Magnetometer Realignment: Theory and Implementation

William Premerlani, October 16, 2011

Problem

A magnetometer that is separately mounted from its IMU partner needs to be carefully aligned with the IMU in order to operate properly. Even a slight misalignment will cause large yaw estimation errors during operation in locations that have a large vertical magnetic field component.

Solution

Enlist the aid of the direction cosine matrix to compute and apply a realignment rotation matrix to magnetic measurements.

Background

First, I wish to thank Ric Kuebler and Peter Hollands for motivating this exploration. Peter was the first to point out the issue during an analysis of the MatrixPilot data log from one of Ric's flights. In particular, Peter discovered a spurious report of a reversal of the direction of the earth's magnetic field, as shown in Figure 1.

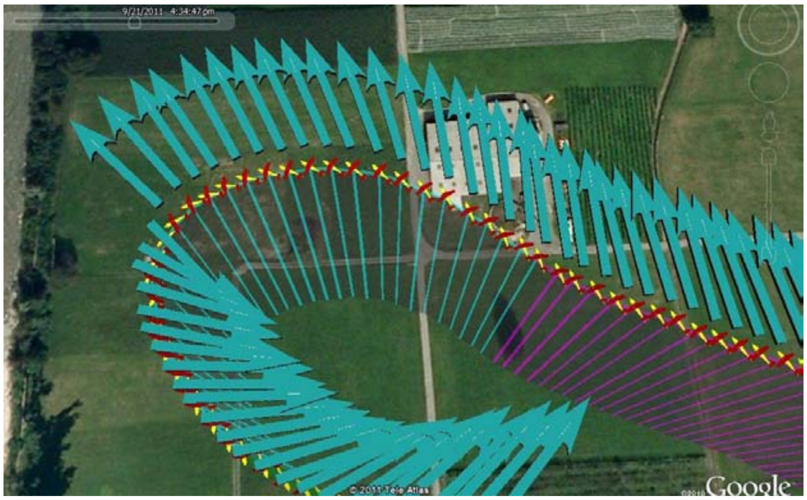

Figure 1 Apparent magnetic field reversal due to magnetometer misalignment

Figure 1 was produced by a flight analysis program that Peter Hollands wrote. The airplane icons indicate the orientation of the aircraft as estimated by the elements of the direction cosine matrix. The blue arrows indicate the direction of the horizontal component of the magnetic field, as seen in the earth frame of reference. If everything is working correctly, the blue arrows should always point in the same direction. In this case, the indicated direction of the magnetic field flipped 180 degrees. How could this be?

A closer look at the raw data revealed that as the plane rolled through the turn, the magnitude of the measured horizontal component of the magnetic field in the earth frame of reference went through zero, and then the vector reversed. With the UAVDevBoard that Ric was using, the magnetometer is mounted separately, so I suspected there might have been an alignment problem. Ric confirmed that there might have been roll misalignment of the magnetometer. He changed the way it was mounted to assure better alignment. During subsequent flights, the issue disappeared, but the incident generated some questions: How closely does the magnetometer need to be aligned with the IMU to get good performance? Is there any way to determine the misalignment in flight and correct for it?

So, the next thing I did was some ground testing of the effects of magnetometer misalignment. The first test was done with a magnetometer that was carefully aligned with a UAVDevBoard. The results are plotted in Figure 2.

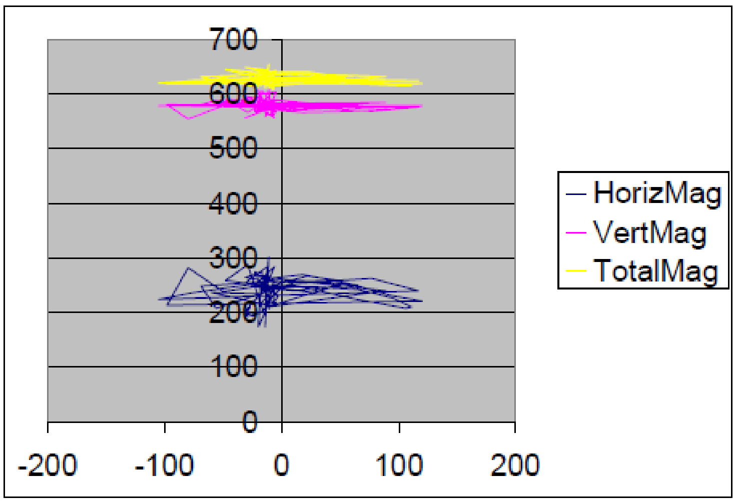

Figure 2 Magnetic field in earth frame versus heading with a carefully aligned magnetometer

In Figure 2 the horizontal, vertical, and the total magnetic field in the earth frame of reference are plotted against heading. Heading is represented with an 8 bit value in which -128 binary indicates -180 degrees, and 127 indicates approximately 180 degrees. The plot shows little variation in the magnetic field.

The next step was to deliberately misalign the magnetometer and record the magnetic field measurements versus heading. The first test was done with a 6 degree tilt misalignment of the magnetometer with the IMU level. The results are shown in Figure 3.

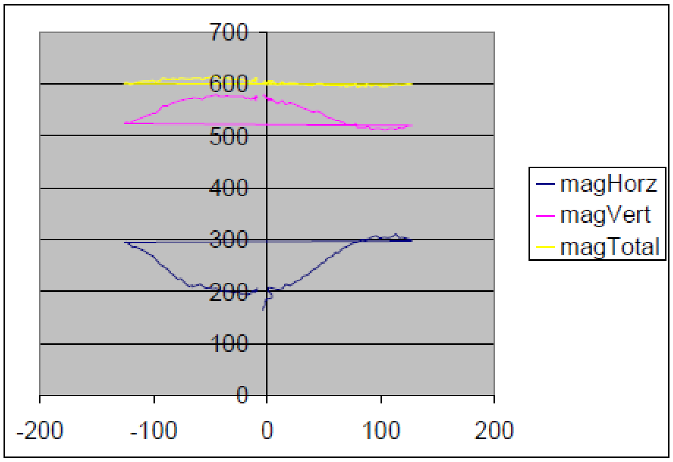

Figure 3 Impact of a 6 degree misalignment on estimated magnetic field versus heading

The next test was done with a 12 degree misalignment. The results, shown in Figure 4, were even more dramatic. The horizontal field is getting close to reversing.

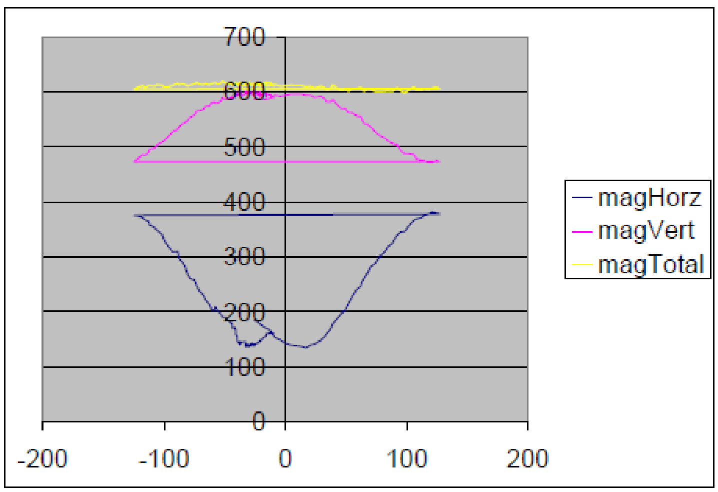

Figure 4 Impact of a 12 degree misalignment on estimated magnetic field versus heading

Figures 3 and 4 show that a slight misalignment can cause significant problems in magnetometer performance. That is because, at my location, the vertical component of the earth's magnetic field is so much stronger than the horizontal component. Notice that in Figure 2 the vertical component is nearly 3 times as large as the horizontal component. Also note that the testing that I did was with a level IMU. The error would be even worse for certain combinations of heading and roll angle.

So, it struck me that there is an operational issue with separate mounting of IMU and magnetometer. Either the magnetometer should be carefully aligned with the IMU, or the software needs to do something about it. It was time to look at the math.

Theory

The math for this problem turned out to be very challenging. If you are not interested in math, you might want to skip to the section on implementation.

I discovered three different ways to solve the problem:

1. Compare the direction of the measured magnetic field in the earth frame of reference with the actual direction.
2. Compare the directions of two magnetic field measurements in the body frame of reference.
3. Compare the directions of two magnetic field measurements in the earth frame of reference.

I worked out all three solutions, and then opted to implement the third one. Interested readers might want to work out solutions 1 and 2 for themselves.

The starting point is the relationship between the magnetic measurement in the body frame and the actual magnetic field in the earth frame, taking into account the misalignment:

$$\mathbf{b_B} = \mathbf{A} \cdot \mathbf{R^T} \cdot \mathbf{b_E}$$

$\mathbf{b_E}$ = magnetic field in earth frame

$\mathbf{b_B}$ = magnetic field measurement in body frame equation 1

$\mathbf{R^T}$ = the transpose of the direction cosine matrix

$\mathbf{A}$ = rotational matrix representing the effect of misalignment

The issue created by the magnetometer misalignment can be seen in the expression for the estimated magnetic field in the earth frame of reference:

$$\hat{\mathbf{b}}_\mathbf{E} = \mathbf{R} \cdot \mathbf{A} \cdot \mathbf{R^T} \cdot \mathbf{b_E}$$

equation 2

$\hat{\mathbf{b}}_\mathbf{E}$ = estimated magnetic field in earth frame of reference

The problem is that the misalignment rotation matrix is between the attitude rotation matrices. If there were no misalignment, the A matrix would be the identity matrix, and the multiplication of the other two matrices would produce the identity matrix, so the estimated field would match the actual field. But the product of the three matrices in equation 2 is not equal to an identity matrix. Students of linear algebra will recognize that the form of the product of the three rotational matrices in equation 2 is a very special one: it represents the transformation of the rotational matrix A from the body frame into the earth frame. In other words, the estimated earth magnetic field becomes rotated by the misalignment, applied in the earth frame rather than the body frame. Equation 2 completely explains figures 3 and 4. There was a tilt misalignment. In the earth frame, the magnetic field measurements appeared tilted. The apparent tilt in the earth frame is equal to the actual tilt misalignment in the body frame, rotated from the body frame to the earth frame. That why the tilt in the earth frame (and the error) depends on the heading. The heading governs the transformation of the tilt from body to earth frame of reference.

Next, consider a pair of measurements:

$$\mathbf{b_{B1}} = \mathbf{A} \cdot \mathbf{R_1^T} \cdot \mathbf{b_E}$$
$$\mathbf{b_{B2}} = \mathbf{A} \cdot \mathbf{R_2^T} \cdot \mathbf{b_E} \qquad \text{equation 3}$$

Transform the measurements into the earth frame of reference:

$$\mathbf{b_{E1}} = \mathbf{R_1} \cdot \mathbf{A} \cdot \mathbf{R_1^T} \cdot \mathbf{b_E}$$
$$\mathbf{b_{E2}} = \mathbf{R_2} \cdot \mathbf{A} \cdot \mathbf{R_2^T} \cdot \mathbf{b_E} \qquad \text{equation 4}$$

Combine and rearrange the pair of equations 4 to relate the two measurements by a rotational transformation:

$$\mathbf{b_{E2}} = \mathbf{R_2} \cdot \mathbf{A} \cdot \mathbf{R_2^T} \cdot \mathbf{R_1} \cdot \mathbf{A^T} \cdot \mathbf{R_1^T} \cdot \mathbf{b_{E1}} \qquad \text{equation 5}$$

It might seem that I am digging myself into a hole with equation 5. There are 6 rotational matrices to be multiplied, and the unknown alignment matrix appears twice! But there is some hope of solving the equation because of the special form of some of the factors.

Equation 5 says that the two measured magnetic fields, when transformed into the earth frame of reference, are related by a rotation. The rotation is not uniquely determined by the two measurements. However, we can gradually determine the alignment rotation with a series of measurements, so at this point we can use the rotational matrix that relates the two measurements with the smallest angle of rotation. Such a minimum angle rotational matrix E can be computed from the cross product of two measurements such that:

$$\mathbf{b_{E2}} = \mathbf{E} \cdot \mathbf{b_{E1}} \qquad \text{equation 6}$$

From equation 5 and 6, we conclude equation 7.

$$\mathbf{E} = \mathbf{R_2} \cdot \mathbf{A} \cdot \mathbf{R_2^T} \cdot \mathbf{R_1} \cdot \mathbf{A^T} \cdot \mathbf{R_1^T} \qquad \text{equation 7}$$

Multiply equation 7 on the left by R2 transpose and on the right by R1 to arrive at equation 8.

$$\mathbf{R_2^T} \cdot \mathbf{E} \cdot \mathbf{R_1} = \mathbf{A} \cdot \mathbf{R_2^T} \cdot \mathbf{R_1} \cdot \mathbf{A^T} \qquad \text{equation 8}$$

Now I am going to take advantage of a very special property of the form of the product of the matrices on the right hand side of equation 8. But before I do that, I need to say something about an operator called "vex". Vex extracts certain information from a rotation matrix. It converts a rotational matrix into a vector representation such that:

$$\mathbf{vex}(\mathbf{R}) = \sin(\alpha) \cdot \begin{bmatrix} n_1 \\ n_2 \\ n_3 \end{bmatrix}$$
$$\alpha = \text{angle of rotation} \qquad \text{equation 9}$$
$$n_1, n_2, n_3 = \text{components of the normalized axis vector}$$

The vex operator is easy to implement. You take one half of the difference between a rotation matrix and its transpose, and then transfer three particular off-diagonal elements of the result into the components of the vector.

The reason I want to use the vex operator is because it will remove one of the factors involving the alignment matrix. It can be shown that for rotation matrices, equation 10 is an identity:

$$\mathbf{vex}\left(\mathbf{A}\cdot\mathbf{R}_2^{\mathbf{T}}\cdot\mathbf{R}_1\cdot\mathbf{A}^{\mathbf{T}}\right)=\mathbf{A}\cdot\mathbf{vex}\left(\mathbf{R}_2^{\mathbf{T}}\cdot\mathbf{R}_1\right) \qquad \text{equation 10}$$

I leave it to the interested reader to prove equation 10. I suspected it was true from some geometrical reasoning, it took me a while to prove it really is true.

Application of the vex operator to both sides of equation 8 yields equation 11.

$$\mathbf{vex}\left(\mathbf{R}_2^{\mathbf{T}}\cdot\mathbf{E}\cdot\mathbf{R}_1\right)=\mathbf{A}\cdot\mathbf{vex}\left(\mathbf{R}_2^{\mathbf{T}}\cdot\mathbf{R}_1\right) \qquad \text{equation 11}$$

Equation 11 implies several interesting things about the relationship between the misalignment matrix and the apparent rotation of the magnetic measurement in the earth frame of reference. First, because rotational matrices preserve vector magnitudes, equation 11 implies that the magnitude of the rotation vector produced by the vex operator on the left hand side of equation 11 is exactly equal to the magnitude of the rotation vector produced by the vex operator on the right hand side. In other words, inserting the error rotation between R2 transpose and R1 does not change the amount of rotation implied by R2 transpose times R1, it only changes its direction.

Equation 11 also implies that the misalignment matrix transforms the vector on the right hand side of equation 11 in exactly the same way that it transforms the magnetic measurement. So all we have to do to discover the misalignment matrix is to build a matrix that rotates the vector on the right hand side of equation 11 to the vector on the left hand side of equation 11. Both vectors in the equation can be computed by performing the indicated operations on matrices that we have or which can be easily computed, and then applying the vex operation. In other words, equation 11 implies a method for computing the misalignment matrix.

Once we have the misalignment matrix, we can multiply the magnetometer measurement vectors by the transpose of the misalignment matrix to undo the effects of the misalignment:

$$\hat{\mathbf{b}}_{\mathbf{B}}=\hat{\mathbf{A}}^{\mathbf{T}}\cdot\mathbf{b}_{\mathbf{B}}$$

$$\hat{\mathbf{b}}_{\mathbf{B}}=\text{realigned measurement in the body frame} \qquad \text{equation 12}$$

$$\hat{\mathbf{A}}^{\mathbf{T}}=\text{transpose of the estimate of the misalignment matrix}$$

Some of you are shaking your head at this point and are asking, what is the implementation.

Implementation

Implementation is easier than it might seem. Much of it can be done by taking vector cross products. The rest is mostly matrix multiplications and the vex operator.

You need to update the estimate of the realignment matrix each time a new measurement becomes available, and apply it to each magnetometer measurement.

I found it easier to update a quaternion representation of the rotation, and then convert it to a direction cosine matrix. (A quaternion is not the same thing as the right side of equation 9, but they are similar. The difference is a quaternion uses the sine of one half of the rotation angle. It is easy to transform either vector representation into an equivalent rotation matarix.) Each time a new magnetometer measurement is available, perform the following steps:

1. Remove the offset vector from the measurement. (Offset compensation is a separate topic.)
2. Convert from a quaternion vector representation of the realignment, **align**, to a direction cosine matrix, **A**, and then apply the matrix to the measurement vector to realign it.
3. Use the realigned measurement for whatever other calculations require a magnetic field vector.
4. Use the direction cosine matrix **R2** to transform the magnetic field vector from the body frame into the earth frame, and normalize the result into a vector with magnitude equal to 1.
5. Take the cross product of the previous normalized earth frame magnetic field vector with the most recent one. Use the result to build the error rotation matrix, **E**, using the inverse of the vex operator.
6. Compute the vector **P** by applying the **vex** operator to **R2** transpose times **E** times **R1**, where **R2** is the direction cosine matrix corresponding to the most recent magnetic measurement, and **R1** corresponds to the previous measurement.
7. Compute the vector **Q** by applying the **vex** operator to **R2** transpose times **R1**.
8. Compute an adjustment vector, **adjust,** by taking the cross product of **Q** times **P**.
9. Update the alignment quaternion, **align**, by setting **align** equal to **align** minus gain times **adjust**. Select gain to strike a balance between stability and speed of convergence. I used a value of 1/16 for the gain.

I integrated the algorithm into MatrixPilot running on the UAV DevBoard. It works quite nicely. Each time the board and the magnetometer rotate a little bit, the estimate of the realignment matrix improves. Gradually the process converges on an exact realignment matrix. It will compensate for any amount of misalignment, including a 180 degree rotation, although the amount of time needed for convergence is proportional to the initial amount of misalignment. It will also compensate for misalignments along any of the three axes, including the yaw axis. In order for the method to compensate for an alignment around a particular axis, there must be rotation around one of the other two perpendicular axes. So, for example, in order to compensate for yaw misalignment, there must be roll or pitch rotations.

Testing

Some of the testing results are shown in the following figures. Actually, there was quite a bit of testing because the development process was not nearly as direct as suggested by the previous, for several reasons. First, I discovered that magnetometer misalignment interfered with offset removal. So I had to interrupt the work on realignment and go off and improve the offset removal, which is the topic of a separate report. Second, I was not entirely sure when I started whether the idea would even work, so some of the initial testing was done with misalignment that was simulated in software. Finally, I gradually realized that the technique would work on any amount of misalignment, but I would have to use a full rotation matrix, rather than an approximate one that I used during early versions of the idea.

Figure 5 is a plot of the realignment quaternion versus sample number for a 90 degree yaw misalignment between the IMU and the magnetometer. Samples were taken 4 times per second. The quaternion components are represented in a binary format equal to 2**14 times the quaternion values. So, a 90 degree yaw should be 11585 for the Z component, with the X and Y components equal to 0. During the test I was randomly rotating the IMU and the magnetometer, which were mounted together on a piece of plastic with the magnetometer rotated by 90 degrees. Convergence took about 30 seconds.

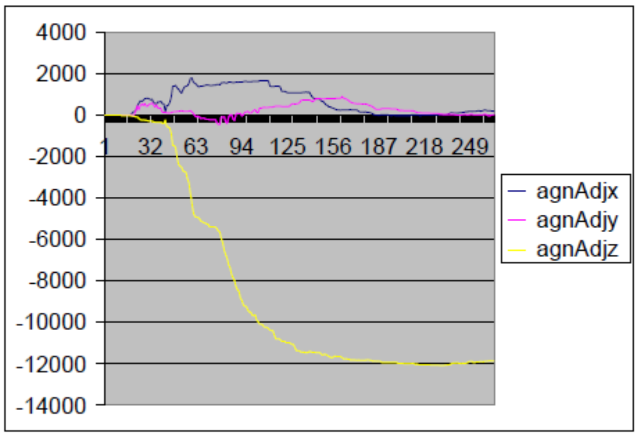

Figure 5 Convergence of realignment parameters for a 90 degree yaw misalignment

A plot of the horizontal component of the magnetic field, after transforming the realigned measurements to the earth frame of reference, is shown in Figure 6.

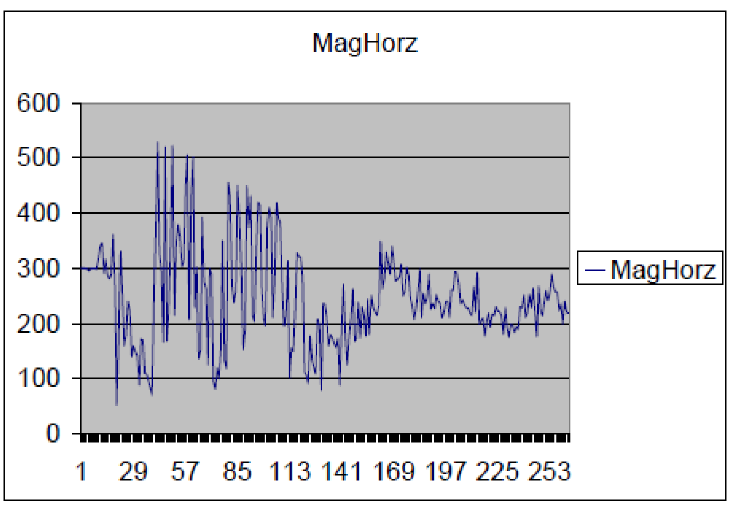

Figure 6 Horizontal magnetic field in earth frame with 90 degree yaw misalignment

The large initial errors in the measured horizontal component of the magnetic field are generated by the large yaw misalignment as I rolled and pitched the system in random directions. However, as the realignment converged, the variations dramatically reduced.

Similar plots for a full 180 degrees yaw misalignment are shown in figures 7 and 8.

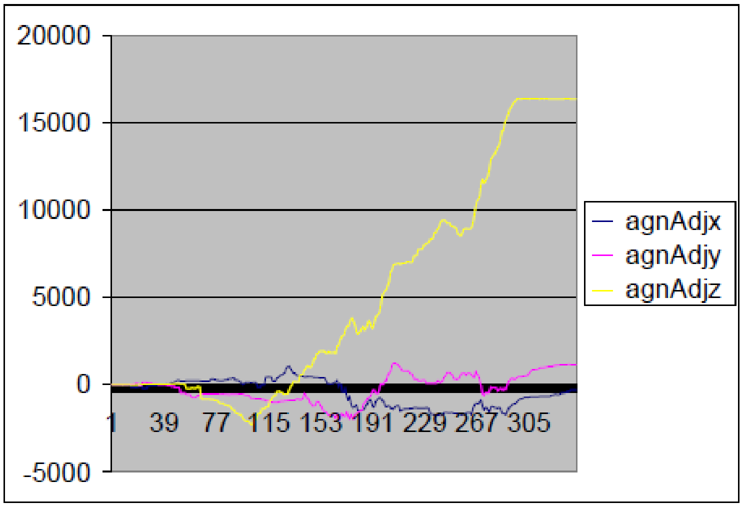

Figure 7 Convergence of realignment parameters for a 180 degree yaw misalignment

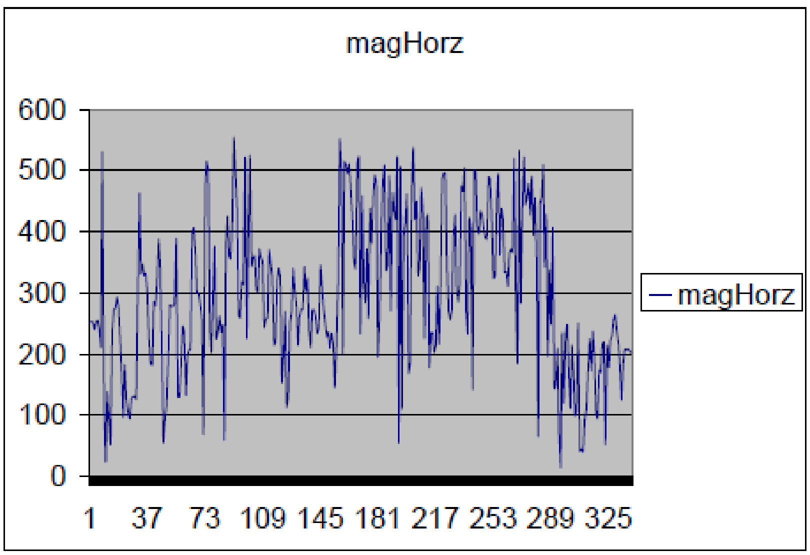

Figure 8 Horizontal component of magnetic field in earth frame with 180 degree yaw misalignment

In the case of the 180 degree misalignment, it took about one minute for the algorithm to converge, it converged near the end of the test. Note that the alignment parameters reached the correct values, and the variance in the horizontal magnetic field reduced, at the end of the test. During the test I had no idea how well it was converging, so I had no idea how long to continue the test.

Another test that I conduct whenever I work on the magnetometer software is a "spin test". I place the IMU and the magnetometer on a record player, and spin them at 78 RPM for several minutes. This is an especially tough test to pass, for several reasons:

1. Gyro gain error tends to break yaw lock. Earlier this year I devised a way for the gyros to become "self-calibrating" to improve performance in this area. Whenever I make changes, I want to make sure this feature still works.
2. The record player distorts the local magnetic field and creates magnetic interference.
3. It is difficult to get the IMU's accelerometer exactly at the center of the turntable, so there is usually some error due to centrifugal effects.

Figure 9 is a plot of the X, Y, and Z component of the measured magnetic field in the earth frame of reference. (Note: the declination of the earth's magnetic field at my location is -14 degrees.) The plot shows the variance reducing due to the action of the offset removal algorithm and the realignment algorithm. It also shows that there was a solid yaw lock during the tests, especially evidenced at the end of the test, when the turntable stopped spinning, there was no sign that either the IMU was "dizzy", or that there was residual offset or misalignment.

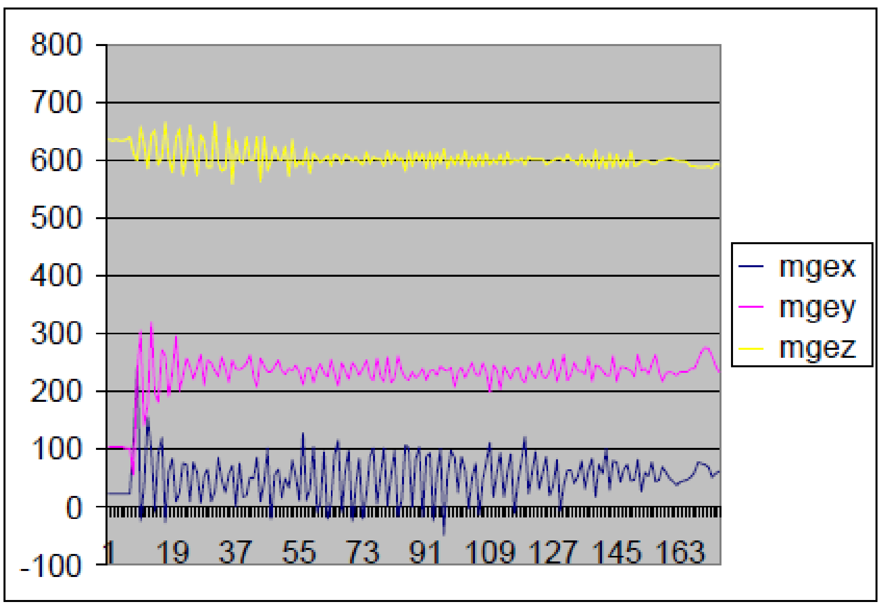

Figure 9 Plot of the X, Y, and Z component of the measured magnetic field in the earth frame, during an extended spin test at 78 RPM