CH Robotics

AN-1008 - Sensors for Orientation Estimation

Document rev. 1.0, updated 10/20/2012

1. Introduction

CH Robotics AHRS and Orientation Sensors combine information from accelerometers, rate gyros, and (in some cases) GPS to produce reliable attitude and heading measurements that are resistant to vibration and immune to long-term angular drift.

Part of what makes reliable orientation estimates possible is the complementary nature of the different kinds of sensors used for estimation. Rate gyros provide good short-term stability and resistance to vibration, accelerometers can provide attitude information that does not become less reliable over time, and magnetic sensors provide heading information in addition to limited attitude information (pitch and roll). By combining data from each type of sensor using a filter such as an Extended Kalman Filter, accurate orientation estimates can be computed.

To fully understand the capabilities and limitations of CH Robotics devices, it helps to have a clear understanding of the operation of the different types of sensors that used to compute attitude and heading. In this application note we describe in detail the operation of accelerometers, rate gyros, and magnetic sensors. We present high level descriptions of how the different types of sensors work, in addition to mathematical models used to quantify the performance characteristics of each type of sensor.

CH Robotics

AN-1008 - Sensors for Orientation Estimation

Document rev. 1.0, updated 10/20/2012

Table of Contents

List of Figures

2. Accelerometers

The output of a three-axis accelerometer mounted to a rigid body can be modeled as

$$\mathbf{a}_m = \frac{1}{m}\left(\mathbf{F} - \mathbf{F}_g\right), \tag{1}$$

where $\mathbf{a}_m$ is the measured acceleration, m is the mass of the body, $\mathbf{F}$ is the sum of all forces on the body (including gravity) expressed in the sensor body frame, and $\mathbf{F}_g$ is the force due to gravity, also expressed in the body frame (See application note *AN1005 - Understanding Euler Angles* for details about different coordinate frames).

The accelerometer model in Equation (1) can be intuitively explained using the highly simplified one-axis accelerometer diagram shown in Figure 1. As shown, an accelerometer can be constructed by attaching a proof-mass to a lever-arm. Any upward or downward acceleration of the sensor in Figure 1 causes deflection of the proof mass, which can be measured to determine acceleration (below the natural frequency of the lever/mass system, deflection is proportional to acceleration). This accounts for the first force term, $\mathbf{F}$, in Equation (1): the sum of all forces on the body, including gravity, produces physical acceleration of the sensor, which in turn causes deflection of the proof mass.
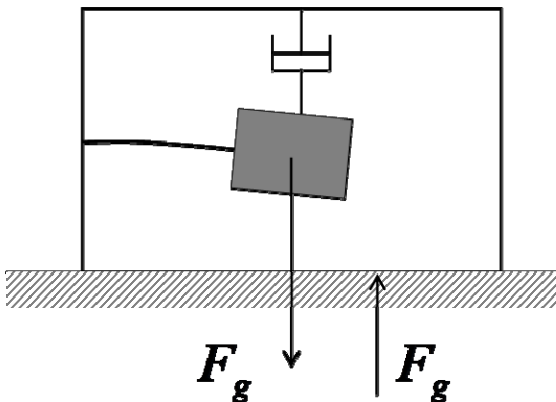

Figure 1 - Simplified One-Axis Accelerometer Diagram

AN-1008 - Sensors for Orientation Estimation

Document rev. 1.0, updated 10/20/2012

The second force term, $\mathbf{F}_g$, is present in the accelerometer model because the force of gravity not only accelerates the sensor body, it also causes deflection of the proof-mass itself. If the accelerometer is not accelerating ($\mathbf{F} = 0$), then gravity produces a downward deflection of the proof mass that appears equivalent to upward acceleration of the sensor at the acceleration of gravity. Similarly, if the accelerometer is in free-fall ($\mathbf{F} = \mathbf{F}_g$) there is no deflection and hence no measured acceleration despite the fact that the accelerometer is accelerating at 9.8 m/s/s toward the center of the earth.

The accelerometer model in Equation (1) suggests that the expected accelerometer measurement can be obtained by drawing a free-body diagram that includes all forces except the force of gravity: the first term, $\mathbf{F}$, includes gravity, while the second term, $\mathbf{F}_g$, removes it. From this perspective, accelerometers never measure gravity directly - they measure the normal forces that prevent the sensor from accelerating toward the center of the Earth (in addition to other external forces).

Going back to Equation (1), let us assume that the external forces acting on the sensor are trivial in comparison with the normal forces. Then the accelerometer output can be approximately modeled as

$$\mathbf{a} = -\frac{\mathbf{F}_g}{m}, \qquad\qquad (2)$$

from which it is easy to back out the pitch and roll angles of the sensor by simply noting which direction gravity the gravity vector points. For example, if it points straight down, then the sensor is level with pitch = roll = 0. This is illustrated in Figure 2. Notice that the x and y accelerometer outputs are nearly zero, while the z-axis output is negative one gravity.
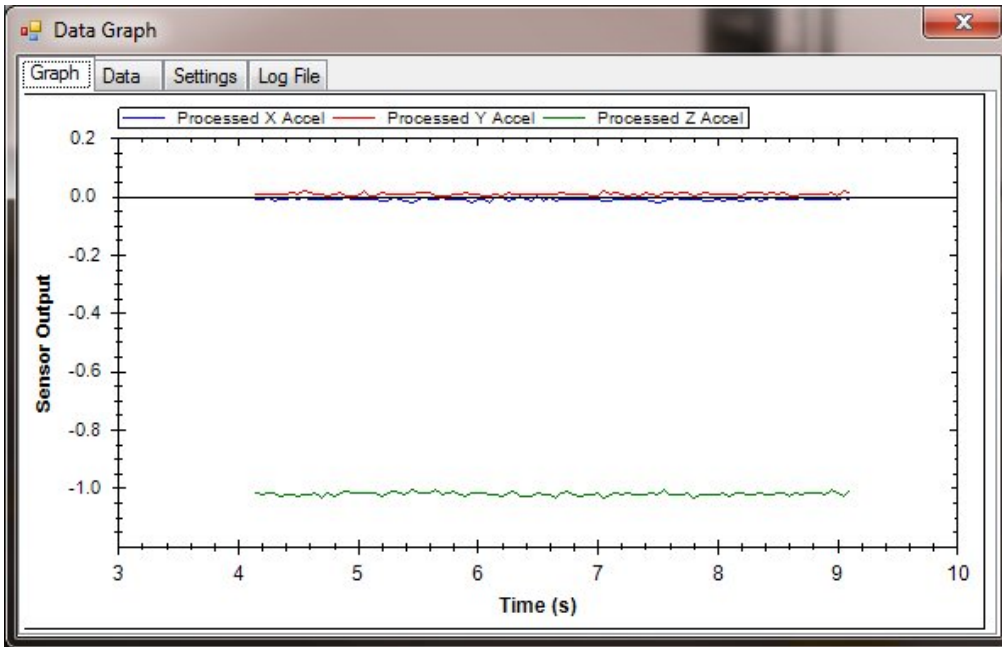
Figure 2 - Accelerometer output when sensor is flat and level

On the UM6 and the UM6-LT, it is assumed that in general the external forces acting on the accelerometer are close to zero, so that the pitch and roll angles can be measured directly using Equation (2). When this assumption is not valid (i.e. on an airplane, a helicopter, or motorcycle - anything that spends a lot of time accelerating aggressively), the angle estimates of the sensor can degrade. This will be true of *any* orientation sensor or AHRS that doesn't utilize an external aid like GPS.

The simplified accelerometer given in Equation (1) does not include terms to account for cross-axis alignment, scale factors, and output bias. All of these items can increase the amount of error in the output of the sensor, affecting the accuracy of orientation estimates and (depending on how the accelerometer is being used) on the accuracy of dead-reckoning velocity and position estimates. A more complete model of the accelerometer output is given by

$$\mathbf{a}_m = M_a \left(\frac{1}{m} S_a(T)(\mathbf{F} - \mathbf{F}_g) - \beta_a(T) \right), \tag{3}$$

where M_a is the accelerometer misalignment matrix, $\beta_a(T)$ is a vector of temperature-variant output biases, and $S_a(T)$ is the diagonal temperature-variant accelerometer sensitivity matrix given by

$$S_a(T) = \begin{pmatrix} S_{a,x}(T) & 0 & 0 \\ 0 & S_{a,y}(T) & 0 \\ 0 & 0 & S_{a,z}(T) \end{pmatrix}. \qquad (4)$$

The sensitivity matrix is what encodes the expected accelerometer raw output for a given measured acceleration. The misalignment matrix describes the effect of cross-axis misalignment and, unlike the bias and sensitivity terms, it is not affected be temperature.

When using accelerometer data, we measure $\mathbf{a}_m$ but what we really want to know is the actual acceleration before scale factors, biases, and misalignment distort the measurement. That is, we want to take the measurement and extract the term $\frac{1}{m}(\mathbf{F} - \mathbf{F}_g)$. Solving Equation (3), we get

$$\mathbf{a}_{m,c} = S_a^{-1}(T)(M_a^{-1}\mathbf{a}_m + \beta_a(T)). \qquad (5)$$

where $\mathbf{a}_{m,c}$ is defined as the corrected accelerometer measurement vector. In order to obtain the best accuracy, the terms $S_a^{-1}(T)$, M_a^{-1}, and $\beta_a(T)$ should be determined experimentally over the full operating temperature of the sensor. At the time of writing, the UM6 and the UM6-LT are not calibrated in-factory so that costs can be kept low. However, new products are scheduled for release in Q2 2013 that will come factory-calibrated for the best possible output accuracy.

3. Rate Gyros

The output of a three-axis rate gyro can be modeled as

$$\mathbf{p}_m = M_g \left(S_g(T)\mathbf{p} - \beta_g(T) + C_{g,a}\mathbf{a}_{m,c} \right), \tag{6}$$

where $\mathbf{p}$ is a vector of the actual body frame angular rates, $\mathbf{p}_m$ is the measured angular rate, M_g is the gyro misalignment matrix, $S_g(T)$ is the gyro sensitivity matrix, $\beta_g(T)$ is a vector of temperature-varying gyro output bias, and $C_{g,a}$ is a matrix that specifies the amount of coupling between physical acceleration of the rate gyro and deviation in the output angular rate. Usually coupling is small enough that $C_{g,a}$ can be neglected, although in cases where high acceleration is expected (e.g. on a rocket), this term must be measured and included in the model for the best accuracy. The angular rate vector $\mathbf{p}$ is more explicitly defined as

$$\mathbf{p} = \begin{pmatrix} p \\ q \\ r \end{pmatrix}, \tag{7}$$

where p is angular rate about the body-frame x-axis, q is the angular rate about the body-frame y-axis, and r is the angular rate about the body-frame z-axis. The measured angular rate vector $\mathbf{p}_m$ is similarly defined. The sensitivity matrix $S_g(T)$ is defined as

$$S_g(T) = \begin{pmatrix} S_{g,x}(T) & 0 & 0 \\ 0 & S_{g,y}(T) & 0 \\ 0 & 0 & S_{g,z}(T) \end{pmatrix}. \tag{8}$$

AN-1008 - Sensors for Orientation Estimation

Document rev. 1.0, updated 10/20/2012

Given a rate gyro measurement $\mathbf{p}_m$, we want to extract the actual angular rate. Solving Equation (6) for $\mathbf{p}$, we get

$$\mathbf{p} = S_g^{-1}(T)\left(M_g^{-1}\mathbf{p}_m + \beta_g(T) - C_{g,a}\mathbf{a}_{m,c}\right)$$

(9)

For best accuracy, the terms $S_g^{-1}(T)$, M_g^{-1}, $\beta_g(T)$, and $C_{g,a}$ should be determined experimentally so that compensation can be performed. However, of all the terms affecting the output error of the gyro measurement, the gyro bias and sensitivity are the largest contributors of error, especially over a wide temperature range. If there is insufficient equipment available to perform a complete calibration, the most important calibration procedure is to remove temperature-induced output bias changes, followed by the removal of temperature-induced scale factor changes.

At the time of writing, CH Robotics keeps costs low by not calibrating the gyros on the UM6 and the UM6-LT in-house. However, beginning Q1 2013, calibration will be provided as an optional feature for applications where higher accuracy is required.

4. Magnetometers

Of all the sensors used by CH Robotics for attitude and heading estimation, magnetometers are perhaps the most tricky to work with for a variety of reasons. First, like accelerometers and rate gyros, magnetic sensors are subject to cross-axis misalignment and temperature-varying scale factors and biases. In addition, the local magnetic field around the magnetometer can be distorted by magnetic and ferrous metal objects. This distortion must be measured and corrected in order for the magnetometer to be useful. The field can also be distorted by time-varying electromagnetic waves from high-power electrical lines and communication lines placed near the sensor. Finally, as the sensor is moved around in practice, the field can be distorted so that, in certain locations (say, next to a concrete wall filled with rebar), the accuracy of the measurement is reduced. All of these factors places some limitations on the use of sensors that rely on magnetometers for heading estimation.

AN-1008 - Sensors for Orientation Estimation

Document rev. 1.0, updated 10/20/2012

We begin with a sensor model essentially equivalent to the gyro and accelerometer models already presented, resulting in the following equation for correcting cross-axis misalignment, biases, and scale factors:

$$\mathbf{b} = S_b^{-1}(T)\left(M_b^{-1}\mathbf{b}_m + \beta_b(T)\right). \tag{10}$$

Similar to the rate gyro and accelerometer cases already discussed, the vector $\mathbf{b}$ represents the actual magnetic field, $\mathbf{b}_m$ is the measured magnetic field, $S_b^{-1}(T)$ is the inverse of the magnetometer sensitivity matrix, M_b^{-1} is the inverse of the magnetometer misalignment matrix, and $\beta_b(T)$ is the magnetometer bias vector.

After the magnetometer measurement is corrected for axis misalignment, scale factor, and bias errors, it must be corrected for additional local magnetic field distortions. There are two types of distortions that affect the measurement: soft-iron distortion and hard-iron distortions.

A soft-iron field distortion is caused by ferrous metal objects that bend the Earth's magnetic field. For example, the screws in the UM6 enclosure bend the local magnetic field. To obtain reasonable accuracy, this distortion must be calibrated out of the measurement. This is actually done in-factory for the UM6. A hard-iron distortion is caused by an object (like the permanent magnets in a motor that has its own permanent magnetic field.

In the absence of any hard or soft-iron distortions, the outputs of the magnetometer as it is rotated should trace out a perfect sphere centered around zero. Soft-iron distortions distort the sphere so that it appears to be an ellipsoid. Hard-iron distortions shift the mean of the sphere away from zero.

One way to calibrate the magnetometer for hard and soft iron distortions is to fit an ellipsoid to a set of magnetometer data collected over a wide variety of different orientations. The center of the ellipsoid is equivalent to the bias caused by hard iron distortions, while the ellipsoid shape on the x, y, and z axes are caused by the soft-iron distortions. A correction matrix can then be calculated that, when multiplied by the magnetometer measurements, alters the ellipsoid so that it looks like a

sphere again. This is essentially the approach taken by the magnetometer calibration algorithm that is included in the CHR Serial Interface software, which can be downloaded from the CH Robotics website (http://www.chrobotics.com).

There are many other methods for calibrating magnetometers, but in most cases, the calibration algorithm produces a matrix D_s and a vector β_h such that the corrected magnetometer measurement $\mathbf{b}_c$ is given by

$$\mathbf{b}_c = D_s \mathbf{b} - \beta_h .\tag{11}$$

Combining Equation (11) with Equation (10), the complete calibration equation is given by

$$\mathbf{b}_c = D_s S_b^{-1}(T)\left(M_b^{-1}\mathbf{b}_m + \beta_b(T)\right) - \beta_h .\tag{12}$$

Terms in Equation (12) can be combined so that there are fewer calibration terms to estimate during calibration, although the process is complicated by the fact that the bias, scale factor, and cross-axis alignment terms will remain constant, while the hard and soft iron calibration terms will change depending on where the end-user mounts the sensor in the end application.

5. Disclaimer

This document is provided as reference only; "typical" device specifications must be evaluated by the end-user. CH Robotics reserves the right to modify this document and the products it describes without notice.

CH Robotics products are not intended for use in weapons systems, aircraft, life-saving or life-sustaining systems, automobiles, or any other application where failure could result in injury, death, property damage, or environmental damage.