Misalignment Calibration Using Body Frame Measurements

Gabriel Hugh Elkaim

Abstract— **Manufacturing flaws and construction realities always impose a misalignment between multiple sensors (even though they remain rigidly coupled together). 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 in body coordinates. 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 matrix ($3\times 3 \in \mathcal{SO}(3)$) such that the second sensor can be rotated into a consistent coordinate frame as the first using the output direction cosine matrix. 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, and shows excellent convergence properties even in the presence of realistic sensor noise. 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.**

I. Introduction

The problem of estimating attitude from vector measurements is well established in both aircraft and satellite usage. 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 and have a rich tradition of real world usage, 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]. However, there is very little in the literature about calibrating the misalignment of one three axis sensor to another so that the sensor frames are coincident. 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 determine 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 has proven effective [3], [4].

When considering the *ensemble averaging* of sensor measurements (using multiple of the same sensors to improve the signal to noise by $\sqrt{n}$), there is an inherent assumption of a common coordinate frame for each of the sensors. If the sensors are misaligned from each other, then the averaging of each axis will degrade the signal to noise worse rather than improve.

Note that this same algorithm can be applied to aligning different sensors (e.g., magnetometer and accelerometer), to aligning multiple of the same sensor (for ensemble averaging), or to align the sensor suite to the vehicle body frame. The paper proceeds as follows: Sec. II will develop the theory of the new algorith and its application, Sec. III will explore Monte Carlo simulations of the algorithm showing convergence metrics and noise sensitivity, Sec. IV will describe two experiments that were used to test the algorithm on real world data and the results of our new algorithm, and Sec. V presents conclusions and future directions.

II. Theory

Of the many solutions to Wahba's Problem, we will use the formulation by Markley in [5] which uses an *SVD* decomposition to solve for the *DCM*. For completeness, that solution is presented here:

$$\min_{\mathcal{R}} J(\mathcal{R}) = \frac{1}{2}\sum_{i=1}^{N} a_i \left\|\vec{w}_i - R\vec{v}_i\right\|^2 \tag{1}$$

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

$$\sum_{i=1}^{N} a_i = 1 \tag{2}$$

From the original solution to Wahba's Problem in [6], we

G. Elkaim is with the Department of Computer Engineering, University of California Santa Cruz, Santa Cruz, CA 95064, USA `elkaim at soe.ucsc.edu`

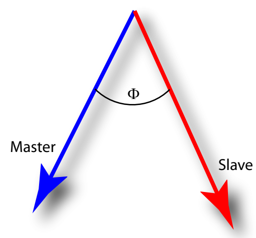

Fig. 1: Unit vectors for two non-collinear non-zero vectors

have the formulation that $J(\mathcal{R})$ can be expressed as:

$$\begin{aligned} J(\mathcal{R}) &= \frac{1}{2}\sum_{i=1}^{N} a_i \left\|\vec{w}_i - R\vec{v}_i\right\|^2 & (3)\\ &= \frac{1}{2}tr(W-\mathcal{R}V)^T(W-\mathcal{R}V) & (4)\\ &= \frac{1}{2}[tr(W^TW)+tr(V^TV)-2trW^T\mathcal{R}V] & (5)\\ &= 1-tr(W\mathcal{R}V^T) & (6)\\ J(\mathcal{R}) &= 1-\sum_{i=1}^{n} a_i {w_i}^T\mathcal{R}v_i & (7)\end{aligned}$$

Markley's solution [5] is based on the *SVD* decomposition, and proceeds as follows:

$$\mathcal{B} = \frac{1}{2}\sum_{i=1}^{N} a_i\vec{w}_i\vec{v}_i^T \quad (8)$$

$$\text{svd}(\mathcal{B}) = U\Sigma V^T \quad (9)$$

$$\mathcal{R} = UMV^T \quad (10)$$

where $M = diag(\left[\begin{array}{ccc} 1 & 1 & \det(U)\det(V)\end{array}\right])$ and is used to enforce that R is a rotation matrix for a right handed coordinate frame. That is:

$$\mathcal{R}_{opt} = U\left[\begin{array}{ccc} 1 & 0 & 0\\ 0 & 1 & 0\\ 0 & 0 & \det(U)\det(V)\end{array}\right]V^T \quad (11)$$

From this solution, we can extract the optimal rotation matrix, $\mathcal{R}$ for any set of non-collinear vectors. We will use this to determine the misalignment matrix between two sensors.

The key observation is that while the unit vectors may be pointed anywhere, there is a constant solid angle between the two (master and slave) unit vectors. When rotating the body through a locus the error in that solid angle will project onto each axis of the basis set and make the misalignment matrix observable.

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.

Given two unit vectors with a solid angle of Φ between them (such as pictured in Fig. 1), the minimal rotation matrix from one to the other is defined as the rotation Φ about the unit vector that is the cross product between the two unit vectors. The well known formula for a direction cosine matrix (*DCM*) from an axis and angle is:

$$\mathcal{R} = \cos\Phi[I]_{3\times 3} + (1-\cos\Phi)\hat{e}\hat{e}^T + \sin\Phi\,[e\times] \quad (12)$$

where $[e\times]$ is the skew symmetric matrix, which also corresponds to the cross product in matrix form:

$$[e\times] = \left[\begin{array}{ccc} 0 & -e_z & e_y\\ e_z & 0 & -e_x\\ -e_y & e_x & 0\end{array}\right] \quad (13)$$

We can use this equation, and the special case of the the fact that both $\hat{\mathbf{m}}$ and $\hat{\mathbf{s}}$ are unit vectors, to derive the minimal rotation matrix from $\mathbf{m} \rightarrow \mathbf{s}$ using some elementary vector identities:

$$\mathcal{R} = m^Ts[I]_{3\times 3} + \frac{(1-m^Ts)}{s^T[m\times]^2 s}[m\times]ss^T[m\times] + [(m\times s)\times] \quad (14)$$

Note that the middle term in Eq. 14 will become undefined if the unit vectors are co-linear, however, then the rotation matrix is simply the identity matrix, $[I]_{3\times 3}$.

A. *Misalignment Calculation*

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

$$\begin{aligned} {}^b\hat{\mathbf{m}}_i &= \mathcal{R}_i[{}^I\hat{\mathbf{m}}] + \nu_m & (15)\\ {}^b\hat{\mathbf{s}}_i &= \mathcal{R}_i[{}^I\hat{\mathbf{s}}] + \nu_s & (16)\end{aligned}$$

where $\mathcal{R}_i$ is the rotation to get the rigid body to that specific orientation, and is not known, and ν is sensor noise. Note that in this case, we can use Wahba's problem to find $\mathcal{R}_i$, but that is only because the sensors both share a common coordinate frame. If we repose the problem with a misalignment matrix $\mathcal{R}_{mis}$ that is *constant in the body frame*, then the equations become:

$$\begin{aligned} {}^b\hat{\mathbf{m}}_i &= \mathcal{R}_i[{}^I\hat{\mathbf{m}}] + \nu_m & (17)\\ {}^b\hat{\mathbf{s}}_i &= \mathcal{R}_{mis}\mathcal{R}_i[{}^I\hat{\mathbf{s}}] + \nu_s & (18)\end{aligned}$$

This constant, unknown misalignment between the coordinate frame of the master and the slave must be estimated before using the sensor (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 aggregated:

$$\begin{aligned} {}^B\mathbf{M} &= \left[\begin{array}{cccc} {}^b\hat{\mathbf{m}}_1 & {}^b\hat{\mathbf{m}}_2 & \cdots & {}^b\hat{\mathbf{m}}_n\end{array}\right] & (19)\\ {}^B\mathbf{S} &= \left[\begin{array}{cccc} {}^b\hat{\mathbf{s}}_1 & {}^b\hat{\mathbf{s}}_2 & \cdots & {}^b\hat{\mathbf{s}}_n\end{array}\right] & (20)\end{aligned}$$

The algorithm for estimating $\mathcal{R}_{mis}$ uses an iterative approach through all of the paired measurements. $\widehat{\mathcal{R}}_{mis}$ 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:

$$\widehat{{}^B\mathbf{S}} = \widehat{\mathcal{R}}_{mis}{}^B\mathbf{S} = \widehat{\mathcal{R}}_{mis}\left[\begin{array}{cccc} {}^b\hat{\mathbf{s}}_1 & {}^b\hat{\mathbf{s}}_2 & \cdots & {}^b\hat{\mathbf{s}}_n \end{array}\right] \tag{21}$$

Eq. 21 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. 1). Thus, for each measurement i we compute $\mathcal{R}_i$ from $\widehat{{}^B\mathbf{S}}_i$ and ${}^B\hat{\mathbf{m}}_i$. With this 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:

$${}^B\mathbf{s}_{Ei} = \mathcal{R}_i{}^T[{}^I\hat{\mathbf{s}}] \tag{22}$$

which are aggregated into a larger $3 \times n$ matrix, ${}^B\mathbf{S}_E$. We then again use the solution to Wahba's problem, this time with ${}^B\mathbf{S}_E$ and $\widehat{{}^B\mathbf{S}}$ with the resulting rotation matrix 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. 21, Eq. 22, and Eq. 18 that ${}^B\mathbf{S}_E$ will equal ${}^B\mathbf{S}$, because $\widehat{\mathcal{R}}_{mis}^T\mathcal{R}_{mis} = I_{3\times3}$. We check the convergence of $\widehat{\mathcal{R}}_{mis}$ by checking the Frobenius norm of the difference between subsequent timesteps and stop when it is below some tolerance (generally set at 10^{-15}).

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

1: $\widehat{\mathcal{R}}_{mis} \leftarrow I_{3\times3}$
2: **while** $\hat{\mathcal{R}}_{mis}$ not converged **do**
3: $\quad\widehat{{}^B\mathbf{S}} \leftarrow \widehat{\mathcal{R}}_{mis}\left[\begin{array}{cccc} {}^b\hat{\mathbf{s}}_1 & {}^b\hat{\mathbf{s}}_2 & \cdots & {}^b\hat{\mathbf{s}}_n \end{array}\right]$
4: $\quad$**for all** measurements i in the body frame **do**
5: $\quad\quad$Solve Wahba's Problem using $\vec{v} = \left[{}^b\hat{\mathbf{m}}_i \quad \widehat{{}^B\mathbf{S}}_i\right]$ and $\vec{w} = \left[{}^i\hat{\mathbf{m}} \quad {}^i\hat{\mathbf{s}}\right]$ for $\mathcal{R}_i$
6: $\quad\quad{}^B\mathbf{S}_{Ei} \leftarrow \mathcal{R}_i{}^T[{}^i\hat{\mathbf{s}}]$
7: $\quad$Aggregate all ${}^B\mathbf{S}_{Ei}$ into ${}^B\mathbf{S}_E$
8: $\quad$Solve Wahba's Problem using $\vec{V} = \widehat{{}^B\mathbf{S}}$ and $\vec{W} = {}^B\mathbf{S}_E$ for $\widehat{\mathcal{R}}_{mis}$
9: $\quad$**if** $\left\|\widehat{\mathcal{R}}_{mis}^k - \widehat{\mathcal{R}}_{mis}^{k-1}\right\|_{Frobenius} \leqslant tol$ **then**
10: $\quad\quad$done
11: $\quad$**else**
12: $\quad\quad$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. This is due to the *SVD* nature of Markley's solution to Wahba's Problem, which returns the minimal answer (Eq. 14) when $\operatorname{Rank}\mathcal{B} = 2$.

III. 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 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 in Eq. 13.

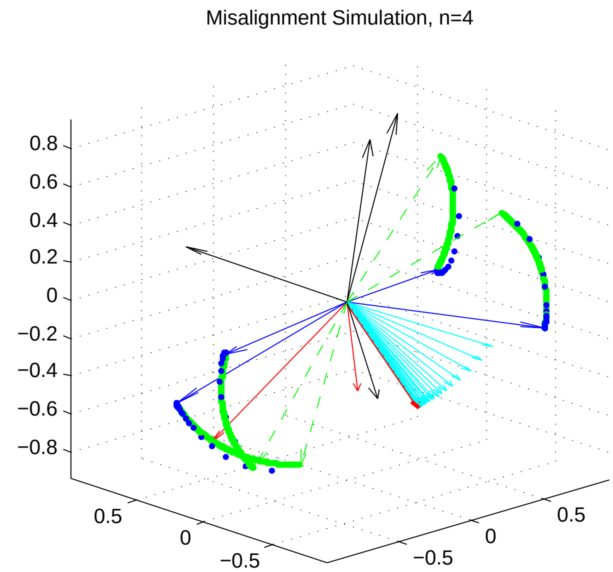

Fig. 2: Simulation of the misalignment problem with $n = 4$ measurements

An additional number of random rigid body rotations (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.

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

Fig. 2 shows a typical run for a scenario with $n = 4$ unique body measurements. The two inertial unit vectors (red), master (black) and true slave (blue) are shown in the body coordinates. The misaligned slave (green dashed) are shown as well (green dots show the misalignment rotation for clarity). The solution for the estimated slave (blue dots) are shown at each iteration. In this case, $\widehat{\mathcal{R}}_{mis}$ does indeed converge to the true value. As another visualization, the axis of $\widehat{\mathcal{R}}_{mis}$ (cyan, in an axis and angle representation) demonstrates the estimate converging to the correct one. 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 $\left[-0.3554 \quad -0.9850 \quad -0.2695\right]$ 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. 3 shows another simulation with $n = 12$ measurements; again the unit vectors (in red), along with the body measured vectors (in black for the master and blue for the slave). The misaligned measurements are shown in green dashed line along with the rotation through the misalignment to show the rotation. Again, here the cyan vectors show the axis of the estimated misalignment and the blue dots show

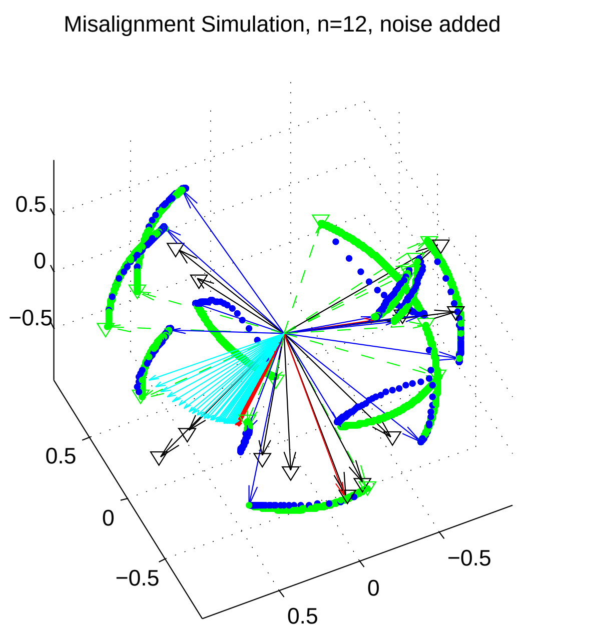

Fig. 3: Simulation of the misalignment problem with noise and $n = 12$.

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.

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 shows in simulation that the converged misalignment is well within the noise covariance of true even when using noisy data.

Using the 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 axes of rotation, $1 - \mathbf{v}^T \hat{\mathbf{v}}$. Since both of these measures are very small when the algorithm converges, the $\log$ of both is taken to quantify convergence; anything less than -2 is for all intents and purposes converged. Fig. 4 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 $n = 50$.

What can be noted by studying the simulation is that the estimate of the misalignment matrix becomes trapped in a local minima. If the algorithm is restarted with a different initial condition ($\mathcal{R}_0 \neq I_{3\times3}$) then it will often converge to the correct solution. This gives an insight on how to validate ones 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

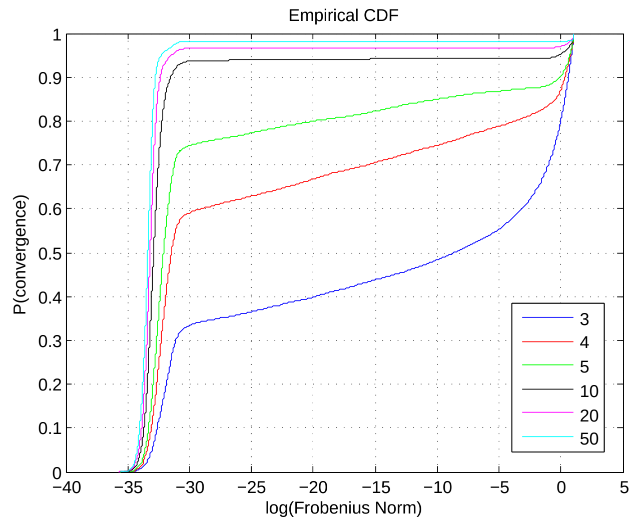

Fig. 4: Cumulative probability distribution of convergence for various measurements (noise free)

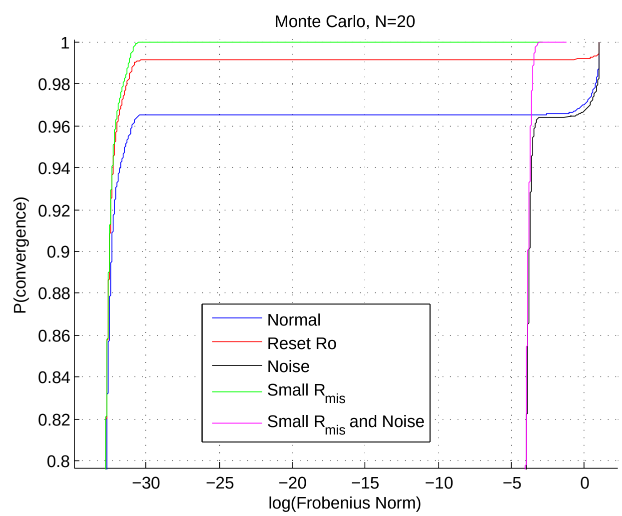

Fig. 5: Cumulative probability distribution for $N = 20$

value with high confidence.

Fig. 5 shows the CDF for several variants of the $N = 20$ Monte Carlo simulation. 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 ($< 6°$ on any axis), then the convergence in 100% for both the noise free and noisy cases (green and cyan respectively).

IV. Experimental Results

In order to validate the algorithm, two different experiments were performed. The first was from a small satellite experiment where three high quality magnetometers are used

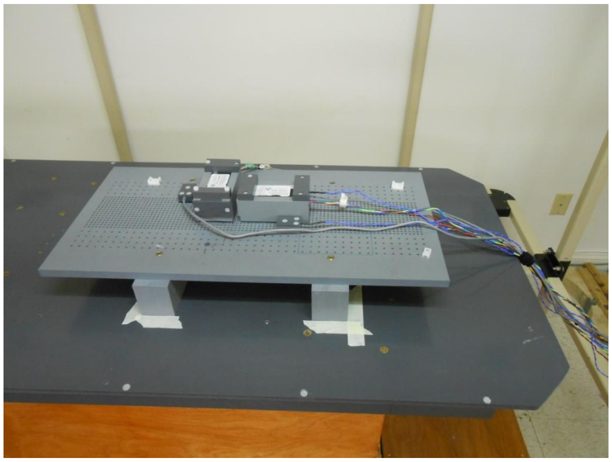

Fig. 6: Three Magnetometers in Helmholz coil

to take an ensemble average to generate better data for a *space weather* application. The second was from a small UAV autopilot[7][8] 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[9].

A. SmallSat Experiment

For the SmallSat experiment, data was provided from three high quality flux-gate magnetometers mounted on a rigid platform pictured in Fig. 6. Rather than rotating the rigid platform, it was inserted into a Helmholz 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 ~$900nT$ and a 1-σ of ~5-10nT.

The first calibration pass was to correct the magnetometers for hard and soft iron biases using our previous *two step* techniques outlined in [3] and [4]. This forced the mean to exactly $900nT$ and resulted in a 1-σ of ~$1.35nT$.

When plotting the post calibration data individually (see Fig. 7a), it can be easily seen that this is a general lack of agreement between some of the axes of the three accelerometers (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.

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

$$\overset{A\to B}{\mathcal{R}_{mis}} = \begin{bmatrix} -0.00286586 & 0.999984 & 0.00495841 \\ -0.999925 & -0.00292483 & 0.0119263 \\ 0.0119406 & -0.00492386 & 0.999917 \end{bmatrix} \tag{23}$$

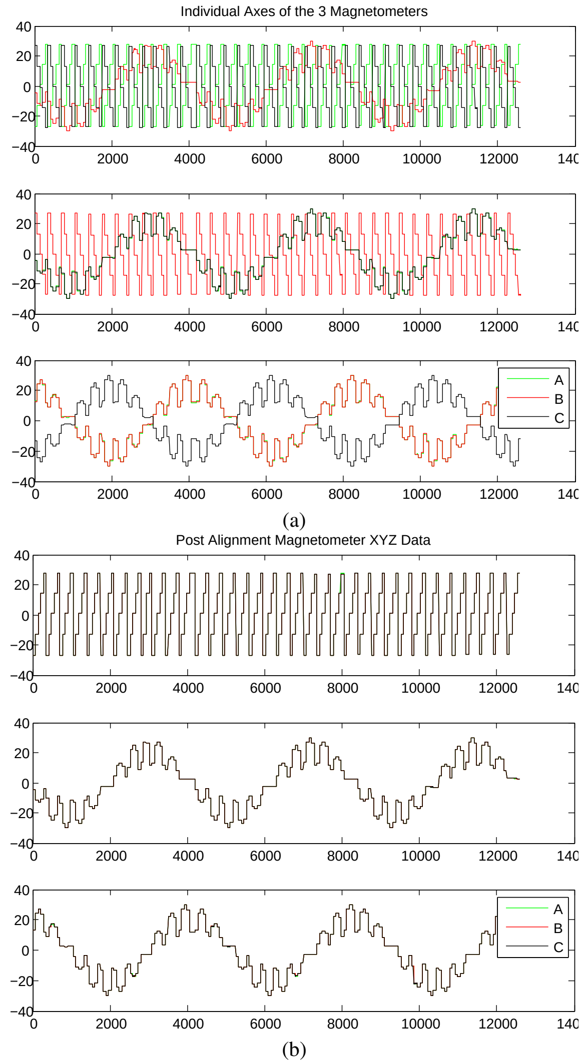

Fig. 7: XYZ data of the three magnetometers: (a) Pre-Alignment and (b) Post-Alignment

and the misalignment between A and C is computed as:

$$\overset{A\to C}{\mathcal{R}_{mis}} = \begin{bmatrix} -0.999919 & -0.0125745 & -0.00213666 \\ -0.0125591 & 0.999896 & -0.00709941 \\ 0.00222571 & -0.007072 & -0.999973 \end{bmatrix} \tag{24}$$

Note that while $\overset{A\to C}{\mathcal{R}_{mis}}$ is somewhat close to the identity matrix (albeit with x and z axes reversed), the one from $A \to B$ is not, indicating that it is mounted differently (and this can be seen in Fig. 6. 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. 7b shows the aligned data, and shows a much better match on the data then the pre-alignment data.

The conclusion from the small improvement from the ensemble average (theoretically the signal to noise should

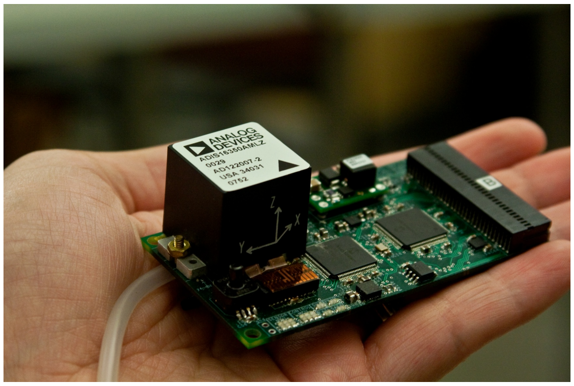

Fig. 8: The SLUGS autopilot

be an improved by $\sqrt{3}$) is that the noise is not independent, but rather from the Helmholz 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.

B. 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. **??**). 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.

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 500 iterations. The final misalignment *DCM* was:

$$\mathcal{R}_{mis} = \begin{pmatrix} 0.9756 & -0.213 & 0.05394 \\ -0.214 & -0.9768 & 0.01293 \\ 0.04993 & -0.02415 & -0.9985 \end{pmatrix} \tag{25}$$

From the misalignment matrix, it can be seen that the internal sensor axes of the magnetometer forms a right handed coordinate frame that is opposite of the accelerometer. That is, the magnetometer is essentially "flipped over" by rotating it 180° around the $x-$axis. More interesting still is the covariance of the misalignment estimate:

$$\mathcal{P}_{body} = \begin{pmatrix} 1.308 & 0.08815 & -0.6806 \\ 0.08815 & 1.916 & -1.497 \\ -0.6806 & -1.497 & 11.11 \end{pmatrix} \tag{26}$$

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.

V. Conclusions

We have developed an algorithm that is capable of estimating the misalignment between multiple three axis sensors from a locus of measurements taken by the sensors in the body frame. This is done via repeatedly applying Wahba's Problem solution to the data and iterating until converged. The algorithm works with either heterogenous sensors 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 large misalignments and quantify them in a mathematically rigorous manner.

Acknowledgements

We would like to thank Creare for providing the magnetometer data from the Helmholz coils. Special thanks goes to Ze-Maria Vasconcelos for countless hours perfecting the algorithm and convincing me that it could work.

References

[1] G. Wahba *et al.*, "Problem 65-1: A least squares estimate of satellite attitude," *Siam Review*, vol. 7, no. 3, p. 409, 1965.

[2] J. Crassidis, F. Markley, and Y. Cheng, "Survey of nonlinear attitude estimation methods," *Journal of Guidance Control and Dynamics*, vol. 30, no. 1, p. 12, 2007.

[3] G. H. Elkaim and C. Foster, "Extension of a non-linear, two-step calibration methodology to include non-orthogonal sensor axes," *IEEE Transactions on Aerospace Electronic Systems*, vol. 44, 07/2008 2008.

[4] J. Vasconcelos, G. H. Elkaim, C. Silvestre, P. Oliveira, and B. Cardeira, "A geometric approach to strapdown magnetometer calibration in sensor frame," *IEEE Transactions on Aerospace Electronic Systems*, 2010.

[5] F. Markley, "Attitude determination using vector observations and the singular value decomposition," *Journal of the Astronautical Sciences*, vol. 36, no. 3, pp. 245–258, July-September 1988.

[6] J. Farrell, J. Stuelpnagel, R. Wessner, J. Velman, and J. Brook, "A least squares estimate of satellite attitude (grace wahba)," *SIAM Review*, vol. 8, no. 3, pp. 384–386, 1966. [Online]. Available: http://epubs.siam.org/doi/abs/10.1137/1008080

[7] M. Lizarraga, R. Curry, and G. Elkaim, "Reprogrammable uav autopilot system (part 1) - system hardware and software," *Circuit Cellar*, no. 249, pp. 24–35, April 2011.

[8] ——, "Reprogrammable uav autopilot system (part 2) - testing and results," *Circuit Cellar*, no. 250, pp. 36–43, May 2011.

[9] M. Lizarraga, G. H. Elkaim, G. Horn, R. Curry, V. Dobrokhodov, and I. Kaminer, "Low cost rapidly reconfigurable uav autopilot for research and development of guidance, navigation and control algorithms," in *ASME/IEEE MESA09*, International Conference on Mechatronic and Embedded Systems and Applications. San Diego, CA: International Conference on Mechatronic and Embedded Systems and Applications, 08/2009 2009.