

A Robust Complementary Filter Approach for Attitude Estimation of Unmanned Aerial Vehicles using AHRS

Johann Meyer, Kreeelan Padayachee, Benjamin A. Broughton

Abstract Most attitude filters utilise the accelerometer as a vector measurement of gravity to estimate pitch and roll angles; however, during accelerated flight, this assumption does not hold. This paper develops a robust attitude filter model that determines when the accelerometers can be used as vector measurements by using a time-dependent model of steadiness. Furthermore, a gyro measurement model is developed using a Gaussian random walk model as the basis for bounding bias estimates and rejecting improbable estimates that arise from slow dynamics or transitioning from steady to dynamic motion. Monte-Carlo simulations using test cases with dynamic motion were performed to verify its performance. The bias is accurately tracked and gyro integration performance during unsteady motion ultimately improved. Moreover, roll dynamics were tracked more accurately than current state-of-the-art complementary filters.

1 Introduction

State estimation is important for guidance, navigation, and control as it enables the operator of a system or the system itself to make appropriate decisions based on the

Johann Meyer
Incomar Aerospace and Defence Systems, Grevillea Street 2, Highveld Techno Park, Centurion, South Africa,
Delft University of Technology, Mekelweg 5, Delft, Netherlands,
e-mail: johann.meyer@gmail.com

Dr Benjamin A. Broughton
Incomar Aerospace and Defence Systems, 2 Grevillea Street Highveld Techno Park, Centurion, South Africa, e-mail: bbroughton@iadsystems.co.za

Kreeelan Padayachee
Incomar Aerospace and Defence Systems, 2 Grevillea Street Highveld Techno Park, Centurion, South Africa, e-mail: kpadayachee@iadsystems.co.za

current state of the vehicle. It is, moreover, a basic requirement during flight testing and for any subsequent system identification analysis. Other applications include pedestrian navigation, gait analysis, and virtual/augmented reality.

Unmanned Aerial Vehicles (UAVs) or, more specifically, Mini and Micro UAVs, have limited payload and power budgets and, therefore, in many cases rely on using small low-power Micro Electro-mechanical Sensors (MEMS). The challenge of these low-cost MEMS sensors is the higher noise levels in the signals and the drifting of the measured signals from the true signals. With these higher levels of noise and faster drift rates, it is even more important to fuse the results from multiple sensors to be able to filter out the attitude signal from the noise signal. If the attitude signal is not correctly ascertained from the state estimator, the control system may send inappropriate commands to the control surfaces and the navigation system will not work as intended.

The ideal state filter is the optimal processing of a system's sensors such that the true state value is obtained. A state filter can be subdivided into two parts: deciding how to integrate measurements and deciding when to integrate measurements. This paper will focus on when to integrate the sensor measurements, with primary application to aircraft state estimation.

Attitude filters using gravity (along with magnetic field measurements) as a vector measurement are referred to as MARG (Magnetometer, Angular Rate and Gravity) filters, which is a subset of AHRS (Attitude and Heading Reference System) filters. These MARG filters assume the accelerometer measurement is an accurate vector measurement of the gravity vector. However, this is only true during steady translational flight. During rotational and accelerated flight, this vector measurement is corrupted by the additional accelerations at the accelerometer location and thus these manoeuvres can be considered as disturbances to the measurement of the true gravity vector. Moreover, the gyro cannot correct for offsets as it does not utilise an external reference and therefore only when the vectors are unperturbed again, is it possible for the attitude estimate to become accurate again.

This paper aims to develop a method for determining when the aircraft is in a steady state such that the accelerometer may only be used when the assumptions surrounding it hold. Furthermore, it will develop a method for robustly detecting biases in the gyro measurements. The basis of the filter is similar to that found in [25]; however, the principles can be extended to almost any complementary attitude filter utilising the accelerometer measurement as a vector measurement or any filter attempting to estimate biases in the gyro. Magnetic distortion will not be addressed in this paper, although the problem has been orthogonalised by using a model that limits magnetic distortion effects to the yaw-axis, thereby making it no longer a safety-critical problem in most applications.

2 Previous Work

The process of obtaining an attitude estimate from vector measurements is termed static attitude estimation. Attitude filters require at least two non-collinear vectors to uniquely estimate the attitude of a vehicle. In spacecraft, often more than two vector observations are available and filters solving Wahba's problem are then required to produce an attitude estimate from the over-constrained set of measurements. The survey in reference [3] mostly focuses on these over-constrained problems. In other words, the solutions in [3] act as a data compression technique for the vector measurements, where the output is a static attitude estimate. Vector measurements are weighted when solving Wahba's problem but the impact of erroneous vector measurements cannot be constrained to a specific axis. In the case of aircraft and most land-based systems, only a magnetometer and accelerometer are usually present, therefore the problem is fully-constrained and no optimisation step is necessary. Abstractly, most attitude filters below are a hybrid between a static attitude estimate and a dead-reckoning approach.

A Kalman filter (KF) is a statistically optimal linear state estimator. This means the weights will optimally weigh the contributions from the different sensors. The Extended KF (EKF) is an extension of the linear KF for nonlinear problems, where a local linearisation is used to enable the use of the KF. The EKF is, however, not optimal for nonlinear problems. It is, nonetheless, the de facto benchmark for state estimation. The optimality of the KF comes at a computational cost. Computing the Kalman gain involves a matrix inversion of an $m \times m$ matrix, where m is the number of measurements. There have been filters that aimed to minimise this cost by reducing the size of the matrix. Examples of these are the Murrell Multiplicative EKF (described in [14, 18]) and the Sequential Multiplicative EKF[18]. The dimensions of the measurement matrix can also be reduced by using a static attitude estimator to reduce the dimensionality of the measurements, as mentioned before. Alternatively, the Information Filter avoids the matrix inversion by propagating the inverse of the covariance matrix, also known as the Fisher Information Matrix.

Most Kalman filters assume error states as these filters enable small angle and linearisation approximations to be made with little impact on the accuracy[11]. These Kalman filters are often referred to as Error-state or Indirect Kalman filters. Error states also enable singular attitude representations to be used which may provide computational and accuracy benefits. The accuracy benefits arise due to the requirement for redundant schemes needing to incorporate additional constraints, like the quaternion norm or the orthogonality of the direction cosine matrix. The choice of the representation of the error state varies between literature[17, 24] and the application, as computational cost and interpretability are in conflict. Another important consideration for KFs, or any attitude filter, is how to define the attitude error. The error can be defined either additively or multiplicatively. Additively is how errors are normally treated in the KF; however, they have no basis for attitude errors[13]. The alternative defines the attitude error as a rotation and is, ultimately, more intuitive. It results in lower computational cost and higher accuracy and is, thus, the basis of most modern KFs[13]. Particle and Unscented Kalman Filters are not pro-

posed as solutions, as they are more computationally expensive than KFs by a factor determined by the number of particles and are still ultimately limited by the gravity vector assumption.

Complementary filters (CF) are an alternative to Kalman filters that offer faster performance at the cost of statistical optimality. Most complementary filters operate using fixed gains[10, 12]; however, adaptive gains have been seen in more recent literature[2, 4, 7, 10, 22, 23, 25, 27]. A Kalman filter with a fixed Kalman gain can be shown to be equivalent to a complementary filter[5]. Therefore, the adaptive gain approaches can be seen as approximations of the Kalman filter without the optimality criterion. The complementary filter is termed so, because of the filter using two sets of measurements with complementary spectral characteristics. Internally, most CFs only differ in the error term that is fed back[10, 12] or the method to handle and/or the heuristic used to detect distortions to the vector measurements[2, 4, 7, 10, 22, 23, 25, 27]. The basis filter for the revised algorithm of [10] is akin to that found in [12].

The problem with attitude filters in general are the strong assumptions surrounding them that often do not hold for most applications. The most common assumptions are that the accelerometers and magnetometers are measuring the Earth's gravitational and magnetic fields, respectively. The problem with the accelerometer assumption, in particular, is that the aircraft is often in a dynamic state, where non-gravitational accelerations are being measured and thereby corrupting the gravitational direction vector. This in turn violates the Gaussian noise assumptions behind the Kalman filter. There are Kalman filters that attempt to correct for these accelerations[4, 9, 22] but they require an independent airspeed velocity measurement to estimate the linear and centripetal accelerations. At the far-side of this spectrum, some approaches, for example [15], do not use the accelerometer as a vector measurement at all but as a part of a larger kinematic model, which enables more states to be recovered but at the cost of more instrumentation to ensure observability and, in turn, more computation.

A common approach to improving the filter performance when these additional measurements are not possible is by reducing the weight given to the accelerometer by either reducing the gain in the CF or increasing the measurement noise associated with the accelerometer in the KF[21, 22, 23, 25, 26]. The problem with these approaches is that after long periods of corruption, the corrections will have completely leaked into the attitude estimate, as described by reference [26]. Effectively, these filters are adapting the cut-off frequency of the signal. The gyro can also not compensate for these generated offsets, therefore the estimate will remain corrupted until the accelerometer becomes uncorrupted again. Some of the CF approaches eventually set the gain to zero; however, the problem arises when misclassification of the gravity occurs. This is particularly a problem for slow accelerated motion (e.g. an aircraft in a gentle, steady coordinated turn) which is difficult to distinguish from the high noise levels present in MEMS accelerometers. It is for these reasons that interpolation and gain-scheduling is deemed computationally unnecessary and problematic. In the revised algorithm in [10], the accelerometer update step is eliminated when the accelerometer exceeds a threshold value for some period of time.

Note that if the update step of the Kalman filter is not performed for an extended period of time the Kalman gain will rely heavily on the next update step due to the growth of the uncertainty in the covariance matrix. This over-reliance on a single update is detrimental when considering the noise levels associated with a MEMS accelerometer.

For gyro bias estimation, the literature is often divided into two main sets: integral compensation (e.g. [12]) and low-pass filter estimation (e.g. [25]). Integral compensation regulates all the errors between the gyro and the static attitude estimate. The downside of this approach is when the static attitude estimate is corrupted by accelerations or magnetic distortion all the errors are assumed to be as a result of the gyro. In contrast, low-pass filtering the gyro measurements enables one to determine only the bias associated with the gyro. However, if the vehicle is in dynamic motion, then the low-pass filter estimate will also be corrupted. The benefit of the separation of the errors in the low-pass filter approach is that it enables the application of a model for bias evolution, as will be presented in this paper.

In order to compensate for magnetic distortion, some filters implement a threshold parameter. However, as of late, the trend is to limit the magnetic distortion effects to the yaw angle estimate [10, 22, 25]. Using a filter structure such as this, enables the question of magnetic distortion compensation to be eliminated from the attitude problem. An alternative is presented in reference [1], where the magnetometer is only checked for quasi-steadiness so that it can be used as a reference for bias removal. Without an inertial reference for yaw, however, the already present drift in the attitude estimate cannot be corrected.

3 Method

This section will explain the structure of the proposed attitude filter model. In addition, a brief explanation of the datasets and the approach taken to model the errors in the sensors is given.

3.1 Filter Model

This section will elaborate on the filter model used. First, a high-level overview of the proposed filter structure will be given and the differences between it and the filter structure used in [25], hereafter referred to as the Valenti filter, will be discussed. Thereafter, the specifics of each part of the proposed filter will be elaborated.

The Valenti filter integrates the gyro results, then computes the complementary filter gain for the accelerometer based on the magnitude error of the accelerometer. If the gain is greater than 0, the accelerometer measurement is fused; otherwise, the path short-circuits to the magnetometer update. To incorporate the measurements, the measurements are rotated according to the best-available attitude estimate and

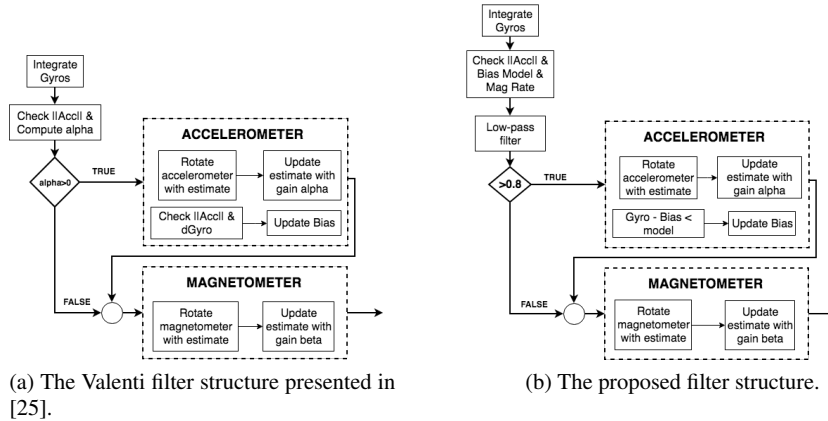


Fig. 1 A comparison of the filter structures.

then the correction is partially applied using interpolation. The bias update is a low-pass filter of the gyro measurements when the magnitude of the accelerometer and differences in the rates measured by the gyros are all within a threshold.

The proposed filter is similar with respect to how the measurements are incorporated, but differs by when they are incorporated. The modified filter is significantly more cautious of incorporating the accelerometer measurement than the original version. It utilises the Swiss cheese model of risk management by stacking fallible checks to produce a more robust check. This approach is akin to the process of boosting in machine learning. Three checks are used at present: accelerometer magnitude error, rate of change of the magnetic field vector, and a model of the gyro bias. If this check passes, the Boolean result is low-pass filtered and only if this filtered result becomes greater than a threshold, the accelerometer results are utilised. The low-pass filter leads to a temporal steadiness check that becomes less sensitive to dynamic motion that transiently satisfies the thresholds. The temporal check is much like the model used in [10] but it differs in that the low-pass filter makes the filter less sensitive to noise and disturbances that may reset the filter counter too quickly. The proposed model is also more averse to using the accelerometer compared to the model in [10], as that model uses the temporal term to stop using the accelerometer as opposed to start using it. Moreover, only if the accelerometer measurement can be trusted, does the model consider performing a bias update. The aforementioned bias model in the check will have determined whether or not the update is realistic by comparing it to a model of the bias evolution. The use of the bias model in this way avoids slow dynamics, such as the aperiodic roll or a slow turn, from being absorbed as a bias. This approach should also be preferred over the integral dynamics approach as the error terms can be more carefully managed, especially during these slow manoeuvres.

For clarity, the quaternion that the filter returns is the passive transformation from the body axis to the inertial axis and therefore the output should be conjugated before being used as an attitude estimate.

3.1.1 Gyro Update

The gyro update, or the so-called prediction step, is performed by integrating the corrected angular rate measurements ($\tilde{\omega}$) using (4) or (5), in combination with (2). The measured angular rate ($\hat{\omega}$) is corrected for bias B (where B is a function of time) using (1). If the 1st-order approximate form (5) is used, then care must be taken to re-normalise the quaternion. For most applications, the approximate form should be sufficient under the small angle assumption. Moreover, the noise levels in the measurements will most likely have a larger impact on the estimate's accuracy.

$$\tilde{\omega} = \hat{\omega} - B(t) \quad (1)$$

$$\dot{q}_{IB}^t = -\frac{1}{2} \begin{bmatrix} 0 \\ \tilde{\omega} \end{bmatrix} \otimes q_{IB}^{t-\Delta t} \quad (2)$$

$$q_{IB}^t = q_{IB}^{t-\Delta t} \otimes \exp\left(\frac{\tilde{\omega}\Delta t}{2}\right) \quad (3)$$

$$= q_{IB}^{t-\Delta t} \otimes \begin{bmatrix} \cos\left(\frac{\|\tilde{\omega}\|\Delta t}{2}\right) \\ \frac{\tilde{\omega}}{\|\tilde{\omega}\|} \cdot \sin\left(\frac{\|\tilde{\omega}\|\Delta t}{2}\right) \end{bmatrix} \quad (4)$$

$$\approx q_{IB}^{t-\Delta t} + \dot{q}_{IB}^t \cdot \Delta t \quad (5)$$

3.1.2 Accelerometer Update

The true normalised gravity vector is given by (6). In typical aircraft setups, the IMU is aligned with the aircraft body-axis, and it is therefore important to remember that the accelerometer measurement needs to be negated for this to be the correct reference vector.

$$\bar{g} = [0 \ 0 \ 1]^T \quad (6)$$

Next, the normalised accelerometer measurement \hat{a} should be rotated with the current estimate of the attitude quaternion (7). The transpose of the rotation matrix means the vector observation is being rotated and not the coordinate frame. If the estimate \tilde{a} is correct then the estimate should equal the reference accelerometer measurement \bar{g} . If drift has occurred, then the required rotation to transform \tilde{a} into \bar{g} is given by (8) and (9).

$$\tilde{a} = R^T(\hat{q}_{IB}) \cdot \hat{a} \quad (7)$$

$$\hat{q}_s = \frac{\sqrt{\tilde{a}_z + 1}}{\sqrt{2}} \quad (8)$$

$$\delta \hat{q}_{acc} = \left[\hat{q}_s \frac{-\tilde{a}_y}{2\hat{q}_s} \frac{\tilde{a}_x}{2\hat{q}_s} 0 \right]^T \quad (9)$$

Since it is desired to combine the gyro and accelerometer estimates, it is necessary to use Linear Interpolation (LERP) or Spherical LERP (SLERP), where the interpolation is between the correction quaternion and the identity quaternion. The gain for the interpolation is denoted as α in [25]. In the paper by Valenti et. al.[25], it was recommended to switch between the two interpolation schemes using the scalar component of the quaternion. The scalar quaternion is related to the projection of the two vectors and is indicative of the magnitude of the required rotation. However, reasonably good results are obtained by only using LERP, since the errors should be small, except if a significant amount of drift in the attitude estimate has occurred.

3.1.3 Magnetometer Update

The reference magnetic field is given by (10), where Γ is the square of the magnetic field strength in the xy-plane. This reference simplifies some of the later equations; however, it is important to note that the reference is now magnetic North and not true North as used by most navigation systems. It is also important to note that the correction that will be applied will only impact yaw, meaning the z-component will remain unaffected by the correction. This is critical to ensure the magnetometer distortion only affects the yaw estimate. It also means that the filter does not need to know the inclination of the local magnetic field; however, as a consequence, the filter will not be able to use this information to its benefit either. At large latitudes, the impact of magnetic distortion and attitude errors will begin to dominate the Γ term.

$$\tilde{m} = \left[\sqrt{\tilde{m}_x^2 + \tilde{m}_y^2} \ 0 \ \tilde{m}_z \right]^T = \left[\sqrt{\tilde{\Gamma}} \ 0 \ \tilde{m}_z \right]^T \quad (10)$$

As with the accelerometer update, the magnetometer measurement must be rotated according to the current estimate of the attitude, which includes the rotation generated by the accelerometer, if it was used during the current time-step. Integration errors in pitch and roll and accelerometer updates that are corrupted by accelerations other than gravity will mean that $\tilde{\Gamma}$ will in general not be equal to $\tilde{\Gamma}$. Therefore, the yaw estimate can only ever be as good as the estimate coming from the accelerometer and the gyro. The estimate of $\delta \hat{q}_{mag}$ is combined using LERP/SLERP with the gain parameter β as in [25].

$$\tilde{m} = R^T(\hat{q}_{IB}) \cdot \hat{m} \quad (11)$$

$$\delta \hat{q}_{mag} = \left[\frac{\sqrt{\tilde{\Gamma} + \tilde{m}_x \sqrt{\tilde{\Gamma}}}}{\sqrt{2\tilde{\Gamma}}} \ 0 \ 0 \ \frac{\tilde{m}_y}{\sqrt{2(\tilde{\Gamma} + \tilde{m}_x \sqrt{\tilde{\Gamma}})}} \right]^T \quad (12)$$

3.1.4 Steadiness Model

During dynamic motion, there may be short periods where the accelerometer will read a magnitude of 1g and potentially consider it a steady state. During a typical phugoid motion, for example, momentary 1g values at the inflection points of the periodic motion are experienced. These corrections tend to drive the estimated pitch angle towards zero and, as a result of the method used to remove the bias, these corrections would leak into the bias estimates of the gyro. These biases would then further corrupt the gyro integration, which is critical for the unsteady portions of the flight.

By using a low-pass filter on the Boolean check, steadiness can no longer occur during dynamic motion. Even though more reliance is being placed on the gyros as a result of there being fewer steady periods, this approach should be preferred as the assumption that the accelerometer is measuring the gravity vector is being violated and should, in any case, not be used. Since the filtering is applied to the Boolean check, it is possible to construct other checks that can augment or replace the current checks, thereby making it more adaptable than simply applying it to the measurements directly. Note that it is not necessary to low-pass filter the Boolean check but alternatives like an increment and decrement approach can work; however, then bounds checking should be implemented to avoid wind-up effects from long periods of steady or unsteady motion.

γ_{steady} and γ_{unsteady} are the low-pass filter dynamics for transitioning in and out of steadiness. t_{steady} is a threshold used to determine when the aircraft is steady. A reasonable value for t_{steady} is 0.7-0.8.

An important question is how to determine what constitutes steady motion. Approximate 1g conditions can occur during accelerated flight as previously stated and therefore additional checks are required. The first considers the use of a gyro bias model and the second the time derivative of the low-pass filtered magnetometer measurements, as described below. The accelerometer measurement is also low-pass filtered prior to being used in the check.




Fig. 2 Illustration of the exponential transition from unsteady to steady and vice-versa as a result of the low-pass filter dynamics.

3.1.5 Gyro Bias Model

Often with accelerometer-based complementary filters, pseudo steady-states result in dynamics or slow motion being erroneously classified as biases. This is particularly common with roll dynamics. This is in part due to the absence of an absolute measure for yaw rates. Such attitude filters commonly employ gyro thresholding (i.e. $\omega^t - \omega^{t-1} < t_\omega$) to detect steady states and, as a consequence, the gyro cannot distinguish a bias from a constant rate turn. Moreover, the noise levels in the accelerometer make it difficult to differentiate between a gentle turn and noise. The rate of change of the magnetometer solves some of these problems but still suffers from transition periods from steady to dynamic motion. Having a model of plausible updates means the filter can then detect and reject erroneous updates.

The most common model for describing the bias of a gyro is through a random walk model (e.g. [8] and [6]). This model assumes that quasi-deterministic drift factors in the bias, for example thermal sensitivity, are already adjusted for, making this model an accurate representation of the bias. A well-known model for the bounds of a Gaussian random walk is given by $c_\sigma \cdot \sigma \sqrt{n}$, where c_σ is the confidence bound for a Gaussian distribution (e.g. $c_\sigma = 2$ gives a 95% confidence interval), σ is the standard deviation of the random walk, and n is the number of iterations since the last bias update. Figure 3 illustrates a Monte-Carlo simulation of a Gaussian random walk and the growth of the assumed bounds for different values of c_σ .

When implemented, the gyro bias is continuously estimated such that the true bias is already being tracked when the first bias update occurs, resulting in an ultimately lower tracking error, especially if the filter becomes unsteady shortly after becoming steady. Note, however, that since no high-frequency estimate of the bias is being provided, the low-pass filter will always lag behind the true value, but it is assumed that the bias evolves slow enough such that the lag is not noticeable or detrimental to filter performance. Furthermore, the model enables a higher bias gain to be used without suffering from the dynamics leaking into the bias estimate. Note that the bias being estimated is only the component that arises from the rate random walk, which means that even if perfect tracking accuracy is obtained, an angle




Fig. 3 The evolution of the uncertainty in the position of a 1D agent over 100 runs, where each step is taken according to a Gaussian distribution.

random walk will still occur due to the integration of the Gaussian noise present in the gyro measurements. The accelerometer and magnetometer measurements are responsible for removing the error caused by the angle random walk.

In order to apply the model, the model bounds need to be reduced after updates. A first approach may be to reduce the bound to zero after an update; however, the bounds also dictate steadiness and usually more updates are desirable since the bias is being estimated from a noisy process. Therefore, low-pass filtering n to zero, with the same dynamics as the gyro bias, is the approach adopted in this paper.

3.2 Simulation Setup

In order to test the filter, data was generated in the JSBSim¹ flight dynamics simulator using the default Cessna C172x aircraft model. Since the data was simulated, the correct attitudes were known. However, to test the filter, the data was passed onto the filter after being realistically corrupted prior to usage. All three instruments were modelled as experiencing Gaussian noise, but bias was only added to the gyro data, which was modelled as a random walk. Results from an aileron doublet, an elevator doublet, and a coordinated turn are shown in the following section.

4 Results & Discussion

This section will present and discuss the results obtained from the simulated measurements described in Section 3.2. The state estimation performance of the proposed filter will be compared with that of the Valenti filter. For the readers reference, two other popular filters have been included for comparison but will not be analysed in depth. The two filters are: the commonly-used original Madgwick filter (chapter 3 in [10]) and the Multiplicative EKF (magnetic distortion not limited to the yaw-axis). All filters in this section assume the accelerometer measures the gravity vector. The attitude filter gains for all filters were kept constant for all manoeuvres to better demonstrate the robustness of the filters in a real flight situation. In addition, all filters are initialised with the known true attitude; however, this is only of significance in the coordinated turn test case. Next, outliers were determined by (13), where x is the data point, Q denotes the quartile, and IQR is the inter-quartile range. Finally, the noise realisations for the single stochastic simulations were not the same but chosen such to illustrate a worst-case scenario.

$$\text{outliers} \in \begin{cases} x > Q_3 + 1.5 \times IQR \\ x < Q_1 - 1.5 \times IQR \end{cases} \quad (13)$$

¹ JSBSim homepage: <http://jsbsim.sourceforge.net/>

4.1 Impact of Time-dependent Steadiness

It is clear from Fig. 4 that the proposed filter can be said to perform on average better than the Valenti filter with a 95% confidence interval for the simulated drift rates. However, it is also clear that the results from the proposed filter have a larger variance. The larger variance arises because the proposed filter is purely integrating the gyro values (Fig. 7), whose biases continue to drift while the bias estimate remains constant (Fig. 5). The angle random walk from integrating the white noise in the gyro measurement is also a contributing factor. The Valenti filter, in contrast, periodically updates using a non-drifting but corrupted source (accelerometer), which thereby reduces this variance.

The root-mean-square error (RMSE) is a useful metric for overall performance. However, the behaviour of the filter should be evaluated to ensure that the filter is producing meaningful results. Figure 7 demonstrates that the Valenti filter damped out the roll response of the aircraft, whereas the proposed filter managed to avoid filtering out these dynamics. It is clear that long periods of dynamic motion would ultimately suffer from drift, as expected for any dead-reckoning-based algorithm. If the accelerometer gain in the Valenti filter were increased, the corrections in the pitch axis will be incorrectly driven towards zero, as in Fig. 6. In cases where the expected behaviour is not well known, the Valenti filter will, as a result, be harder to tune and may perform substantially worse.

Filters based on regulating an error term tend to merge the acceleration distortion with the bias estimate. The Valenti filter solved this problem with the low-pass filter on the gyro, but it can be observed in Fig. 5 that without a robust scheme to detect dynamic motion, dynamics will leak into the estimate.

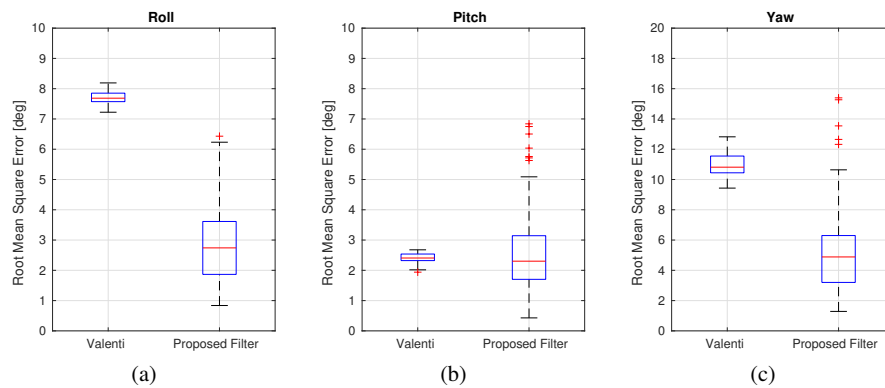
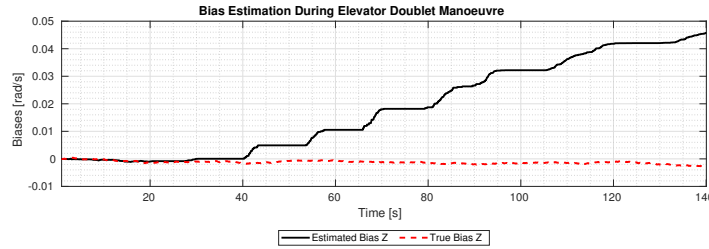
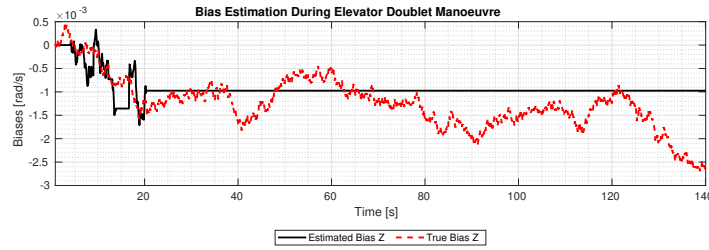


Fig. 4 A Monte-Carlo simulation of 50 runs for the elevator doublet test case. The wider spread for the proposed filter arises because of the stochastic drift.



(a) Valenti filter



(b) Proposed filter

Fig. 5 A comparison of the bias tracking performance for the different filter models. The Valenti filter bias tracking error is an order of magnitude larger than that of the proposed filter.

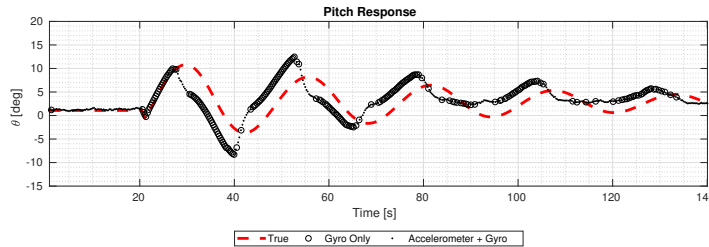



Fig. 6 Valenti Filter pitch response when a larger accelerometer gain is used.


4.2 Impact of Gyro Model

One problem with slow dynamic motion is that it may be incorrectly classified as a bias. This was seen in the previous set of results; however, it can be more clearly determined from Fig. 8 at around 20 seconds. Note that the accelerometer is not used after 5 seconds due to the inclusion of the rate of change of the magnetic field, which is assessed in more detail in Section 4.3.

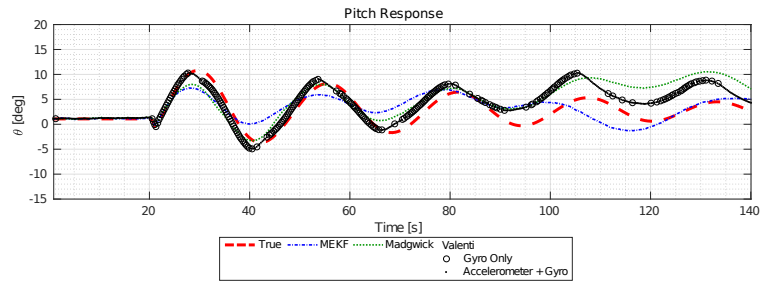
Furthermore, it is clear that the injected doublet at the 5 second mark is also being partially classified as a bias, which ultimately deteriorates the subsequent gyro integration performance. The integration performance would be substantially worse if the gain for bias estimation was increased. This result means the gain would need to be partially selected based on the manoeuvres performed and not purely based




(a) Reference filters



(b) Proposed filter



(c) Reference filters




(d) Proposed filter


Fig. 7 A comparison of the attitude estimates for a single stochastic simulation of the filters to a elevator doublet. The proposed filter error is dominated by drift. The Valenti filter drives the roll angle towards zero and corrects at the pseudo-static inflection points of the angular rates.

on the properties of the gyro. Figure 11 illustrates the model's ability to reject the doublet motion and, ultimately, the pseudo steady-state at 20 seconds. If the bias model is only used for updating the bias and not in the steadiness check, Fig. 9 will result. This small offset in roll is usually not detrimental but the impact would be dependent on the accelerometer gain used, meaning another parameter would need to be selected based on the manoeuvres performed.

Finally, Fig. 10 shows the impact the removed roll response has on the heading estimate. In order for the assumption to treat the heading and attitude determination as a decoupled problem to be valid, the error impact on yaw must be limited to avoid the heading performance from being bounded by the attitude determination algorithm. This conclusion supports the use of the proposed filter.



(a)



(b)

Fig. 8 The impact of not using a bias model. Note this includes the magnetic field rate. Furthermore, the bias' low-pass filter gain was selected to minimise the impact of the command input on the bias estimate.

4.3 Impact of Rate of Change of Magnetic Field

The aforementioned results demonstrated the issue of roll angle determination; however, it is also necessary to assess a shortcoming of the steady determination in the Valenti filter for this particular application. During a steady turn, the measured rates will be constant and the acceleration noise will also make it difficult to differentiate

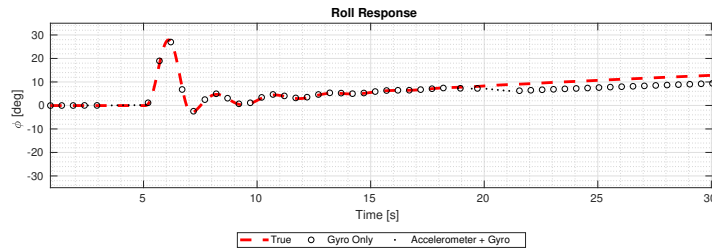


Fig. 9 The impact of using the bias model only for bias estimate updates. The filter is unable to distinguish a constant rate turn from a steady state. The unsteady state post 20 seconds arises due to the inclusion of the magnetic field rate.

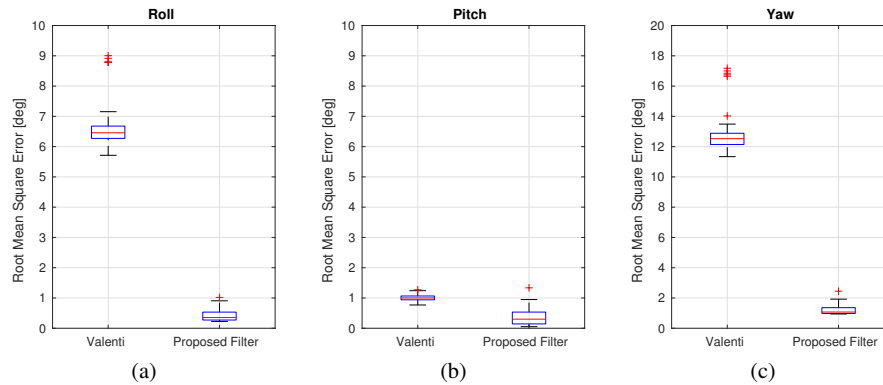




Fig. 10 A Monte-Carlo simulation of 50 runs for the aileron doublet test case.

the noise from the increased norm due to the centripetal acceleration. This is the reason why the Valenti filter performs corrections in Fig. 12. Moreover, the linear interpolation between the thresholds t_1 and t_2 in the Valenti filter enable the slow dynamics to leak through in any case. The gyro model and/or the magnetic field rate is responsible for the correct detection of unsteadiness in this case. The use of the magnetic field rate can be seen as a more robust (Fig. 13) alternative to the almost de facto standard of thresholding the difference in the gyro values and an alternative when the bias cannot be accurately modelled by a Gaussian random walk.




(a) Reference filters




(b) Proposed filter

Fig. 11 A comparison of the attitude estimates for a single stochastic simulation of the filters to a aileron doublet. The Valenti filter incorrectly continues to use the accelerometer vector measurement during times of dynamic motion.



(a) Reference filters



(b) Proposed filter

Fig. 12 A comparison of the roll estimates for a single stochastic simulation of the filters during a coordinated turn. The Valenti filter incorrectly drives the roll angle towards zero.

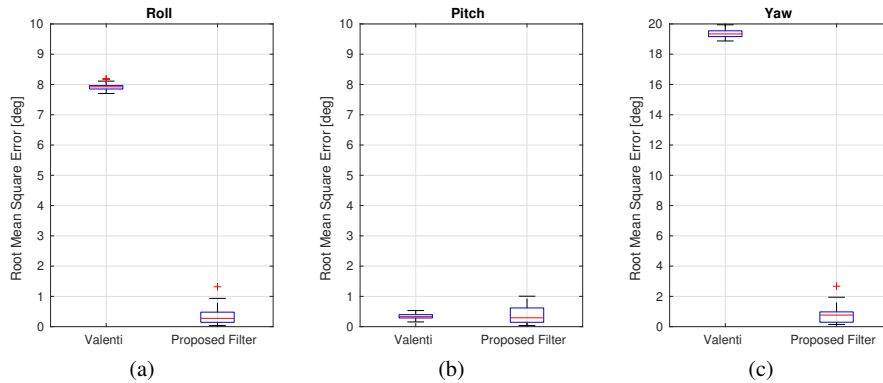


Fig. 13 A Monte-Carlo simulation of 50 runs for the coordinated turn test case.

5 Conclusions & Recommendations

In this paper, a novel robust dynamic-state detection algorithm was developed for use with aircraft state filters using an accelerometer as a vector measurement. The algorithm consisted of a time-dependent check that was made more robust by modelling the evolution of the gyro biases and including the rate of change of the magnetic field vector. The model is able to successfully distinguish slow dynamic manoeuvres from biases and eliminate the problem of pseudo steady-states. The filter is, however, not suitable for extended periods of dynamic motion as the gyro drift will begin to determine the filter performance. The filter was tested on a fixed-wing aircraft but the model can be applied equally to other applications, such as pedestrian navigation, rotary aircraft, and gait analysis.

Evaluation of the filter performance using real sensors to validate it in more realistic scenarios is still necessary. Moreover, magnetic distortion has not been treated at length and is an important concern for electric propulsion vehicles and pedestrian navigation. It may be possible to generate a magnetic distortion model by using the gyros to estimate the distortion vector during the distortion transition period. Furthermore, to limit the impact of extended accelerations on the bias estimation, an approach similar to the quasi-static magnetic field for bias removal present in [1] could be used to determine quasi-static accelerations during rectilinear motion. Additionally, the proposed filter can be compared with the revised filter in [10]. Finally, the angular rate of the aircraft can be obtained from a set of two vector measurements. During a sustained turn, the rate of change of the accelerometer vector measurement is zero and the gravity vector direction can be estimated from the gyro attitude estimate once the aircraft is determined to be in a sustained turn and before significant drift has occurred. This will enable the vehicle to use a non-drifting reference for angular rate and potentially enable the drift of the gyro to be estimated during turns. A mathematical proof-of-concept using least-squares regression to in-

vert the skew-symmetric matrix has been made but practical testing and feasibility with respect to noise levels in measurements should still be evaluated.

Acknowledgements

I would like to thank Incomar Aerospace and Defence Systems for supporting this research.

References

- [1] Afzal MH, Renaudin V, Lachapelle G (2011) Use of earth's magnetic field for mitigating gyroscope errors regardless of magnetic perturbation. *MDPI Sensors* 11(11390-11414), DOI 10.3390/s111211390
- [2] Calusdian J, Yun X, Bachmann E (2011) Adaptive-gain complementary filter of inertial and magnetic data for orientation estimation. In: *IEEE International Conference on Robotics and Automation*, IEEE, pp 1916–1922, DOI 10.1109/ICRA.2011.5979957
- [3] Crassidis J, Markley F, Cheng Y (2007) Survey of nonlinear attitude estimation methods. *Journal of Guidance, Control, and Dynamics* 30(1):12–28, DOI 10.2514/1.22452
- [4] Euston M, Coote P, Mahony R, Kim J, Hamel T (2008) A complementary filter for attitude estimation of a fixed-wing uav. In: *International Conference on Intelligent Robots and Systems*, IEEE/RSJ
- [5] Higgins WT (1975) A comparison of complementary and kalman filtering. *IEEE Transactions on Aerospace and Electronic Systems* 11(3):321–325, DOI 10.1109/TAES.1975.308081
- [6] IEEE (1997) Ieee standard specification format guide and test procedure for single-axis interferometric fiber optic gyros. Online, DOI 10.1109/IEEESTD.1998.86153
- [7] Lee SC, Hong SK (2016) Velocity-aided attitude estimation for helicopter aircraft using microelectromechanical system inertial-measurement units. *MDPI Sensors* 16(2102), DOI 10.3390/s16122102
- [8] Lefferts E, Markley F, Shuster M (1982) Kalman filtering for spacecraft attitude estimation. *Journal of Guidance, Control, and Dynamics* 5(5):417–429, DOI 10.2514/3.56190
- [9] Lima RR, Tôrres LA (2012) Performance evaluation of attitude estimation algorithms in the design of an ahrs for fixed wing uavs. In: *2012 Brazillian Robotics Symposium and Latin American Robotics Symposium*, IEEE, Fortaleza, Brazil, pp 255–260, DOI 10.1109/SBR-LARS.2012.61
- [10] Madgwick SO (2014) Ahrs algorithms and calibration solutions to facilitate new applications using low-cost mems. PhD thesis, University of Bristol

- [11] Madyastha VK, Ravindra VC, Mallikarjunan S, Goyal A (2011) Extended kalman filter vs. error state kalman filter for aircraft attitude estimation. In: AIAA Guidance, Navigation, and Control Conference, AIAA, DOI 10.2514/6.2011-6615
- [12] Mahony R, Hamel T, Pfimlin JM (2008) Nonlinear complementary filters on the special orthogonal group. *IEEE Transactions on Automatic Control* 53(5):1203–1218, DOI 10.1109/TAC.2008.923738
- [13] Markley F (2004) Multiplicative vs additive filtering for spacecraft attitude determination. *Dynamics and Control of Systems and Structures in Space*
- [14] Markley F, Crassidis J (2014) *Fundamentals of Spacecraft Attitude Determination and Control*. Springer and Microcosm Press, DOI 10.1007/978-1-4939-0802-8
- [15] Mulder J, Chu Q, Sridhar J, Breeman J, Laban M (1999) Non-linear aircraft flight path reconstruction review and new advances. *Progress in Aerospace Sciences* 35:673–726
- [16] NASA/NAIF (2003) Quaternions white paper. Tech. rep., NASA/NAIF
- [17] Pittelkau ME (2003) Rotation vector in attitude estimation. *Journal of Guidance, Control, and Dynamics* 26(6):855–860, DOI 10.2514/2.6929
- [18] Qin F, Chang L, Jiang S, Zha F (2018) A sequential multiplicative extended kalman filter for attitude estimation using vector observations. *MDPI Sensors* 18(5), DOI 10.3390/s18051414
- [19] Shuster MD (1993) A survey of attitude representations. *Astronautical Sciences* 41(4):439–517
- [20] Solà J (2017) Quaternion kinematics for the error-state kalman filter. *CoRR*
- [21] Suh YS (2010) Orientation estimation using a quaternion-based indirect kalman filter with adaptive estimation of external acceleration. *IEEE Transactions on Instrumentation and Measurement* 59(12):3296–3305, DOI 10.1109/TIM.2010.2047157
- [22] Teague H (2016) Comparison of attitude estimation techniques for low-cost unmanned aerial vehicles. *arXiv.org*
- [23] Tian Y, Wei H, Tan J (2013) An adaptive-gain complementary filter for real-time human motion tracking with marg sensors in free-living environments. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 21(2)
- [24] Trawny N, Roumeliotis SI (2005) Indirect kalman filter for 3d attitude estimation. Tech. Rep. 2005-002, University of Minnesota
- [25] Valenti RG, Dryanovski I, Xiao J (2015) Keeping a good attitude: A quaternion-based orientation filter for imus and margs. *MDPI Sensors* DOI 10.3390/s150819302
- [26] Xing L, Hang Y, Xiong Z, Liu J, Wan Z (2016) Accurate attitude estimation using ars under conditions of vehicle movement based on disturbance acceleration adaptive estimation and correction. *MDPI Sensors* DOI 10.3390/s16101716
- [27] Yoo TS, Hong SK, Yoon HM, Park S (2011) Gain-scheduled complementary filter design for a mems based attitude and heading reference system. *MDPI Sensors* 11(3816-3830), DOI 10.3390/s110403816