Control Engineering Practice 18 (2010) 712–722

Contents lists available at ScienceDirect

Control Engineering Practice

journal homepage: www.elsevier.com/locate/conengprac

Design and implementation of a low-cost observer-based attitude and heading reference system

Philippe Martin [a,*], Erwan Salaün [b]

[a] Centre Automatique et Systèmes, Mines ParisTech, 60 boulevard Saint-Michel, 75272 Paris Cedex 06, France
[b] School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0150, USA

ARTICLE INFO

Article history:
Received 31 March 2009
Accepted 27 January 2010
Available online 11 February 2010

Keywords:
Attitude estimation
Strapdown systems
Nonlinear observers
Extended Kalman filter
Inertial navigation
Symmetries

ABSTRACT

A nonlinear observer (i.e. a "filter") is proposed for estimating the attitude of a flying rigid body, using measurements from low-cost inertial and magnetic sensors. It has by design a nice geometrical structure appealing from an engineering viewpoint; it is easy to tune, computationally very thrifty, and with guaranteed (at least local) convergence around every trajectory. Moreover it behaves sensibly in the presence of acceleration and magnetic disturbances.

Experimental comparisons with a commercial device illustrate its good performance; an implementation on an 8-bit microcontroller with very limited processing power demonstrates its computational simplicity.

1. Introduction

Aircraft, especially unmanned aerial vehicles (UAV), commonly need to know their attitude to be operated, whether manually or with computer assistance. When cost or weight is an issue, using very accurate inertial sensors for "true" (i.e. based on the Schuler effect due to a non-flat rotating Earth) inertial navigation is excluded. Instead, low-cost systems—often called attitude and heading reference systems (AHRS)—rely on light and cheap strapdown gyroscopes, accelerometers and magnetometers. The various measurements are "merged" according to the motion equations of the aircraft assuming a flat non-rotating Earth, usually with some kind of "filter"; for more details about avionics, various inertial navigation systems and sensor fusion, see for instance Kayton and Fried (1997) and Grewal, Weill, and Andrews (2007).

The attitude estimation problem has received a lot of attention especially in the aerospace engineering community, see the recent survey Crassidis, Markley, and Cheng (2007) and the references therein. By far the most widely used approach is the extended Kalman filter (EKF) and its variants, see e.g. Shuster and Oh (1981), Lefferts, Markley, and Shuster (1982) and Markley (2003). While it is a general method capable of good performance when properly tuned, the EKF suffers from several drawbacks: it is not easy to choose the numerous parameters; it is computationally expensive, which is a problem in low-cost embedded systems; it is usually difficult to prove the convergence, and the designer has to rely on extensive simulations.

An alternative route is to use a dedicated nonlinear observer as proposed in Thienel and Sanner (2003) and Mahony, Hamel, and Pflimlin (2008). The present paper follows the same lines; it uses the rich geometric structure of the attitude-heading problem to derive an observer by the method developed in Bonnabel, Martin, and Rouchon (2008), building up on the preliminary work Martin and Salaün (2007, 2008a). The proposed observer has by design a nice geometrical structure appealing from an engineering viewpoint; it is easy to tune, computationally very thrifty, and with guaranteed (at least local) convergence around every trajectory. Moreover it behaves sensibly in the presence of acceleration and magnetic disturbances. Experimental comparisons with a commercial device illustrate its good performance; an implementation on an 8-bit microcontroller with very limited processing power demonstrates its computational simplicity.

As any other AHRS the proposed observer assumes the linear acceleration is small so that the accelerometers measurements are close to the gravity vector, which limits its use to "quasi-hover" situations. The relevance of this assumption in the context of a rotary wing UAV is discussed in Martin and Salaün (2010) and Pflimlin, Binetti, Souères, Hamel, and Trouchet (2010). When velocity measurements are available, observers based on the same approach can also be designed, see Martin and Salaün (2008c, 2008b).

The paper first presents the physical model used and proceeds with the construction of the observer. The choice of the tuning

* Corresponding author.
E-mail addresses: philippe.martin@mines-paristech.fr (P. Martin), erwan.salaun@gatech.edu (E. Salaün).

0967-0661/$ - see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.conengprac.2010.01.012

parameters, taking into account possible magnetic disturbance, and the ensuing convergence is then studied. Finally, experimental results on a very low-cost implementation are reported.

2. The physical system

2.1. Motion equations

The motion of a flying rigid body (assuming the Earth is flat and defines an inertial frame) is described by

$$\dot{q} = \tfrac{1}{2} q * \omega, \tag{1}$$

$$\dot{V} = A + q * a * q^{-1}, \tag{2}$$

where:

- q is the unit quaternion representing the orientation of the body-fixed frame with respect to the Earth-fixed frame;
- ω is the angular velocity vector expressed in the body-fixed frame;
- $A = ge_3$ is the (constant) gravity vector expressed in the Earth-fixed frame (the unit vectors e_1, e_2, e_3 point, respectively, North, East, Down);
- V is the velocity vector of the center of mass expressed in the Earth-fixed frame;
- a is the specific acceleration vector, in this case the aerodynamic forces divided by the body mass, expressed in the body-fixed frame.

The first equation describes the kinematics of the body, the second is Newton's force law. It is customary to use quaternions instead of Euler angles since they provide a global parametrization of the body orientation, and are well-suited for calculations and computer simulations. For more details see Stevens and Lewis (2003) or any other good textbook on aircraft modeling, and Appendix A for useful formulas used in this paper.

2.2. Measurements

In an AHRS there are no velocity measurements. Three triaxial sensors providing nine scalar measurements are used: three gyroscopes measure ω; three magnetometers measure the magnetic field in the body-fixed frame $y_B = q^{-1} * B * q$, where $B = B_1 e_1 + B_3 e_3$ is the Earth magnetic field in the Earth-fixed frame; three accelerometers measure a.

Clearly, the velocity V is not observable. A simple first-order analysis shows it is moreover not detectable. Indeed, linearizing (1)–(2) around the equilibrium point $(\overline{V}, \overline{q}, \overline{\omega}, \overline{a}) = (0, 1, 0, -A)$ yields

$$\delta\dot{q} = \tfrac{1}{2}\delta\omega$$

$$\delta\dot{V} = \delta a + 2A \times \delta q,$$

$$\delta y_B = 2B \times \delta q,$$

where $\delta q = \delta q_1 e_1 + \delta q_2 e_2 + \delta q_3 e_3$; but no observer

$$\begin{pmatrix} \delta\dot{\hat{q}} \\ \delta\dot{\hat{V}} \end{pmatrix} = \begin{pmatrix} \frac{1}{2}\delta\omega \\ \delta a + 2A \times \delta\hat{q} \end{pmatrix} + L(\delta y_B - \delta\hat{y}_B),$$

$$\delta\hat{y}_B = 2B \times \delta\hat{q},$$

where L is a freely chosen 6×3 matrix, is able to estimate δV: there will be a linearly growing error due to a double zero eigenvalue. The conclusion is the same when linearizing around any other equilibrium point.

For that reason, it is customary to assume the linear acceleration $\dot{V}$ small, hence to approximate the specific acceleration vector by $a = -q^{-1} * A * q$ using (2). This yields the new output $y_A = -a = q^{-1} * A * q$ (the sign is reversed for convenience).

Therefore the physical system (1)–(2) is seen as

$$\dot{q} = \tfrac{1}{2} q * \omega, \tag{3}$$

with output measurements

$$\begin{pmatrix} y_A \\ y_B \end{pmatrix} = \begin{pmatrix} q^{-1} * A * q \\ q^{-1} * B * q \end{pmatrix}. \tag{4}$$

2.3. Sensor imperfections

The sensors are of course not perfect, in particular they are usually biased. A reasonable assumption is to consider these biases constant but otherwise unknown. It would then be desirable to estimate them online together with the attitude and heading. While this is doable for the gyro biases, this is impossible for the accelero biases (though up to six unknown constants can be estimated since there are six output measurements). Indeed, assume the accelerometers measure in fact $a_m = a + a_b$, where a_b is a constant vector bias; (3)–(4) then becomes

$$\dot{q} = \tfrac{1}{2} q * \omega,$$

$$\dot{a}_b = 0,$$

$$\begin{pmatrix} y_A \\ y_B \end{pmatrix} = \begin{pmatrix} q^{-1} * A * q + a_b \\ q^{-1} * B * q \end{pmatrix}.$$

But this system is clearly unobservable: a first-order analysis as in the previous section reveals one combination of the components of a_b cannot be estimated. In a similar way, it is also impossible to completely estimate a bias vector on the magnetic measurements.

Another issue is the possible local perturbation of the magnetic field B. Once again a linear analysis shows it is not possible to estimate the three components of the magnetic field B (hence the perturbation), but only the North and Down components. Moreover only one imperfection on a_m can be estimated without relying on the possibly disturbed magnetic measurements. In an AHRS it is usually desirable to use the magnetic measurements to estimate only the heading, so that a magnetic disturbance does not affect the estimated attitude. As seen later this decoupling can be achieved by considering $y_C := y_A \times y_B = q^{-1} * C * q$, where $C := A \times B = gB_1 e_2$, rather than the direct measurement y_B. Notice that $\langle y_A, y_C \rangle = \langle A, C \rangle = 0$, hence only eight independent measurements out of nine are left; as a consequence only five unknown constants can now be estimated. This is not a drawback and is even beneficial since the observer will then not depend on the latitude-varying B_3.

Finally, the sensors are modeled as follows: the three gyros measure $\omega_m = \omega + \omega_b$, where ω_b is a constant vector bias; the three accelerometers measure $a_m = a_s a$, where $a_s > 0$ is a constant scaling factor; the three magnetometers measure $y_B = b_s q^{-1} * B * q$, where $b_s > 0$ is a constant scaling factor, which implies $y_C := c_s q^{-1} * C * q$, where $c_s := a_s b_s > 0$. There are therefore five unknown constants, which can all be estimated, see next section.

Noise also corrupts all the measurements; it is dealt with indirectly through the tuning of the observer gains.