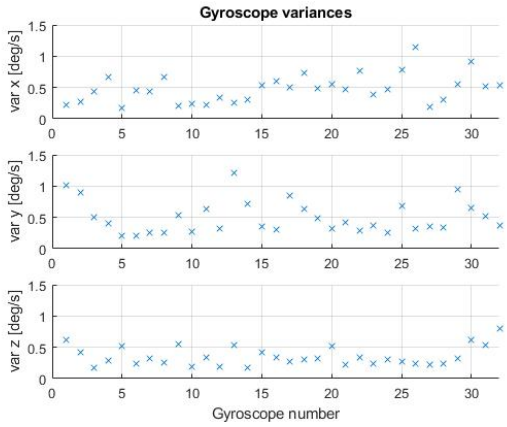

Figure B-1: Gyroscope covariance of the MIMU4884 array. The x- (top), y- (middle) and z-(axes) are shown with their respective gyroscope number.

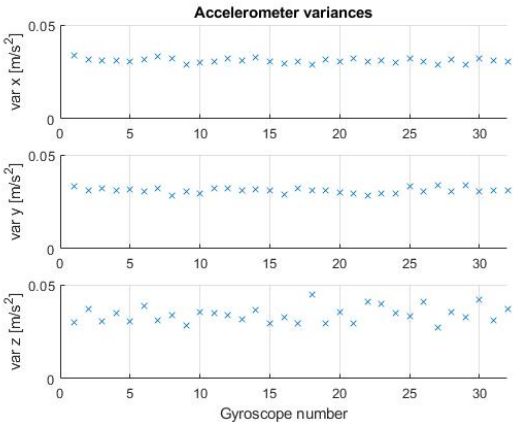

Figure B-2: Accelerometer covariance of the MIMU4884 array. The x- (top), y- (middle) and z-(axes) are shown with their respective accelerometer number.

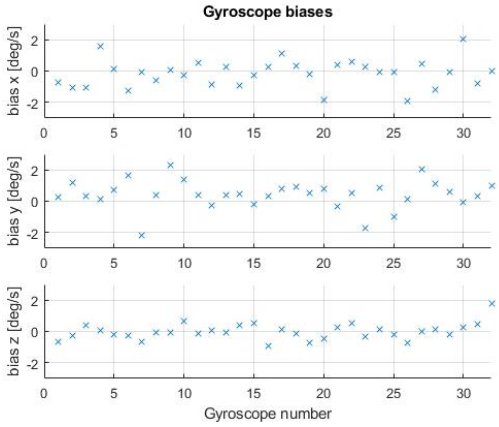

Figure B-3: Gyroscope biases of the MIMU4884 array. The x- (top), y- (middle) and z-(axes) are shown with their respective gyroscope number.

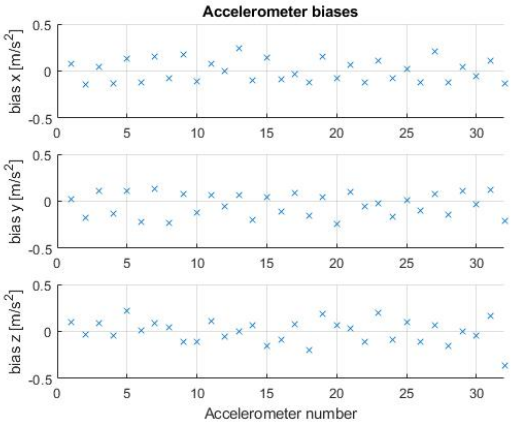

Figure B-4: Accelerometer biases of the MIMU4884 array. The x- (top), y- (middle) and z-(axes) are shown with their respective gyroscope number.

B-2 Additional Results

motion model	Angular velocity magnitude in x- y- and z- axis [rad/s]				
Angular velocity	$\omega = 10^0$	$\omega = 10^1$	$\omega = 10^2$	$\omega = 10^3$	$\omega = 10^4$
ω no motion model	11.191	11.16	8.579	1.431	0.148
ω random walk	8.798	8.774	6.745	1.125	0.116
$\dot{\omega}$ motion model	10.582	10.556	8.285	1.429	0.148
$\dot{\omega}$ random walk	10.303	10.279	8.139	1.428	0.148
Specific force					
s no motion model	5.390				
s random walk	4.237				

Table B-1: Performance comparison different motion models of the MIMU4844 IMU array of Inertial Elements. The performance is indicated by the x-axis ratio of the RMSE with respect to the MTi 300 IMU from Xsens.

motion model	Angular velocity magnitude in x- y- and z- axis [rad/s]				
Angular velocity	$\omega = 10^0$	$\omega = 10^1$	$\omega = 10^2$	$\omega = 10^3$	$\omega = 10^4$
ω no motion model	11.444	11.412	8.715	1.434	0.148
ω random walk	8.997	8.971	6.852	1.128	0.117
$\dot{\omega}$ motion model	10.813	10.786	8.416	1.432	0.148
$\dot{\omega}$ random walk	10.525	10.5	8.268	1.431	0.148
Specific force					
s no motion model	5.399				
s random walk	4.245				

Table B-2: Performance comparison different motion models of the MIMU4844 IMU array of Inertial Elements. The performance is indicated by the y-axis ratio of the RMSE with respect to the MTi 300 IMU from Xsens.

motion model	Angular velocity in x- y- and z- axis [rad/s]				
Angular velocity	$\omega = 10^0$	$\omega = 10^1$	$\omega = 10^2$	$\omega = 10^3$	$\omega = 10^4$
ω no motion model	11.191	9.794	1.644	0.161	0.016
ω random walk	8.798	7.699	1.292	0.126	0.013
$\dot{\omega}$ motion model	6.352	5.854	1.547	0.161	0.016
$\dot{\omega}$ random walk	5.761	5.324	1.503	0.161	0.016
Specific force					
s no motion model	5.390				
s random walk	4.237				

Table B-3: Performance comparison different motion models of the IMU cylinder. The performance is indicated by the ratio of the x-axis RSME with respect to the MTi 300 IMU from Xsens

motion model	Angular velocity in x- y- and z- axis [rad/s]				
Angular velocity	$\omega = 10^0$	$\omega = 10^1$	$\omega = 10^2$	$\omega = 10^3$	$\omega = 10^4$
ω no motion model	11.444	8.665	1.636	0.162	0.017
ω random walk	8.997	6.811	1.286	0.127	0.013
$\dot{\omega}$ motion model	4.489	3.954	1.362	0.161	0.017
$\dot{\omega}$ random walk	4.025	3.559	1.286	0.161	0.017
Specific force					
s no motion model	5.399				
s random walk	4.245				

Table B-4: Performance comparison different motion models of the IMU cylinder. The performance is indicated by the ratio of the y-axis RSME with respect to the MTi 300 IMU from Xsens

B-3 Additional results

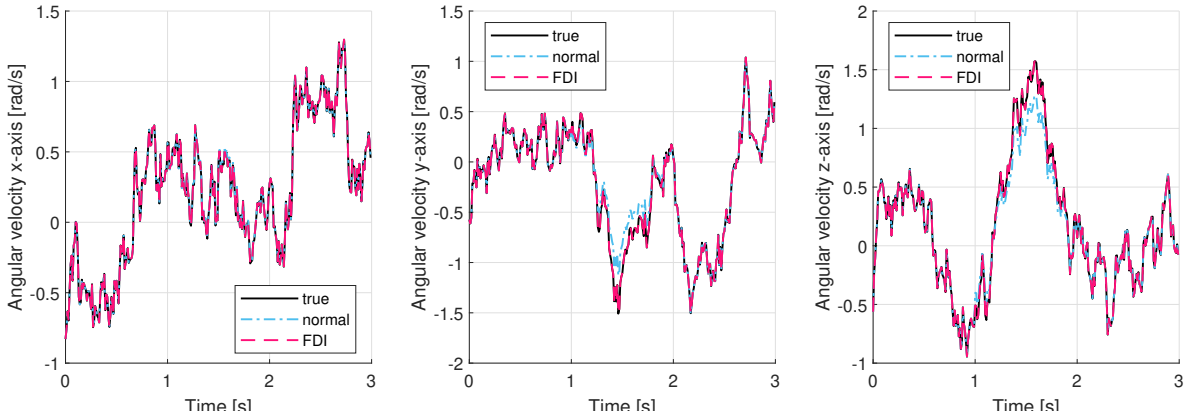

Figure B-5: True angular velocity states (black) and mean angular velocity states of the synthetic IMU with (blue) and without (pink) FDI using 1000 MC realizations for the x-axis (left), y-axis (mid) and z-axis (right) of the body frame.

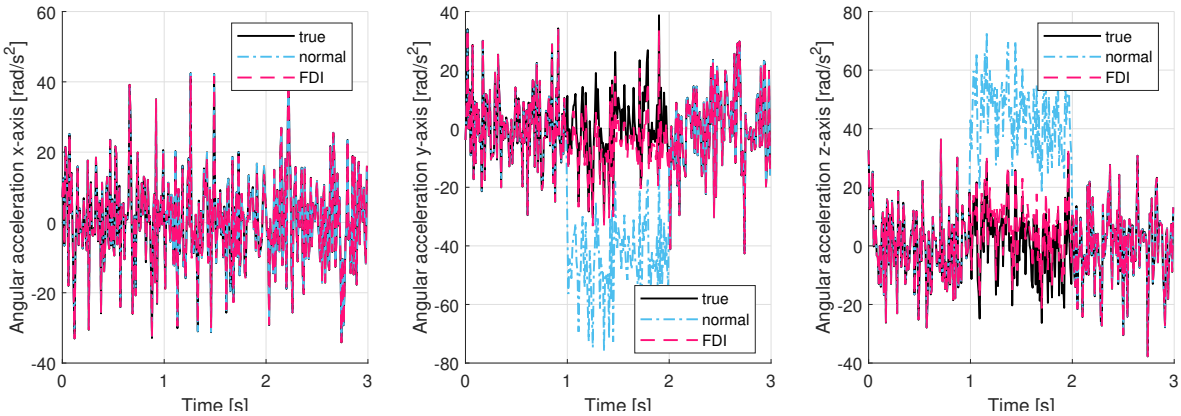

Figure B-6: True angular acceleration states (black) and mean angular acceleration states of the synthetic IMU with (blue) and without (pink) FDI using 1000 MC realizations for the x-axis (left), y-axis (mid) and z-axis (right) of the body frame.

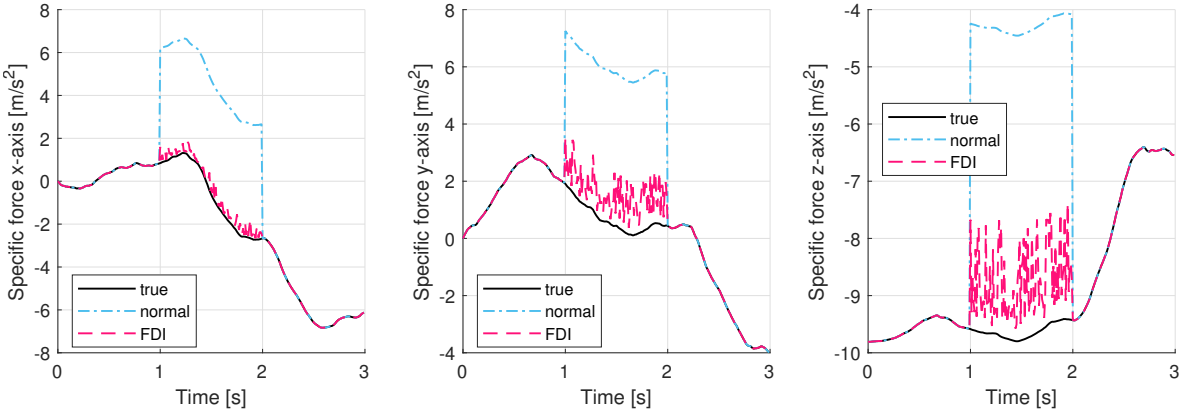

Figure B-7: True specific force states (black) and mean specific force states of the synthetic IMU with (blue) and without (pink) FDI using 100 MC realizations for the x-axis (left), y-axis (mid) and z-axis (right) of the body frame.

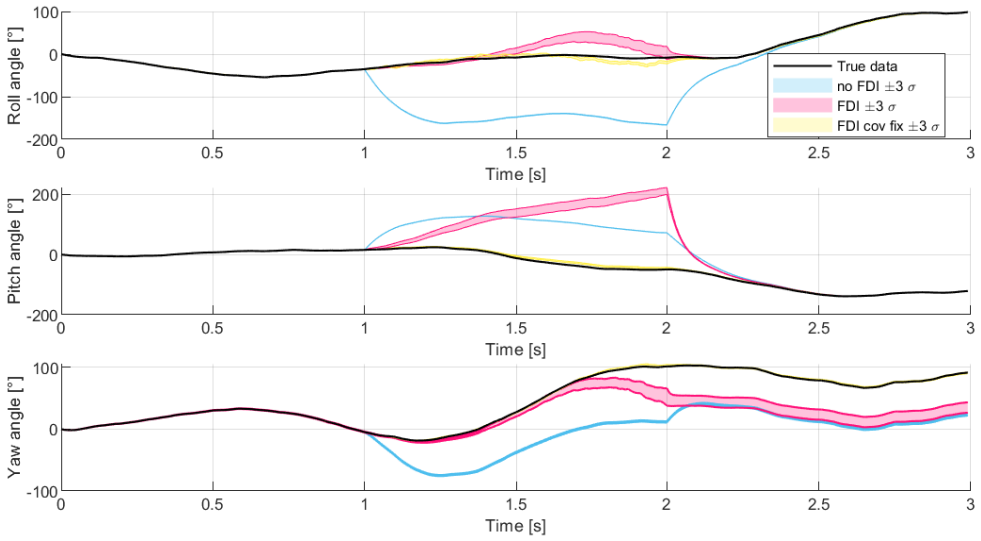

Figure B-8: MC simulation of 1000 realizations with an applied disturbance. The synthetic IMUs undergo no FDI (blue), and undergo FDI (pink and yellow). Pink has the origin at the centre of the cylinder and yellow at the centre of the accelerometers. The estimated and true angles (left) and the error of the angles (right) are the roll (top), pitch (mid) and yaw (bottom).

Appendix C

Additional Algorithms

Algorithm 6: Position calibration for a sensor array

Set reference position ρ_1 ;
Set initial estimate for $\hat{\rho}_{2:N_a}$, set $j = 1$;
while *Not converged* **do**
  for $k = 1 : N$ **do**
    ML estimate for $\hat{x}_k$ can be found using Algorithm 5.

$$\hat{x}_k = \arg\max_x p(Y_k|x_k, \hat{\rho}_{1:N_a}),$$
$$\Omega_k = [\hat{\omega}_k\times]^2 + [\dot{\hat{\omega}}_k\times].$$

  end

$$C = \left(\sum_{k=1}^{N} \Omega_k^T \Omega_k\right)^{-1}$$

  for $i = 2 : N_a$ **do**

$$\hat{\rho}_i = C * \left(\sum_{k=1}^{N} \Omega_k^T \left(y_{a_i,k} - \hat{s}_k\right)\right).$$

  end
  $j = j + 1$
end

Algorithm 4: Simulation of rigid body motion to generate accelerometer and gyroscopes measurement data.

Set $g^n = \begin{bmatrix} 0 & 0 & 9.81 \end{bmatrix}^T$ and initial condition $q_0 = \begin{bmatrix} 1 & 0 & 0 & 0 \end{bmatrix}^T$;
Assign values for $a^b_{1:N}$, $\omega^b_{1:N+1}$, $b_{a,1:N_a}$ and $b_{\omega,1:N_\omega}$;
for $k = 1 : N$ **do**
  Calculate the angular acceleration;

$$\dot{\omega}^b_k = \frac{\omega^b_{k+1} - \omega^b_k}{T}.$$

  Update quaternion, ;

$$q^{nb}_k = \exp_q(\frac{\omega^b_k}{2}) \odot q^{nb}_{k-1},$$

  Find rotation matrix ;

$$R^{bn}_k = \left[q_v q_v^T + q_0^2 \mathcal{I}_3 + 2q_0[q_v\times] + [q_v\times]^2 \right].$$

  Ω^b_k is constant for every accelerometer;

$$\Omega^b_k = \left([\omega^b_k\times]^2 + [\dot{\omega}^b_k\times] \right).$$

  Set accelerometer measurements;
  for $i = 1 : N_a$ **do**

$$y_{a_i,k} = R^{s_i b} \left(\left(a^b_k - R^{bn}_k g^n \right) + \Omega^b_k \rho_i \right) + b^{s_i}_{a_i} + e^{s_i}_{a_i,k}.$$

  end
  Set gyroscope measurements;
  for $j = 1 : N_\omega$ **do**

$$y_{\omega_j,k} = R^{s_j b} \omega^b_k + b^{s_j}_{\omega,j} + e^{s_j}_{\omega_j,k}.$$

  end
end

Algorithm 5: Fusion of the MIMU measurements to estimate ω_k and θ_k

Pre-process the data ;
Initialize ω_0, $k = 0$;

$$\hat{\omega}_0 = \left((\mathbf{1}_{N_\omega}^T \otimes \mathcal{I}_3) Q_\omega^{-1} (\mathbf{1}_{N_\omega} \otimes \mathcal{I}_3)\right)^{-1} (\mathbf{1}_{N_\omega}^T \otimes \mathcal{I}_3) Q_\omega^{-1} Y_\omega,$$
$$H_{s,i} = \begin{bmatrix} [-\rho_i \times] & \mathcal{I}_3 \end{bmatrix}, \qquad H_\omega = \mathbf{0}_{3N_\omega,6}, \qquad J_{\omega,h} = \mathbf{1}_{N_\omega} \otimes \mathcal{I}_3.$$

Stack ;

$$H = \begin{bmatrix} H_{s_1}^T & \dots & H_{s_{N_s}}^T & H_\omega^T \end{bmatrix}^T,$$

while *Not converged* **do**
  for $i = 1 : N_s$ **do**

$$J_{s_i,h}(\hat{\omega}_k) = \left([[\rho_i \times]\hat{\omega}_k \times] - [\hat{\omega}_k \times][\rho_i \times]\right),$$
$$h_{s,i}(\hat{\hat{\omega}}_k) = [\hat{\omega}_k \times]^2 \rho_i.$$

  end

$$h_{\omega,i}(\hat{\omega}_k) = \mathbf{1}_{N_\omega} \otimes \hat{\omega}_k.$$

  Stack ;

$$J_h(\hat{\omega}_k) = \begin{bmatrix} J_{s_1,h}^T(\hat{\omega}_k) & \dots & J_{s_{N_s},h}^T(\hat{\omega}_k) & J_{\hat{\omega},h}^T(\hat{\omega}_k) \end{bmatrix}^T,$$
$$h(\hat{\omega}_k) = \begin{bmatrix} h_{s_1}^T(\hat{\omega}_k) & \dots & h_{s_{N_s}}^T(\hat{\omega}_k) & h_\omega^T(\hat{\omega}_k) \end{bmatrix}^T.$$

  Gauss-Newton iteration ;

$$\hat{\omega}_{k+1} = \hat{\omega}_k + \left((J_h^T(\hat{\omega}_k) P J_h(\hat{\omega}_k)\right)^{-1} J_h^T(\hat{\omega}_k) P \left(Y - h(\hat{\omega}_k)\right).$$

  $k = k + 1$
end

$$\hat{\theta} = (H^T Q^{-1} H)^{-1} H^T Q^{-1} (Y - h(\hat{\omega}_k))$$

Bibliography

[1] Xsens Technologies B.V., 2016. Accessed on 2020 September 29.

[2] Inertial Elements, 2016. Accessed on 2020 September 30.

[3] Jared B. Bancroft and Gérard Lachapelle. Data fusion algorithms for multiple inertial measurement units. *Sensors*, 11(7):6771–6798, 7 2011.

[4] RB Bhat, RV Dukkipati, and M Pascal. *Advanced Dynamics*, volume 55. Cambridge University Press, 2002.

[5] Yoram Bresler. Maximum likelihood estimation of a linearly structured covariance with application to antenna array processing. pages 172–175. Publ by IEEE, 1988.

[6] Carmen M.N. Brigante, Nunzio Abbate, Adriano Basile, Alessandro Carmelo Faulisi, and Salvatore Sessa. Towards miniaturization of a MEMS-based wearable motion capture system. *IEEE Transactions on Industrial Electronics*, 58(8):3234–3241, 2011.

[7] Philippe Cardou and Jorge Angeles. Angular velocity estimation from the angular acceleration matrix. *Journal of Applied Mechanics, Transactions ASME*, 75(2):0210031–0210038, 3 2008.

[8] Håkan Carlsson, Isaac Skog, and Joakim Jaldén. On-the-fly geometric calibration of inertial sensor arrays. *2017 International Conference on Indoor Positioning and Indoor Navigation, IPIN 2017*, 2017-Janua(September):1–6, 2017.

[9] Generalitat De Catalunya. Redundant Imus for Precise Trajectory Determination. Technical report, 2004.

[10] James Chaffee and Jonathan Abel. GDOP and the Cramer-Rao bound. In *Record - IEEE PLANS, Position Location and Navigation Symposium*, pages 663–668. Publ by IEEE, 1994.

[11] Ward Cheney and David Kincaid. Linear Algebra: Theory and Applications. page 624, 2010.

[12] Ronald Christensen. Testing fisher, neyman, pearson, and bayes. *American Statistician*, 59(2):121–126, 5 2005.

[13] Naser El-Sheimy and Xiaoji Niu. The Promise of MEMS to the Navigation Community. *Inside GNSS*, pages 46–56, 2007.

[14] Thomas Forsberg, Niklas Grip, and Natalia Sabourova. Non-iterative calibration for accelerometers with three non-orthogonal axes, reliable measurement setups and simple supplementary equipment. *Measurement Science and Technology*, 24(3), 2013.

[15] Eric Gagnon, Alexandre Vachon, and Yanick Beaudoin. Data fusion architectures for orthogonal redundant inertial measurement units. *Sensors (Switzerland)*, 18(6), 2018.

[16] Janos J. Gertler. *Fault detection and diagnosis in engineering systems*. 2017.

[17] Philip E. Gill and Walter Murray. Algorithms for the Solution of the Nonlinear Least-Squares Problem. *SIAM Journal on Numerical Analysis*, 15(5):977–992, 10 1978.

[18] Stéphane Guerrier. Improving accuracy with multiple sensors: Study of redundant MEMS-IMU/GPS configurations. Technical report, 2009.

[19] Stéphane Guerrier, Dr. Jan Skaloud, and Adrian Wägli. Integration of Skew-Redundant MEMS-IMU with GPS for Improved Navigation Performance. Technical Report February 2015, 2008.

[20] Fredrik Gustafsson. *Statistical Sensor Fusion*. 2013.

[21] Hashim A. Hashim. Special Orthogonal Group SO(3), Euler Angles, Angle-axis, Rodriguez Vector and Unit-Quaternion: Overview, Mapping and Challenges. (3), 2019.

[22] Saïd HILOUT. GaussNewton method. *Numerical Methods for Equations and its Applications*, pages 220–251, 2012.

[23] Intertial Elements. Osmium MIMU4844.

[24] Amin G. Jaffer. Maximum Likelihood Direction Finding of Stochastic Sources: a Separable Solution. In *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, pages 2893–2896. IEEE, 1988.

[25] Manon Kok, Jeroen D. Hol, and Thomas B. Schön. Using inertial sensors for position and orientation estimation. *Foundations and Trends in Signal Processing*, 11(1-2):1–153, 2017.

[26] Manon Kok and Thomas B. Schon. Magnetometer calibration using inertial sensors. *IEEE Sensors Journal*, 16(14):5679–5689, 7 2016.

[27] Zhiping Lin, Qiyue Zou, E. Sally Ward, and Raimund J. Ober. Cramer-Rao lower bound for parameter estimation in nonlinear systems. *IEEE Signal Processing Letters*, 12(12):855–858, 12 2005.

[28] F. Landis Markley. Attitude error representations for Kalman filtering. *Journal of Guidance, Control, and Dynamics*, 26(2):311–317, 2003.

[29] John Olof Nilsson and Isaac Skog. Inertial sensor arrays - A literature review. In *2016 European Navigation Conference, ENC 2016*. Institute of Electrical and Electronics Engineers Inc., 8 2016.

[30] Fredrik Olsson, Manon Kok, Kjartan Halvorsen, and Thomas B. Schon. Accelerometer calibration using sensor fusion with a gyroscope. In *IEEE Workshop on Statistical Signal Processing Proceedings*, volume 2016-Augus. IEEE Computer Society, 8 2016.

[31] G. Panahandeh, I. Skog, and M. Jansson. Calibration of the accelerometer triad of an inertial measurement unit, maximum likelihood estimation and Cramér-Rao bound. In *2010 International Conference on Indoor Positioning and Indoor Navigation, IPIN 2010 - Conference Proceedings*, 2010.

[32] Michael Perlmutter and Laurent Robin. High-performance, low cost inertial MEMS: A market in motion! pages 3–7.

[33] Mark E. Pittelkau. Calibration and attitude determination with Redundant Inertial Measurement Units. *Journal of Guidance, Control, and Dynamics*, 28(4):743–752, 2005.

[34] Umar Qureshi and Farid Golnaraghi. An Algorithm for the In-Field Calibration of a MEMS IMU. *IEEE Sensors Journal*, 17(22):7479–7486, 2017.

[35] Ryan Hanson. Using Multiple MEMS IMUs to form a Distributed Inertial Measurement Unit. Technical report, 2005.

[36] Paul G. Savage. Strapdown inertial navigation integration algorithm design part 1: Attitude algorithms. *Journal of Guidance, Control, and Dynamics*, 21(1):19–28, 1998.

[37] Sailes K. Sengupta and Steven M. Kay. *Fundamentals of Statistical Signal Processing: Estimation Theory*, volume 37. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1995.

[38] Derek K. Shaeffer. MEMS inertial sensors: A tutorial overview. *IEEE Communications Magazine*, 51(4):100–109, 2013.

[39] Young Young Shen and Allison P. Anderson. Extended Kalman Filter for Magnetometer-Free Estimation of Spacesuit Wearer Joint Kinematics Using Inertial Measurements. Technical report, 2019.

[40] Isaac Skog, John Olof Nilsson, Peter Handel, and Arye Nehorai. Inertial Sensor Arrays, Maximum Likelihood, and Cramér-Rao Bound. *IEEE Transactions on Signal Processing*, 64(16):4218–4227, 9 2016.

[41] Petre Stoica and Arye Nehorai. On the concentrated stochastic likelihood function in array signal processing. Technical Report 5, 1995.

[42] MARK STURZA. Skewed axis inertial sensor geometry for optimal performance. American Institute of Aeronautics and Astronautics (AIAA), 10 1988.

[43] Salah Sukkarieh, Peter Gibbens, Ben Grocholsky, Keith Willis, and Hugh F. Durrant-Whyte. A low-cost, redundant inertial measurement unit for unmanned air vehicles. *International Journal of Robotics Research*, 19(11):1089–1103, 11 2000.

[44] Shuyan Sun, Xiaoli Meng, Lianying Ji, Zhipei Huang, and Jiankang Wu. Adaptive Kalman filter for orientation estimation in micro-sensor motion capture. *Fusion 2011 - 14th International Conference on Information Fusion*, 2011.

[45] David Törnqvist. Statistical Fault Detection with Applications to IMU Disturbances, 2006.

[46] David Törnqvist and Fredrik Gustafsson. Eliminating the initial state for the generalized likelihood ratio test. *IFAC Proceedings Volumes (IFAC-PapersOnline)*, 6(PART 1):599–604, 1 2006.

[47] David Törnqvist, Fredrik Gustafsson, and Inger Klein. GLR tests for fault detection over sliding data windows. *IFAC Proceedings Volumes (IFAC-PapersOnline)*, 16(1):743–748, 1 2005.

[48] Michel Verhaegen and Vincent Verdult. *Filtering and system identification: A least squares approach*, volume 9780521875. Cambridge university press, 2007.

[49] Adrian Waegli, Stéphane Guerrier, and Jan Skaloud. Redundant MEMS-IMU integrated with GPS for performance assessment in sports. Technical report, 2008.

[50] Adrian Waegli, Jan Skaloud, Stéphane Guerrier, Maria Eullia Parés, and Ismael Colomina. Noise reduction and estimation in multiple micro-electro-mechanical inertial systems. *Measurement Science and Technology*, 21(6), 2010.

[51] Johan Wahlström, Isaac Skog, and Peter Händel. Inertial Sensor Array Processing with Motion Models. *2018 21st International Conference on Information Fusion, FUSION 2018*, pages 788–793, 2018.

[52] Oliver J Woodman. Number 696 An introduction to inertial navigation An introduction to inertial navigation. 2007.

Glossary

List of Acronyms

IMUs	Inertial Measurement Units
MEMS	Micro-Electromechanical System
KF	Kalman Filter
EKF	Extended Kalman Filter
MEKF	Multiplicative Extended Kalman Filter
ML	Maximum Likelihood
MAP	Maximum A Posteriori
CRLB	Cramér-Rao Lower Bound
FDI	Fault Detection and Isolation
IID	Independent and Identically Distributed
GLRT	Generalised Likelihood Ratio Test
MLRT	Marginalised Likelihood Ratio Test
FIM	Fisher Information Matrix
SLS	Separable Least Squares
PDF	Probability Density Function

List of Symbols

n	Navigation frame
b	Body frame
s	Sensor frame
R^{bn}	Rotation matrix from navigation to body frame
q^{bn}	Quaternion from navigation to body frame
η	Orientation deviation
T	Sampling time
ω	Angular velocity
$\dot{\omega}$	Angular acceleration
s	Specific force
g	Gravity vector
Σ	Covariance matrix
y	Measurement output
b	Measurement bias
e	Measurement noise
x	State vector
D	Misalignment error matrix
Y	Stacked measurements
$h(\omega)$	Stacked non-linear measurements
H	Stacked linear measurements
E	Stacked errors
ρ	accelerometer positions
$\mathcal{I}(x)$	Fisher Information Matrix of x
Q	Diagonally stacked covariance matrices of measurements
J	Jacobian
N	Number of sensors
$\mathcal{H}_0$	Null Hypothesis
$\mathcal{H}_1$	Alternate Hypothesis
P_D	Detection probability
P_C	Rejection probability
P_{FA}	False alarm probability
P_M	Missed detection probability
ϵ	Prediction error
r	Residual
$\mathcal{P}$	Projection matrix
$\mathcal{B}_{H_{\omega,\perp}}$	Orthonormal basis
F	Fault
$[x\times]$	skew-symmetric matrix
$\mathrm{Det}(x)$	Deterinant of x
$\mathrm{Cov}(x)$	Covariance matrix of x
$\mathcal{N}$	Normally distributed
$\sim$	Distributed as
$\mathcal{I}$	Identity matrix