

Batch Misalignment Calibration of Multiple Three-Axis Sensors

Gabriel Hugh Elkaim¹

Author is with the Autonomous Systems Lab at the University of California in Santa Cruz, Santa Cruz, CA 95060, USA. `elkaim at soe.ucsc.edu`

Abstract. When using multiple three-axis sensors, misalignments between the sensors can corrupt the measurements even when each individual sensor is well calibrated. An algorithm was developed to determine the misalignment between any two three-axis sensors whose unit vectors are known in the inertial frame and can be measured by the sensor in body coordinates. This is a batch algorithm, assuming each sensor has been individually calibrated, and requires no extraneous equipment. The input to the misalignment algorithm is a locus of measurements, the paired body-fixed measurements of both sensors, and the known unit vectors for each in inertial coordinates. The output of the algorithm is the misalignment direction cosine matrix such that the second (slave) sensor can be rotated into a consistent coordinate frame as the first (master) sensor. The algorithm works both with heterogenous and homogenous sensors (e.g., accelerometer and magnetometer or multiple magnetometers). The algorithm was validated using Monte Carlo simulations for both large and small misalignments, with and without realistic sensor noise, and shows excellent convergence properties. The algorithm is demonstrated experimentally on a small UAV sensor package and on a small satellite equipped with three high quality magnetometers. In both cases, the algorithm identifies large misalignments created during installation, as well as residual small misalignments when sensing axes are aligned.

1 Introduction

The problem of estimating attitude from vector measurements is well studied for aircraft and satellites. Often known as Wahba's problem due to her 1965 problem in the SIAM Journal [1], there have been myriad solutions proposed and implemented [2]. While these methods all work very well in practice, they assume that the sensors are in fact perfectly aligned and calibrated. That is, the assumption is that the two independent sensors are rigidly coupled and have collinear sensing axes. Furthermore, all null shifts, scale factors, and cross coupling on the axes are also assumed to be fully calibrated.

The literature is rich with calibration solutions for determining the null shifts, scale factor errors, and cross coupling [3][4][5]. However, there are far fewer works specifically about calibrating the misalignment of one three axis sensor

to another so that the sensor frames are coincident. [6] solves a misalignment problem for multiple IMUs using a manufactured platonic solid; [7] (and its related works) proposes three different related methods for “harmonizing” the sensor axes of homogenous sensors; [8] addresses a similar problem with Rotation Averaging for cameras in the Computer Vision domain. It can be assumed that the sensor suite is *never* perfectly aligned with the body axes. This is also true of robotics where accelerometers mounted at the joints are misaligned due to inherent manufacturing variations.

There are many applications where three axis sensors are used for attitude measurements (e.g., aircraft, satellites, UAVs, and underwater vehicles). For example, small UAVs will often use a combination of accelerometers and magnetometers to provide aiding which is then used to estimate biases on the gyros. Previous work on calibrating the three axis sensors for inherent biases and scale factor errors based on a two-step solution [3, 4] or an iterated least squares solution [9] have proven effective. Note that while these solutions calibrate the individual sensor triads, *they provide no information about relative sensing axes.*

When considering the *ensemble averaging* of sensor measurements (using multiple of the same sensors to improve the signal to noise), there is an inherent assumption of a common coordinate frame for each of the sensors. If the sensor axes are misaligned, then the averaging of each axis will degrade the signal to noise ratio (SNR). When using multiple sensors to measure the *gradient* of the field as an additional measurement, again each sensor must be in a consistent coherent axis system or the results will be off [10].

The contribution of this new algorithms is to extract the misalignment between two sensors without any need for an external reference. This greatly simplifies the calibration procedures, and can easily be implemented in the field.

Note that this same algorithm can be applied to align heterogenous sensors (e.g., magnetometer and accelerometer), to align multiples of the same sensor (for ensemble averaging or gradient measurement), or to align the sensor suite to the vehicle body frame. The paper proceeds as follows: Sec. 2 will develop the theory of the new algorithm and its application, Sec. 3 will explore numerical simulations of the algorithm showing convergence metrics and noise sensitivity, Sec. 4 will describe two experiments that were used to test the algorithm on real world data and the results of our new algorithm, and Sec. 5 presents conclusions and future directions.

2 Theory

Of the many solutions to Wahba’s Problem, we will use the formulation by Markley in [11] which uses the singular value decomposition (*SVD*) to solve for the direction cosine matrix (*DCM*). For completeness, that solution is presented here:

$$\min_{\mathcal{R}} J(\mathcal{R}) = \frac{1}{2} \sum_{i=1}^n a_i \|\mathbf{w}_i - \mathcal{R}\mathbf{v}_i\|^2 \quad (1)$$

where \mathbf{w}_i are a set of vectors in the inertial frame, \mathbf{v}_i are the corresponding set of vectors in the body frame, and \mathcal{R} is the *DCM* that transforms a vector from the body to the inertial frame. The a_i 's are an optional set of weights, which for normalization purposes, $\sum a_i = 1$.

From the original solution to Wahba's Problem in [12], we have the formulation that $J(\mathcal{R})$ can be expressed as:

$$J(\mathcal{R}) = \frac{1}{2} \sum_{i=1}^n a_i \|\mathbf{w}_i - \mathcal{R}\mathbf{v}_i\|^2 \quad (2)$$

$$= \frac{1}{2} \text{tr}(\mathbf{W} - \mathcal{R}\mathbf{V})^T (\mathbf{W} - \mathcal{R}\mathbf{V}) \quad (3)$$

$$= 1 - \text{tr}(\mathbf{W}\mathcal{R}\mathbf{V}^T) \quad (4)$$

$$J(\mathcal{R}) = 1 - \sum_{i=1}^n a_i \mathbf{w}_i^T \mathcal{R}\mathbf{v}_i \quad (5)$$

Markley's solution [11] is based on the *SVD* decomposition, and proceeds as follows:

$$\mathcal{B} = \frac{1}{2} \sum_{i=1}^n a_i \mathbf{w}_i \mathbf{v}_i^T \quad (6)$$

$$\text{svd}(\mathcal{B}) = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T \quad (7)$$

$$\mathcal{R} = \mathbf{U}\mathbf{M}\mathbf{V}^T \quad (8)$$

where $\mathbf{M} = \text{diag}([1 \ 1 \ \det(\mathbf{U}) \ \det(\mathbf{V})])$ and is used to enforce that \mathcal{R} is a rotation matrix for a right handed coordinate frame. That is:

$$\mathcal{R}_{opt} = \mathbf{U} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \det(\mathbf{U}) \det(\mathbf{V}) \end{bmatrix} \mathbf{V}^T \quad (9)$$

From this solution, the optimal rotation matrix, \mathcal{R} , can be extracted for any set of non-collinear vectors. Markley's solution will be used extensively throughout the development of the algorithm.

Before continuing with the solution, it is useful to visualize the problem. Consider two unit vectors that are known in inertial coordinates, defined as $\hat{\mathbf{m}}$ and $\hat{\mathbf{s}}$ for *master* and *slave* respectively. The notation we will use shows the coordinate frame to the right of the vector such that ${}^I\hat{\mathbf{m}}$ is the inertial master unit vector, expressed in the inertial frame, and ${}^B\hat{\mathbf{m}}$ is the same unit vector expressed in the body frame.

The key observation is that while the individual unit vectors may be pointed anywhere (e.g., the body frame is not aligned to the inertial frame), there is a constant solid angle between the two (master and slave) unit vectors *in the body frame*. When rotating the body through a set of orientations, the error in that solid angle will project onto each axis of the basis set and make the misalignment matrix observable. That is, a rich set of paired measurements in the body frame is sufficient to reconstruct the misalignment matrix between master and slave sensors.

2.1 Misalignment Calculation

Consider a rigid body with two three axis sensors (one master, \mathbf{m} , and one slave, \mathbf{s}) which measure the three components of their respective signals. That is, for any arbitrary rotation of the rigid body, \mathcal{R} , the sensors will measure:

$${}^B\hat{\mathbf{m}}_i = \mathcal{R}_i^T [{}^I\hat{\mathbf{m}}] + \nu_m \quad (10)$$

$${}^B\hat{\mathbf{s}}_i = \mathcal{R}_i^T [{}^I\hat{\mathbf{s}}] + \nu_s \quad (11)$$

where \mathcal{R}_i is the rotation matrix for that specific measurement, and is not known, and ν is sensor noise. Note that in this case, Wahba's problem can be used to find \mathcal{R}_i , but that is only because the sensors both share a common coordinate frame. If, however, the master and slave do *not* have the same coordinate frames, then there will be a misalignment matrix, \mathcal{R}_{mis} that effectively rotates measurements in the *body-slave* frame to the *body-master*. That is:

$${}^{BM}\hat{\mathbf{s}}_i = \mathcal{R}_{mis} {}^{BS}\hat{\mathbf{s}}_i \quad (12)$$

Where BM and BS refer to body-master and body-slave coordinates. The misalignment matrix \mathcal{R}_{mis} is *constant in the body frame*, then Eq. 10 becomes:

$${}^B\hat{\mathbf{m}}_i = \mathcal{R}_i^T [{}^I\hat{\mathbf{m}}] + \nu_m \quad (13)$$

$${}^B\hat{\mathbf{s}}_i = \mathcal{R}_{mis} \mathcal{R}_i^T [{}^I\hat{\mathbf{s}}] + \nu_s \quad (14)$$

This constant, unknown misalignment between the coordinate frame of the master and the slave must be estimated before using the sensors (either for ensemble average or for attitude estimation). Data from each sensor is collected at different attitudes of the body to generate the paired measurements which can be collected into matrix form:

$${}^B\mathbf{M} = [{}^B\hat{\mathbf{m}}_1 \quad {}^B\hat{\mathbf{m}}_2 \quad \dots \quad {}^B\hat{\mathbf{m}}_n] \quad (15)$$

$${}^B\mathbf{S} = [{}^B\hat{\mathbf{s}}_1 \quad {}^B\hat{\mathbf{s}}_2 \quad \dots \quad {}^B\hat{\mathbf{s}}_n] \quad (16)$$

The algorithm for estimating \mathcal{R}_{mis} uses an iterative approach through all of the paired measurements. $\hat{\mathcal{R}}_{mis}$ is initialized to the identity matrix (a valid DCM). We rotate the body fixed measurements of the slave by the estimate of the misalignment matrix:

$$\hat{{}^B\mathbf{S}} = \hat{\mathcal{R}}_{mis} {}^B\mathbf{S} = \hat{\mathcal{R}}_{mis} [{}^B\hat{\mathbf{s}}_1 \quad {}^B\hat{\mathbf{s}}_2 \quad \dots \quad {}^B\hat{\mathbf{s}}_n] \quad (17)$$

Eq. 17 is used to rotate all slave measurements by the estimate of \mathcal{R}_{mis} . Each measurement pair is then plugged into a Wahba's problem to solve for the individual rotation matrix \mathcal{R}_i for each measurement (see Eq. 9). Thus, for each measurement i we compute \mathcal{R}_i from $\hat{{}^B\mathbf{S}}_i$ and ${}^B\hat{\mathbf{m}}_i$. With this \mathcal{R}_i estimate, we

calculate a new ${}^B\hat{\mathbf{s}}$ by rotating the slave unit vector ${}^I\hat{\mathbf{s}}$ into the body frame. That is:

$${}^B\mathbf{s}_i = \mathcal{R}_i^T [{}^I\hat{\mathbf{s}}] \quad (18)$$

which are aggregated into a larger $3 \times n$ matrix, ${}^B\mathbf{S}$. We then again use the solution to Wahba's problem, this time with ${}^B\mathbf{S}$ and ${}^B\hat{\mathbf{S}}$ with the resulting solution being a new estimate of the misalignment matrix, $\widehat{\mathcal{R}}_{mis}$. This is repeated until the misalignment matrix converges.

It is easy to see that when $\widehat{\mathcal{R}}_{mis}$ is equal to \mathcal{R}_{mis} , then by Eq. 17, Eq. 18, and Eq. 14 that ${}^B\mathbf{S}$ will equal ${}^B\hat{\mathbf{S}}$, because $\widehat{\mathcal{R}}_{mis}^T \mathcal{R}_{mis} = I_{3 \times 3}$. We check the convergence of $\widehat{\mathcal{R}}_{mis}$ by checking the Frobenius norm of the difference between subsequent steps and stop when it is below some tolerance (generally set at 10^{-15}).

Indeed, the algorithm can be shown to always converge to a local minima by expanding Eq. 1 with the definitions of the body measurements, Eq. 14, ignoring the noise terms and rearranging:

$$J(\mathcal{R}, \mathcal{R}_{mis}) = \frac{1}{2} \sum_{i=1}^n \left\| {}^I\mathbf{s} [\mathbf{I} - \mathcal{R}_i \mathcal{R}_{mis} \mathcal{R}_i^T] \right\|^2 \quad (19)$$

The cost function is everywhere differentiable, thus the algorithm breaks down into a block-coordinate descent approach which alternately optimizes for the \mathcal{R}_i 's with \mathcal{R}_{mis} held fixed (which decouples into n Wahba's problems), and then for \mathcal{R}_{mis} with the \mathcal{R}_i 's fixed (which is a simple Procrustes problem on the vectors \mathbf{s}). Since each sub-problem can be optimally solved, the block coordinate descent always converges to a local minima.

Algorithm 1 Compute \mathcal{R}_{mis}

- 1: $\widehat{\mathcal{R}}_{mis} \leftarrow I_{3 \times 3}$
 - 2: **if** $\widehat{\mathcal{R}}_{mis}$ not converged **then**
 - 3: ${}^B\hat{\mathbf{S}} \leftarrow \widehat{\mathcal{R}}_{mis} [{}^B\hat{\mathbf{s}}_1 \ {}^B\hat{\mathbf{s}}_2 \ \dots \ {}^B\hat{\mathbf{s}}_n]$
 - 4: **for all** measurements i in the body frame **do**
 - 5: Solve Wahba's Problem using $\mathbf{v} = [{}^B\hat{\mathbf{m}}_i \ {}^B\hat{\mathbf{S}}_i]$ and $\mathbf{w} = [{}^i\hat{\mathbf{m}} \ {}^i\hat{\mathbf{s}}]$ for \mathcal{R}_i
 - 6: ${}^B\mathbf{s}_i \leftarrow \mathcal{R}_i^T [{}^i\hat{\mathbf{s}}]$
 - 7: Collect all ${}^B\mathbf{s}_i$ into ${}^B\mathbf{S}$
 - 8: Solve Wahba's Problem using $\mathbf{V} = {}^B\hat{\mathbf{S}}$ and $\mathbf{W} = {}^B\mathbf{S}$ for $\widehat{\mathcal{R}}_{mis}$
 - 9: **if** $\left\| \widehat{\mathcal{R}}_{mis} - \widehat{\mathcal{R}}_{mis}^{prev} \right\|_{Frobenius} \leq tol$ **then**
 - 10: done
 - 11: **else**
 - 12: goto step [3]
 - 13: $\mathcal{R}_{mis} \leftarrow \widehat{\mathcal{R}}_{mis}$
-

Note that the algorithm still works when $\hat{\mathbf{m}}$ and $\hat{\mathbf{s}}$ are the same (e.g., homogeneous sensors). This is due to the SVD nature of Markley's solution to Wahba's

Problem, which returns the minimal answer (Eq. 9) even when $\text{Rank } \mathcal{B} = 2$; as long as there are sufficiently diverse measurements observe the misalignment, the estimate for \mathcal{R}_{mis} will converge to a local minima.

3 Simulation

In order to determine the convergence and stability of the algorithm, as well quantify the error in estimates of \mathcal{R}_{mis} , a set of Monte Carlo numerical simulations were performed.

Two unit vectors, $\hat{\mathbf{m}}$ and $\hat{\mathbf{s}}$, and \mathcal{R}_{mis} were chosen at random, using $\mathcal{R}_{mis} = e^{[\omega \times]}$ where ω is a randomly chosen rotation $[3 \times 1]$ and $[\cdot \times]$ is defined as the skew symmetric matrix.

An additional number of random rigid body rotations, \mathcal{R}_i (varied from 3 to 50) were used to generate ${}^B\mathbf{M}$ and ${}^B\mathbf{S}$. The simulations were run with both with and without noise added to the body measurements. Typical noise variance on each measurement was set to 0.01 (though anecdotal observations show the algorithm is quite tolerant of larger noise variance in simulation).

For observability of the $\hat{\mathcal{R}}_{mis}$ a minimum of three non-coplanar body fixed measurements are required. However, with only three measurements, the likelihood of converging to the correct $\hat{\mathcal{R}}_{mis}$ is poor. More measurements increase the probability of converging to the correct answer.

Fig. 1 shows a typical run for a scenario with $n = 4$ unique \mathcal{R}_i 's. The two true inertial unit vectors (orange), master (black) and true slave (blue) unit vector triads are shown in the body coordinates. The misaligned slave (green dashed) are shown as well (green dots show the misalignment iterations for clarity). The solution for the estimated slave (blue dots) are shown at each iteration. In this case, $\hat{\mathcal{R}}_{mis}$ does indeed converge to the true value. As another visualization, the axis of $\hat{\mathcal{R}}_{mis}$ (cyan, in an axis and angle representation) demonstrates the estimate converging to the correct one (red). In this case, the difference between the true and estimated misalignment matrix (Frobenius norm) was 4.29×10^{-15} .

The true misalignment matrix had an axis of $[-0.3554 \ -0.9850 \ -0.2695]$ and an angle of 51.11° . Which is to note that this is *not* a particularly small misalignment. This was a noise free case, and converged to the true value in 208 iterations.

Fig. 2 shows another simulation with $n = 12$ measurements; again the unit vectors (in orange), along with the body measured vectors (in black with triangles for the master and blue for the slave). The misaligned measurements are shown in green dashed line with triangles with the dots showing the evolution of the $\hat{\mathcal{R}}_{mis}$. Again, here the cyan vectors show the axis of the estimated misalignment and the blue dots show the convergence of the estimate. In this simulation, noise has been added to *all* body measurements. The estimate converges to within 10^{-2} of the exact misalignment, even in the presence of noise (this corresponds to an angular error of less than $\frac{1}{2}^\circ$).

The tolerance was set to 1×10^{-15} . Note that due to the use of the Markley SVD solution, the covariance of the misalignment matrix is also available, and

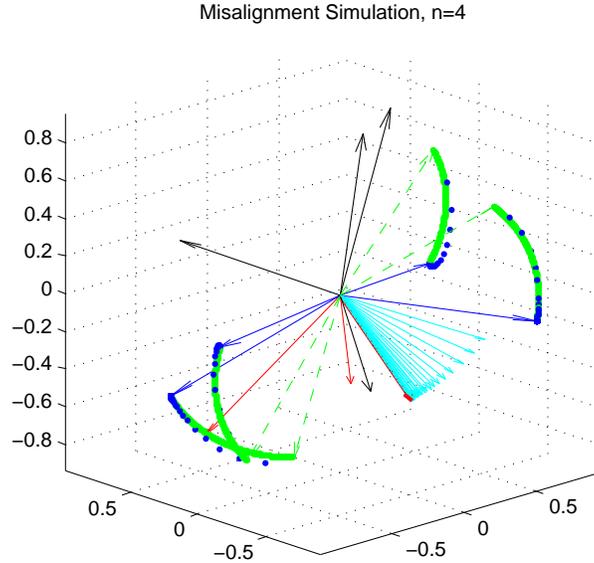


Fig. 1: Simulation of the misalignment problem with $n = 4$ measurements, showing the master (black) and true slave (blue) unit vector triads. The misaligned slave unit vector triads (green dashed), and the iterations that walk in the solutions to \mathcal{R}_{mis} , showing convergence to the true axes. The inertial sensor vectors are shown in orange, and the axis of \mathcal{R}_{mis} is shown in red, with the steps shown in cyan.

shows in simulation that the converged misalignment is well within the noise covariance of true even when using noisy data. See [11] for details of the covariance matrix.

Using Monte Carlo simulations with 10000 runs for each of $n = [3, 4, 5, 10, 20, 50]$ demonstrates that the more data available in terms of measurements, the more likely convergence will occur. There are two available metrics to quantify convergence: (i) $\left\| \mathcal{R}_{mis} - \widehat{\mathcal{R}}_{mis} \right\|_F$ and (ii) the dot product of the true axis of rotation and the estimated one: $1 - \mathbf{v}^T \hat{\mathbf{v}}$. Since both of these measures are very small when the algorithm converges, the log of both is taken to quantify convergence; anything less than -6 is for all intents and purposes converged (corresponding to angular errors of less than $\frac{1}{100}^\circ$). Fig. 3 shows the cumulative probability distribution from the Monte Carlo runs for various measurement numbers (in the noise free case). The probability of converging on the correct misalignment is greater than 98% when $n = 50$. The small number of iterations that failed to converge were due to the measurement rotations \mathcal{R}_i being insufficiently diverse to allow reconstruction of \mathcal{R}_i from Whaba's problem.

Misalignment Simulation, n=12, noise added

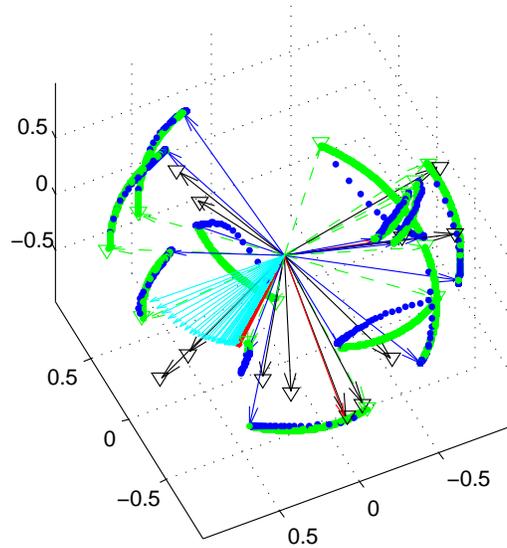


Fig. 2: Simulation of the misalignment problem with noise and $n = 12$, showing that the misaligned slave vectors (green dashed) converge to their true values (blue) even in the presence of noise on all measurements.

It is noted from the simulation results that the estimate of the misalignment matrix can become trapped in a local minima. If the algorithm is restarted with a different initial condition ($\mathcal{R}_0 \neq I_{3 \times 3}$) then it will often converge to the correct solution. This indicated how to validate the global nature of the estimate in practice. Segregate the data into two or more segments of at least 20 points each. Run the algorithm on each from at least two separate \mathcal{R}_0 . If all converge to the same estimate, then the estimate has converged to the true value with high confidence.

Fig. 4 shows the CDF for several variants of the $n = 20$ Monte Carlo simulation (note that this is a relatively small number of individual measurements). The nominal noise free case (blue) has a 96.5% chance of success. If a simple retry of resetting the initial guess on \mathcal{R}_0 is allowed, then the probability goes to 99.2% (red). In the case of noise added to the measurements, it simply moves the lowest Frobenius norm, but does not actually affect convergence (black). This is because the noise is uniformly distributed and thus averages out through the Wahba's problem solution. When the misalignment matrix is constrained to be small ($< 6^\circ$ on any axis), then the convergence is 100% for both the noise free and noisy cases (green and cyan respectively). Thus, for small misalignments, this algorithm always finds the true solution.

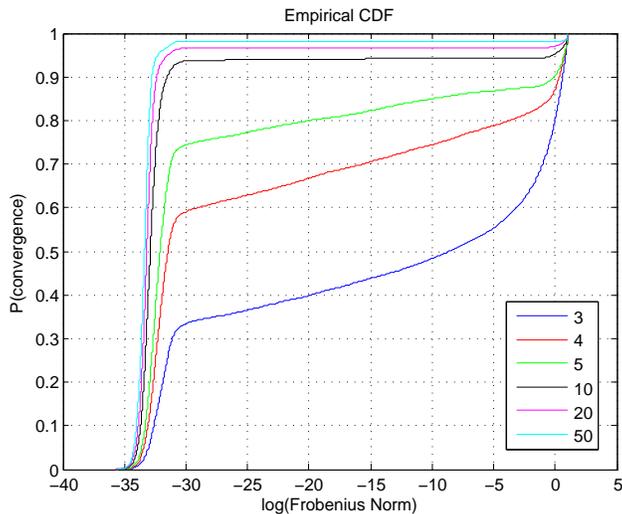


Fig. 3: Cumulative probability distribution of convergence for various measurements (noise free), showing the error between true \mathcal{R}_{mis} and the estimate.

4 Experimental Results

In order to demonstrate applications of the algorithm, two different experiments were performed. The first was from a small satellite experiment where three high quality magnetometers are used to take an ensemble average to generate better data for a *space weather* application. The second was from a small UAV autopilot (see Fig. 6) from its accelerometer and magnetometer sensors, where the data comes from a tumble test used to calibrate the sensors for scale factor, biases, and non-orthogonality [13].

4.1 SmallSat Experiment

For the SmallSat experiment, data was provided from three high quality flux-gate magnetometers mounted on a rigid platform. Rather than rotating the rigid platform, it was inserted into a Helmholtz coil which could manipulate the magnetic field simulating body rotations in a very controlled manner. Three experiments were run, with data collected at each instance. Before any calibration was performed, the magnitude of the total magnetic field was calculated for each of the three sensors showing a mean of $\sim 900nT$ and a standard deviation of $\sim 5-10nT$.

The first calibration pass was to correct the magnetometers for bias, scale factor, and non-orthogonality errors using our previous *two step* techniques outlined in [3] and [4]. This forced the mean to exactly $900nT$ and resulted in a standard deviation of $\sim 1.35nT$.

When plotting the post calibration data individually (see Fig. 5a), it can be easily seen that this is a general lack of agreement between some of the

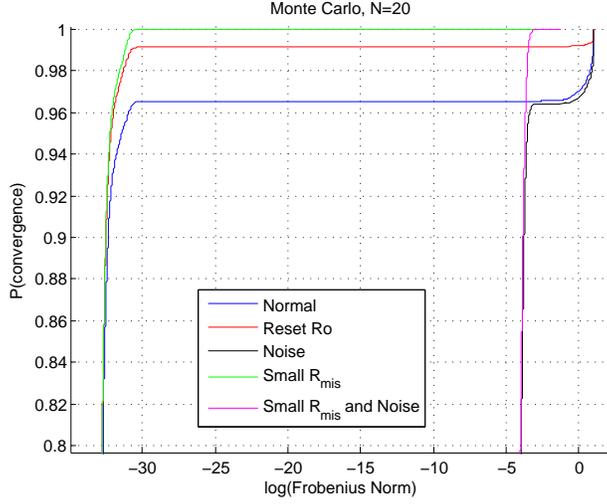


Fig. 4: Cumulative probability distribution for $n = 20$

axes of the three magnetometers (this was not known until the algorithm was run to estimate \mathcal{R}_{mis}). That is, there is a large rotation between one of the magnetometers and the other two due to the installation of the magnetometers on the platform.

Using Magnetometer “A” as the master, the misalignment matrix between A and B is computed as:

$$\mathcal{R}_{mis}^{A \rightarrow B} = \begin{bmatrix} -0.00287 & 0.99998 & 0.00496 \\ -0.99993 & -0.00292 & 0.0119 \\ 0.0119 & -0.00492 & 0.99992 \end{bmatrix}$$

and the misalignment between A and C is computed as:

$$\mathcal{R}_{mis}^{A \rightarrow C} = \begin{bmatrix} -0.99992 & -0.01257 & -0.00214 \\ -0.01256 & 0.9999 & -0.0071 \\ 0.00223 & -0.0071 & -0.99997 \end{bmatrix}$$

Note that while $\mathcal{R}_{mis}^{A \rightarrow C}$ is somewhat close to a 180° rotation (x and z axes reversed), the one from $A \rightarrow B$ is not, indicating that it is mounted differently (the third magnetometer is mounted underneath the test platform). Finally, magnetometers B and C are rotated into the coordinate frame of A, and an ensemble average is taken. For the ensemble measurement, the norm of the magnetic field has a mean of $900nT$ and a standard deviation of $1.31nT$. While this is only a small improvement, it nevertheless caught the large misalignment of magnetometer B which was mounted on its side. Fig. 5b shows the aligned data, and shows a much better match on the data than the pre-alignment data.

The conclusion from the small improvement from the ensemble average (theoretically the signal to noise should be improved by $\sqrt{3}$) is that the noise is not independent, but rather from the Helmholtz coil itself. This means that the magnetometers are actually capable of better performance than is indicated by this experiment. This is a result that would not have been possible to determine without the misalignment correction.

4.2 UAV Tumble Experiment

The second experiment is based on data from the SLUGS autopilot developed at UCSC for UAV research into guidance, navigation, and control (GNC). The SLUGS has a MEMS-based Analog Devices three axis accelerometer, and a Honeywell three axis magnetometer on its circuit board (Fig. 6). While every care has been made to align the axes during manufacture, assembly, and mounting to the aircraft, this alignment cannot be trusted to be exact. Furthermore, the manufacturers of the sensors disagree on coordinate frames for sensing axes resulting in ambiguity of the polarity of the measurements. The magnetometer is run through the misalignment calibration process twice: first to correct the coordinate frame ambiguity, and then again with the appropriate axes flipped to do the fine alignment.

In order to calibrate these sensors, the aircraft was tumbled in a field by hand to collect data for processing in our *two step* calibration routines. This same data was then used to run the misalignment algorithm to determine the actual misalignment. Note that the final “realigned” sensor will still be misaligned from the body axes of the aircraft, but both magnetometer and accelerometer will be on a single coherent “sensor” frame.

The data from the tumble test was 25,633 pairs of accelerometer and magnetometer data that were run through the algorithm. The algorithm converged after 241 iterations. The final misalignment *DCM* was:

$$\mathcal{R}_{mis} = \begin{bmatrix} 0.9952 & -0.822 & 0.0537 \\ 0.0816 & 0.9966 & 0.0129 \\ -0.0546 & -0.0085 & 0.9985 \end{bmatrix}$$

From the misalignment matrix, it is observed to have pitch and roll misalignment on the order of $3-4^\circ$. More interesting still is the covariance of the misalignment estimate:

$$\mathcal{P}_{body} = \begin{bmatrix} 1.603 & 0.7256 & -4.576 \\ 0.7256 & 3.013 & -9.802 \\ -4.576 & -9.802 & 60.42 \end{bmatrix}$$

Which shows that the algorithm had the hardest part in determining the $[3, 3]$ term of the matrix. Rerunning the algorithm using decimated data and different initial \mathcal{R}_0 always converges to the same value. This gives confidence that this is the true misalignment matrix for the magnetometer relative to the accelerometer.

5 Conclusions

A batch algorithm for estimating the misalignment between multiple three axis sensors from sensor measurements in the body frame has been developed. This is done via repeatedly applying Wahba's Problem solution to the data and iterating until converged. The algorithm works with either heterogenous or homogenous sensors, and is shown to converge well with sufficient number of points. The misalignment matrix is from one *master* to a *slave* sensor. Monte Carlo simulations show convergence probabilities in both the noise free and noisy data cases. The algorithm was run on two real experiments: a SmallSat (homogenous) and a UAV (heterogenous). In both cases, the algorithm was able to identify both large and small misalignments and quantify them in a mathematically rigorous manner.

References

1. Wahba, G., et al.: Problem 65-1: A least squares estimate of satellite attitude. *Siam Review* **7**(3) (1965) 409
2. Crassidis, J., Markley, F., Cheng, Y.: Survey of nonlinear attitude estimation methods. *Journal of Guidance Control and Dynamics* **30**(1) (2007) 12
3. Elkaim, G.H., Foster, C.: Extension of a non-linear, two-step calibration methodology to include non-orthogonal sensor axes. *IEEE Transactions on Aerospace Electronic Systems* **44** (07/2008 2008)
4. Vasconcelos, J., Elkaim, G.H., Silvestre, C., Oliveira, P., Cardeira, B.: A geometric approach to strapdown magnetometer calibration in sensor frame. *IEEE Transactions on Aerospace Electronic Systems* (2010)
5. Dorveaux, E.: Magneto-inertial navigation: Principles and application to an indoor pedometer. PhD thesis, École Nationale Supérieure des Mines de Paris (2011)
6. Nilsson, J.O., Skog, I., Handel, P.: Aligning the forceseliminating the misalignments in imu arrays. *IEEE Transactions on Instrumentation and Measurement* **63**(10) (2014) 2498–2500
7. Dorveaux, E., Vissiere, D., Petit, N.: On-the-field calibration of an array of sensors. In: *American Control Conference (ACC)*, 2010, IEEE (2010) 6795–6802
8. Hartley, R., Trumpf, J., Dai, Y., Li, H.: Rotation averaging. *International Journal of Computer Vision* **103**(3) (2013)
9. Dorveaux, E., Vissiere, D., Martin, A.P., Petit, N.: Iterative calibration method for inertial and magnetic sensors. In: *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on, IEEE* (2009) 8296–8303
10. Dorveaux, E., Boudot, T., Hillion, M., Petit, N.: Combining inertial measurements and distributed magnetometry for motion estimation. In: *American Control Conference (ACC)*, 2011, IEEE (2011) 4249–4256
11. Markley, F.: Attitude determination using vector observations and the singular value decomposition. *Journal of the Astronautical Sciences* **36**(3) (July-September 1988) 245–258
12. Farrell, J., Stuelplnagel, J., Wessner, R., Velman, J., Brook, J.: A least squares estimate of satellite attitude (grace wahba). *SIAM Review* **8**(3) (1966) 384–386

13. Lizarraga, M., Elkaim, G.H., Curry, R.: Slugs uav: A flexible and versatile hardware/software platform for guidance navigation and control research. In: American Control Conference (ACC), 2013, IEEE (2013) 674–679

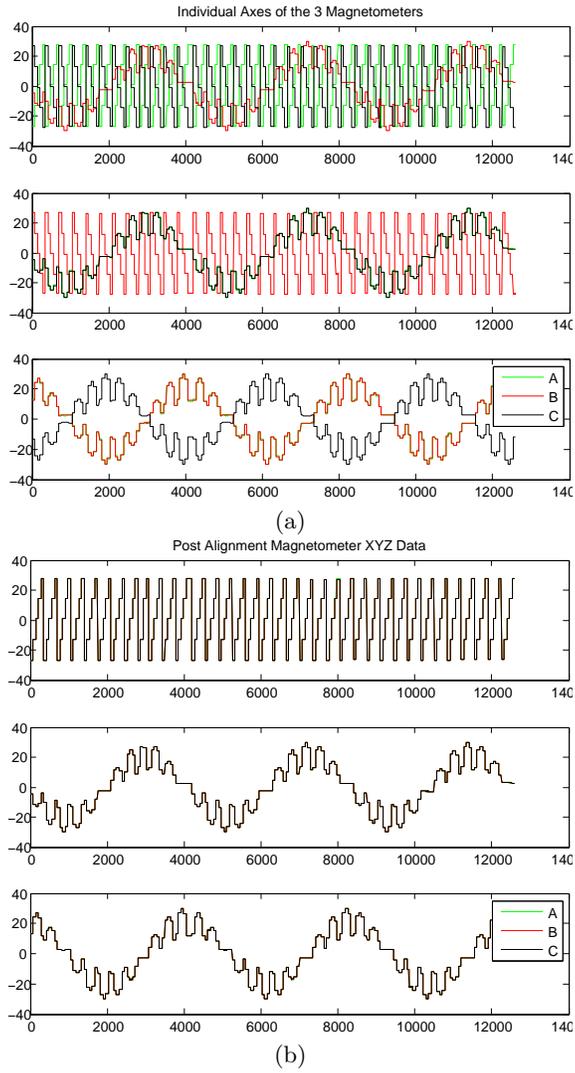


Fig. 5: XYZ data of the three magnetometers: (a) Pre-Alignment and (b) Post-Alignment



Fig. 6: The SLUGS autopilot