UAV Attitude Estimation Using Low Frequency Radio Polarization Measurements

Sean T. G. Maguire and Paul A. Robertson

***Abstract*—A method of attitude determination, which makes use of measurements of the polarization of the magnetic field of Low Frequency (LF) radio signals, is presented and evaluated. This approach offers advantages relative to existing accelerometer-based systems in high-acceleration, cost-constrained environments such as small fixed-wing Unmanned Aerial Vehicles (UAVs). Flight test results are presented which demonstrate that LF polarization measurements can be used to obtain a significantly more accurate result than traditional approaches.**

***Index Terms*—Accelerometer, AHRS, attitude, low frequency radio, magnetic sensors, polarization, UAV.**

I. Introduction

UNMANNED Aerial Vehicles are being deployed in an increasingly wide range of activities. Small, low-cost UAVs in particular are ideally suited to a range of applications such as surveying [1], disaster relief [2], search and rescue [3] and precision agriculture [4].

One of the challenges of developing such UAVs is providing an accurate estimate of the aircraft's attitude (orientation) for control purposes. This function is typically performed by an Attitude and Heading Reference System (AHRS), which uses a filter to combine high-bandwidth gyroscope data with low-bandwidth vector measurements in order to produce an attitude estimate that offers both a high bandwidth and low drift. A three-axis magnetometer and three-axis accelerometer are generally used to provide the vector measurements by measuring the Earth's magnetic field and the acceleration due to gravity respectively.

A significant research effort has been devoted to the development of suitable filters for performing this sensor fusion in attitude estimation applications, including linear [5] and non-linear [6], [7] complementary filters, Extended Kalman Filters [8], [9], Unscented Kalman Filters [10] and Particle Filters [10], and to the development and application of suitable algorithms to produce the low-bandwidth vector estimate including Singular Value Decomposition (SVD) [11], QUEST [12], and others [13], [10].

However, all such AHRS architectures suffer from a limitation in high-acceleration environments, such as fixed-wing UAVs, where the accelerometer's estimate of the acceleration due to gravity is corrupted by the unknown dynamic acceleration due to aircraft motion. Methods have been proposed to reduce this effect, generally by applying corrections using information on platform motion provided by the Global Positioning System (GPS), such as velocity [14].

The authors have previously presented a sensor capable of making measurements of the axis of magnetic field polarization of Low Frequency (LF) radio signals, and shown that a number of existing European LF signals generally have a constant, known, and approximately linear polarization [15], [16]. Vector measurements of the polarization of such signals have further been shown to be suitable for use in attitude determination applications.

In this paper, the performance of attitude estimates based on LF polarization measurements is evaluated, both in static ground tests and in flight tests. In particular, it is shown that the limitation in the case of large and sustained dynamic accelerations due to motion can be eliminated by replacing the accelerometer used in a traditional AHRS with an LF polarization sensor.

Section II summarizes the limitations of using an accelerometer for attitude determination in high-acceleration environments; section III briefly presents the LF polarization sensor and outlines its proposed use for attitude determination; section IV outlines the measurement model used for the LF sensor, and the signal processing steps employed to estimate the axis of polarization and to perform attitude determination; and section V presents the results of ground and flight tests.

II. Low-Bandwidth Attitude Estimation

A platform's attitude is defined as the orientation of a coordinate frame fixed to the platform (the body-fixed frame) with respect to a reference coordinate frame (the Earth-fixed frame). Figure 1 defines the body-fixed frame and the Euler angles used throughout this paper. Note that the standard z-y-x rotation order convention is used, with intrinsic rotations.

Attitude can be estimated by integrating the angular velocity measured by a three-axis rate gyroscope (using Poisson's kinematic equation [17]). Equation (1) is a simple gyroscope model giving the measured angular velocity vector, $\boldsymbol{\omega}_m(t)$, in terms of the true angular velocity $\boldsymbol{\omega}(t)$, a gain matrix K_1 consisting of a scale factor for each axis, a time-varying bias $\mathbf{b}(t)$ and noise $\mathbf{n}_1(t)$ [18]. The change in the bias over time can be modeled as in (2) [8].

$$\boldsymbol{\omega}_m(t) = K_1\boldsymbol{\omega}(t) + \mathbf{b}(t) + \mathbf{n}_1(t) \tag{1}$$

Manuscript received May 4, 2015; revised October 25, 2015 and April 2 2016. This work was supported by an Engineering and Physical Sciences Research Council Doctoral Training Account award.

S. T. G. Maguire is with the Engineering Department, University of Cambridge, Cambridge, UK (e-mail: sean.maguire@cantab.net)

P. A. Robertson is with the Engineering Department, University of Cambridge, Cambridge, UK (e-mail: par10@cam.ac.uk)

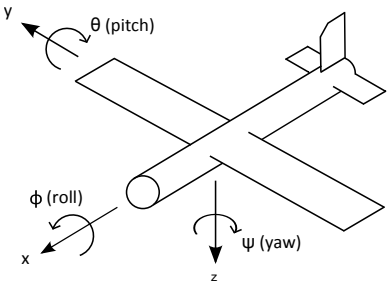

Fig. 1. Body-fixed coordinate frame and Euler angle conventions

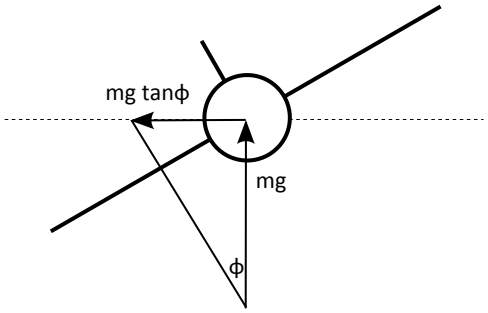

Fig. 2. Aircraft in coordinated turn showing forces corresponding to proper accelerations

$$\dot{\mathbf{b}}(t) = \mathbf{n}_b(t) \tag{2}$$

Drift of the bias $\mathbf{b}(t)$ over time is generally corrected by combining a gyroscope-based attitude estimate with an alternative low-bandwidth estimate that exhibits no drift [18]. This is commonly a vector attitude solution based on solving Wahba's problem [19], [20], [21], using the Earth's magnetic field and the acceleration due to gravity as reference vectors, and taking measurements of those vectors in the body-fixed frame using a magnetometer and an accelerometer respectively.

Alternatively, simple estimates of roll ϕ and pitch θ are given by (3) and (4) respectively, given an accelerometer measurement in the body-fixed frame assumed to correspond to the acceleration due to gravity, $\mathbf{g}_{bff} = \begin{bmatrix} g_x & g_y & g_z \end{bmatrix}^T$. Similarly, the heading can be estimated by using the magnetometer as a tilt-compensated compass.

$$\tan\hat{\phi} = \frac{g_y}{g_z} \tag{3}$$

$$\tan\hat{\theta} = \frac{-g_x}{\sqrt{g_y^2 + g_z^2}} \tag{4}$$

Fusion of these low- and high-bandwidth estimates can be achieved using, for example, a complementary filter (in the simplest case) [5], [18] or, more commonly, an Extended Kalman Filter (EKF) [8], [9].

However, there is a well-known limitation of relying on accelerometer measurements in this way. Accelerometers measure specific force (acceleration relative to free fall) as given by (5), where $\mathbf{a}_m(t)$ is the measured acceleration, $\mathbf{a}_{dynamic}(t)$ is the dynamic acceleration due to motion, $\mathbf{g}$ is the acceleration due to gravity, K_2 is a gain matrix consisting of a scale factor for each axis, and $\mathbf{n}_2(t)$ is a noise term.

$$\mathbf{a}_m(t) = K_2(\mathbf{a}_{dynamic}(t) + \mathbf{g}) + \mathbf{n}_2(t) \tag{5}$$

Attitude determination using accelerometer measurements relies on the assumption that $\mathbf{a}_{dynamic}(t) \approx 0$. When the accelerometer data is used as a low-bandwidth estimate, large dynamic accelerations can be tolerated for short periods. However, if large and sustained dynamic accelerations are present then large errors will be introduced into the attitude estimate.

In the case of small fixed-wing UAVs, large dynamic accelerations due to motion would be expected during flight, and could persist for long periods of time. For example, Fig. 2 shows an aircraft in a coordinated turn with a roll angle of $-\phi$, including the non-gravitational forces. Here the resultant lift force acts perpendicular to the wings, with magnitude $mg\sqrt{1+\tan^2(\phi)}$. It is the acceleration due to this lift force that is measured by the accelerometer, and therefore the orientation of the measured acceleration does not change, relative to the aircraft, in a coordinated turn. This would lead to a roll estimate of 0°, based on (3), in an arbitrarily steep coordinated turn. Such a turn could persist for an arbitrarily long time, for example in a loitering UAV.

The use of GPS corrections has been suggested to reduce this effect [14]. For example, the aircraft's coordinate acceleration can be estimated from GPS velocity measurements, and can in turn be used to apply centripetal acceleration corrections to an AHRS estimate [14] or indeed to directly estimate "pseudo-attitude" [22], [23], [24], [25]. However, GPS can only provide an indirect attitude input, generally requiring assumptions about the aircraft's flight path, flight characteristics and the wind speed to be made. There are thus advantages to an approach that replaces measurements of acceleration with measurements of an alternative vector that is not sensitive to dynamic acceleration and does not require GPS corrections.

III. Low Frequency Radio Polarization and Attitude

One example of the use of polarization measurements for attitude determination has been reported in the literature [26]. In reference [26], it was demonstrated that the attitude of a satellite can be partially determined by taking measurements of the polarisation of Radio Frequency (RF) signals broadcast by the satellite and received by ground stations. In particular, given a satellite broadcasting linearly polarised RF signals from an omnidirectional antenna, the antenna orientation can be determined using 2D measurements of electric (E-) field polarization made at two geographically separated ground stations. This partially determines the satellite attitude, with the rotation about the axis of polarisation being unobservable. It was further shown that measurements at only one ground station were sufficient to provide an estimate of yaw to within approximately 2° in practice.

A related method is proposed here for small UAV applications. If a receiver carried by a UAV is capable of fully determining the polarization of one or more radio signals broadcast by ground stations, each having a known linear polarization, then each measurement provides one vector (subject to a sign ambiguity) that can be used as an input to an attitude determination algorithm (without introducing acceleration-dependence). The signals from multiple geographically separated transmitters can be measured to produce multiple vector measurements. A full, independent attitude solution could be achieved as long as at least two non-collinear vectors are available. Additional signals could be used to improve accuracy and introduce redundancy. All measurement and calculation can also be performed on board the UAV, avoiding the need for communication between the ground stations and the vehicle. In addition, in many regions suitable LF signals are already used for other applications including navigation, timing and AM radio broadcasts, and are therefore available for use in attitude determination applications without the need to construct additional transmitters.

The authors have previously reported the development of an LF radio polarization sensor and proposed its use in attitude determination systems in high-acceleration environments [15], [16]. The sensor comprises an array of three orthogonal air-cored loop antennas, three radio receiver circuits, and a data logging and control module. This sensor was shown to be capable of measuring the magnetic (H-) field due to LF radio signals in three dimensions, and determining the corresponding H-field polarization. A number of LF signals used for AM radio broadcasts in Europe were shown to have a magnetic field with an approximately linear polarization, with a known orientation in the horizontal plane. This orientation was further shown to exhibit little variation over time. Measurements of these vectors can therefore be used to give a low-bandwidth attitude estimate, either alone or in combination with other sensors such as magnetometers. Simple heading estimation using a sensor in a fixed location was demonstrated with an accuracy of $\pm 1.6^\circ$ using an LF signal measured at a range of approximately 150 km from the transmitter and a fully-calibrated sensor, with reduced accuracy (as low as $\pm 8.2^\circ$) in non-ideal cases.

In addition, LF polarization measurements were found to be subject to distortion when large conductive loops were present near the receiver. In a small Remotely Piloted Aircraft (RPA) with a non-metal frame, operating outdoors, this effect is not prohibitive. However, some distortion was predicted (see section IV) and observed experimentally (see section V), and is likely to be dependent on the design and layout of the RPA.

The measured H-field polarization vectors, which lie in the horizontal plane, perpendicular to a line between the transmitter and receiver, have an unavoidable sign ambiguity. This can be resolved by taking measurements of the electric field as well as the magnetic, as is traditionally done in the field of Radio Direction Finding (RDF), or by comparing the polarization vectors with a known reference such as the Earth's magnetic field vector (assuming they are not close to orthogonal).

IV. Measurement Model and Signal Processing

The authors have previously presented a measurement model for the LF sensor of section III [15], [16]. A brief discussion is repeated here, and extended to show the additional processing steps used in the AHRS implementation.

The sensor includes three orthogonal antennas, whose outputs are amplified and band-pass filtered in hardware to produce time-varying voltages $\mathbf{V}_m(t) = \left[V_x(t)\, V_y(t)\, V_z(t)\right]^T$, which are related to the vector magnetic field $\mathbf{h}(t)$ of the LF signal of interest by (6). Here s is a scale factor converting from magnetic field strength to voltage, K is a gain matrix containing a scale factor for each axis (k_{11}, k_{22} and k_{33}) and cross-coupling terms, and $\mathbf{n}(t) = \left[n_x(t)\, n_y(t)\, n_z(t)\right]^T$ is Gaussian noise with a spectrum shaped by the hardware's frequency response.

$$\mathbf{V}_m(t) = sK\mathbf{h}(t) + \mathbf{n}(t) \tag{6}$$

The analogue voltage outputs from each receiver circuit ($V_x(t), V_y(t)$ and $V_z(t)$) are sampled simultaneously by a set of three parallel ADCs at regular time steps $t = k\tau$ ($k = 1, 2 \ldots, n$), resulting in a series of discrete samples given by the vectors $\mathbf{V}_x$, $\mathbf{V}_y$ and $\mathbf{V}_z$. $\mathbf{V}_x$ is shown in (7). $\mathbf{V}_y$ and $\mathbf{V}_z$ are similarly constructed. These samples are stored by the datalogging and control module and subsequently processed off-line in MATLAB to perform filtering, polarization extraction and attitude determination.

$$\mathbf{V}_x = \left[V_x(0)\, V_x(\tau)\, V_x(2\tau)\, \ldots\, V_x(n\tau)\right]^T \tag{7}$$

A second LF signal was measured by using additional parallel hardware filters and ADCs. In practice digitally filtering wideband signals or time-multiplexing the receiver's tuned frequency would allow reduced hardware complexity. The two LF signals measured were BBC Radio 4 longwave, transmitted from Droitwich, UK at 198 kHz with a power of 500 kW, and France Inter, broadcast from Allouis, France at 162 kHz with a power of 2 MW. The comparable Europe 1 station was not used due to poor observed polarization stability over time. This resulted in six digitized time-series $\mathbf{V}_x|_{198}$, $\mathbf{V}_y|_{198}$, $\mathbf{V}_z|_{198}$, $\mathbf{V}_x|_{162}$, $\mathbf{V}_y|_{162}$, and $\mathbf{V}_z|_{162}$.

Figure 3 gives an overview of the digital processing applied to these signals. First, each of the six time series is filtered with a digital elliptic Band-Pass Filter, centered on the appropriate under-sampled carrier frequency with a 100 Hz bandwidth, in order to isolate the carrier. Note that since the signal of interest is at a single frequency in each case, and the same filter is applied to the signal from each axis, passband variation in the filter response can be tolerated. For each LF signal, Singular Value Decomposition (SVD) is then performed on a matrix containing a batch of three-axis filtered data in order to estimate (subject to a sign ambiguity) the vectors $\mathbf{v}_i$ corresponding to the three axes ($i = 1, 2, 3$) of an ellipsoid fitted to the polarisation ellipse of each signal [16]. This results in six estimated vectors ($\hat{\mathbf{v}}_i|_{198}$ and $\hat{\mathbf{v}}_i|_{162}$). It is primarily the major axis estimates $\hat{\mathbf{v}}_1$ that are of interest.

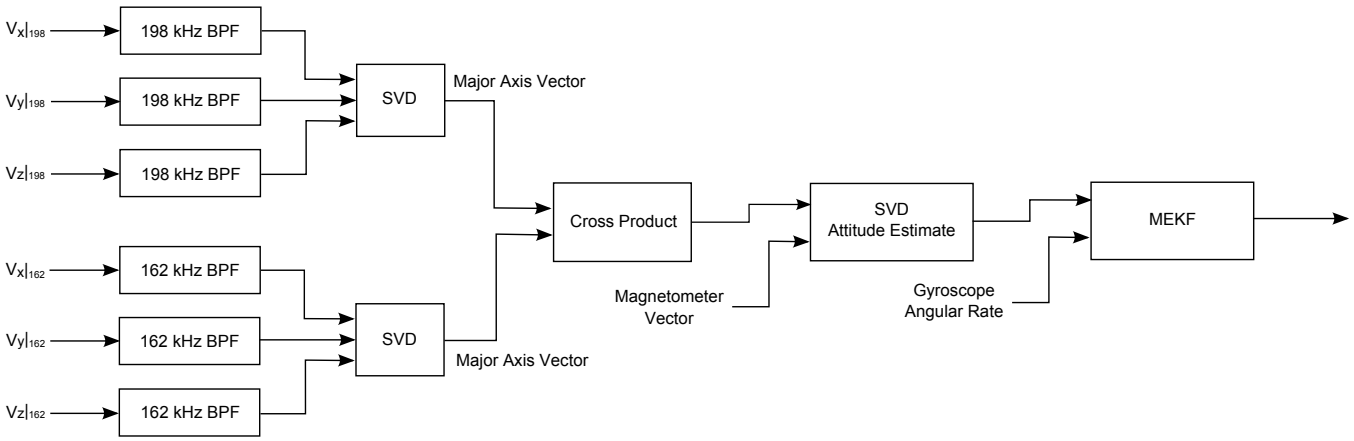

Fig. 3. Overview of digital processing used to produce Multiplicative Extended Kalman Filter (MEKF) attitude estimate, including band-pass filtering, polarisation determination, vertical vector synthesis, vector attitude estimation and fusion of low- and high-bandwidth sensor data in the MEKF

The sensor is calibrated to ensure each axis has equal gain and phase shift, so the diagonal terms of the sensor gain matrix K in (6) are equal at the frequency of interest. If the total gain and phase shift of the system is incorporated in the scaling term s, it follows that $k_{11} = k_{22} = k_{33} = 1$. It will be assumed that the off-diagonal cross-coupling terms are zero ($k_{ij} = 0, i \neq j$), so that K is equal to the identity matrix. Under these assumptions, the sensor does not distort the polarization of the signal being measured. Errors due to the cross-coupling terms are addressed later in section IV.

The vertical vector of (8) can be synthesized as the cross product of the estimated major axes of polarization of the two signals of interest, given that both vectors are assumed to lie in the horizontal plane. The sign ambiguity is resolved at this stage by multiplying any estimates of $\hat{\mathbf{v}}_v$ with a negative z-component by -1 (ie assuming the aircraft is not upside down). Section III suggests more satisfactory approaches for properly resolving the sign ambiguities in the underlying measurements.

$$\hat{\mathbf{v}}_v = \hat{\mathbf{v}}_1|_{198} \times \hat{\mathbf{v}}_1|_{162} \tag{8}$$

In addition to permitting a simple (if restrictive) resolution of the sign ambiguities, synthesis of a vertical vector has the benefit of producing a vector with a known orientation without knowledge of the relative locations of the transmitters and the receiver. It is also directly comparable with the vertical acceleration due to gravity, and permits simple trigonometric calculations of pitch and roll to be performed. It discards the yaw information present in the polarization measurements, but this may be acceptable in many applications given that yaw can generally be measured easily using a magnetometer.

Simple trigonometric estimates of pitch and roll can be formed based on this estimate of $\hat{\mathbf{v}}_v = \begin{bmatrix} v_{vx} \; v_{vy} \; v_{vz} \end{bmatrix}^T$ in the body-fixed frame using (9) and (10). This is analogous to using measurements of the acceleration due to gravity in (3) and (4).

$$\tan \hat{\phi} = \frac{v_{vy}}{v_{vz}} \tag{9}$$

$$\tan \hat{\theta} = \frac{-v_{vx}}{\sqrt{v_{vy}^2 + v_{vz}^2}} \tag{10}$$

Singular Value Decomposition (SVD) was used to calculate a full 3D attitude solution [11], represented as a rotation matrix or Direction Cosine Matrix, based on measurements of the Earth's magnetic field and the vertical vector $\hat{\mathbf{v}}_v$. These vectors are measured (or estimated) in the body-fixed frame and are known in the Earth-fixed frame. The LF "vertical vector" could be assumed to be vertical in the reference frame. However, improved accuracy was achieved by measuring the true value of this reference vector on the ground prior to flight. An equivalent SVD attitude estimate was formed using traditional accelerometer and magnetometer measurements for comparison.

In order to determine the weighting factor applied to each sensor, and to estimate the covariance of the attitude estimate, the SVD approach takes as inputs the variances of the two vectors measured in the body-fixed frame $\sigma_{bi}, (i = 1, 2)$, and the variances of the corresponding reference vectors $\sigma_{ri}, (i = 1, 2)$. σ_{bi} thus corresponds to errors in the measurements, and σ_{ri} corresponds to errors in the model for the orientation of the reference vectors. For each sensor, σ_{bi} was given by the combined variance of the elements of a unit vector aligned with the sensor's three-axis measurement vector during straight and level flight. σ_{ri} was given by the variance of long-term measurements of the reference vectors made on the ground with a stationary sensor. Some additional reference vector noise was added to account for errors in the model, including sensor alignment and, for the accelerometer, the presence of dynamic accelerations.

A Multiplicative Extended Kalman Filter (MEKF), based on reference [9], was implemented in order to fuse the vector attitude estimates obtained using SVD with the gyroscope data, both for the novel LF-based estimate and the traditional estimate. The filter's six-element state vector is given by (11). Here $\mathbf{a}$ is the Gibbs vector representation of the attitude error relative to the reference quaternion, and $\mathbf{b}$ is a vector of gyroscope biases. The Gibbs vector representation was cho-

sen, somewhat unusually, following reference [9]. This filter operates on a minimal three-parameter representation of the attitude error relative to a reference quaternion. The attitude error is assumed to be small and therefore non-singular, while the reference quaternion provides a non-singular representation of the global attitude. Using the Gibbs vector as the three-parameter representation permits the propagation of a normalized reference quaternion without the accumulation of numerical errors [9].

$$\mathbf{x} = \begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix} \tag{11}$$

The process noise covariance matrix Q is given by (12), (13) and (14), where σ_g is the standard deviation of the gyroscope noise $\mathbf{n}_1$ of (1) and σ_b is the standard deviation of the gyroscope bias noise $\mathbf{n}_b$ of (2) [9]. σ_g was estimated from a sample of gyroscope data during straight and level flight of the RPA. σ_b was estimated as the minimum Allan variance, measured on each axis of the gyroscope when stationary, as a function of Allan variance time period τ. Some additional process noise was added to account for error sources which are not modeled.

$$Q = \begin{bmatrix} Q_1 & 0 \\ 0 & Q_b \end{bmatrix} \tag{12}$$

$$Q_1 = \begin{bmatrix} \sigma_{gx} & 0 & 0 \\ 0 & \sigma_{gy} & 0 \\ 0 & 0 & \sigma_{gz} \end{bmatrix} \tag{13}$$

$$Q_b = \begin{bmatrix} \sigma_{bx} & 0 & 0 \\ 0 & \sigma_{by} & 0 \\ 0 & 0 & \sigma_{bz} \end{bmatrix} \tag{14}$$

At each time step, the measurement noise covariance matrix R is given by the SVD attitude error covariance matrix P in the body-fixed frame [9], [11].

There are two main sources of error in the LF polarization estimates and the resulting attitude estimates. Firstly, there are measurement errors due to the gain matrix K of (6). Although the diagonal terms are accurately calibrated at the frequency of interest, small unknown phase shifts can be introduced at other frequencies. Given that the two LF signals (at 198 kHz and 162 kHz) were recorded simultaneously, and phase calibration was only performed at 198 kHz, phase shift errors are present in the 162 kHz measurements.

Similarly, errors are introduced by the presence of cross-coupling terms, which are non-zero in practice. With careful sensor design, it was found that $|k_{ij}| < 0.05, i \neq j$. These errors could be further reduced by performing an extra calibration step to estimate the off-diagonal terms of the gain matrix K. However, this is likely to be impractical for most applications, as it would most likely require calibration measurements to be made of accurate reference signals in a controlled environment, and the calibration would need to be repeated periodically to account for drift. The actual matrix K could also vary due to changes in the nearby environment of the sensor.

The second major source of error is introduced by the reference vectors. The actual orientation of the major axis of the LF signals may deviate from the horizontal by a few degrees, and drifts over time [15], [16]. There is clearly a resultant deviation of the synthesized LF vertical vector $\hat{\mathbf{v}}_v$ from the true vertical. A calibration step is performed, in the case of the SVD estimate, to account for this deviation by measuring the actual value of $\hat{\mathbf{v}}_v$ on the ground prior to flight. However, errors could be reduced by extending the filter to track changes in the reference vector $\hat{\mathbf{v}}_v$ over time.

There are also errors in the LF estimates due to wideband noise. However, these are much less problematic as the LF estimate is intended only as a low-bandwidth estimate, ensuring most of the noise is filtered out in the sensor fusion process.

V. Experimental Results

Ground and flight tests were performed in order to demonstrate the feasibility of attitude estimation using LF polarization measurements, and to evaluate the performance of this approach relative to traditional methods. All measurements were made in Cambridge, UK at a range of approximately 150 km from the BBC Radio 4 transmitter and 600 km from the France Inter transmitter.

A. Ground Tests

The LF H-field polarization sensor outlined in sections III and IV was used to gather data in a number of ground tests. It was mounted to a platform alongside an ADXL345 three-axis MEMS accelerometer and an HMC5883L three-axis magnetometer. Attitude estimates were calculated in post-processing. Fig. 4 shows the roll angle estimated during a series of step changes in roll. Estimates were made using the simple trigonometric relation of (3) operating on accelerometer measurements, and similarly using (9) operating on the synthesized vertical vector of section IV. Given dynamic accelerations were negligible, the accelerometer provides an accurate reference roll angle against which the LF estimate can be compared. It can be seen that the LF estimate agrees reasonably closely with the accelerometer reference, but that there are sustained steady-state errors of up to 10° in some orientations. The LF estimate also has significantly higher noise than the accelerometer estimate. The steady-state errors are due both to measurement errors and to errors in the reference vector orientations. As discussed in section IV, these can be partially mitigated by calibration and design choices.

In order to assess heading (yaw) estimation, the platform was rotated, while level, in an outdoor location. Fig. 5 shows the yaw angle of the platform estimated using the HMC5883L magnetometer as a compass, and compares it against an estimate of yaw obtained using LF H-field polarization measurements. For a horizontal sensor measuring a signal with a known transmitter location, the measured angle of H-field polarization in the horizontal plane is offset from the known bearing to the transmitter by 90°, and therefore allows the yaw angle to be simply estimated (subject to a sign ambiguity). It

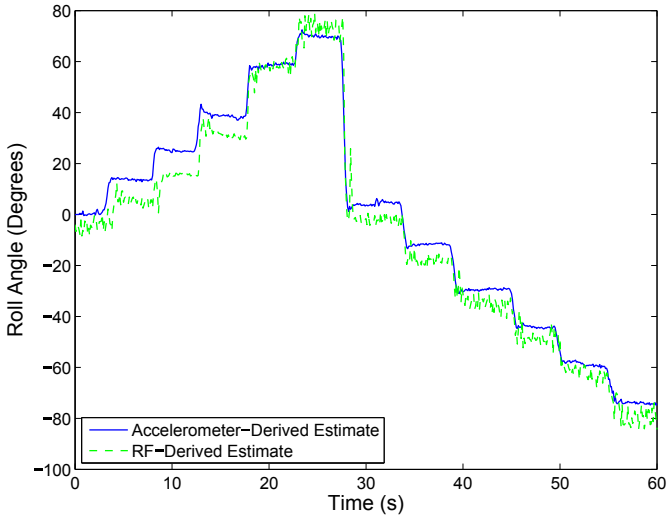

Fig. 4. Roll angle estimated from LF polarization of 198 kHz Radio 4 and 162 kHz France Inter stations, with accelerometer reference angle, with sensor on stationary platform at various roll angles

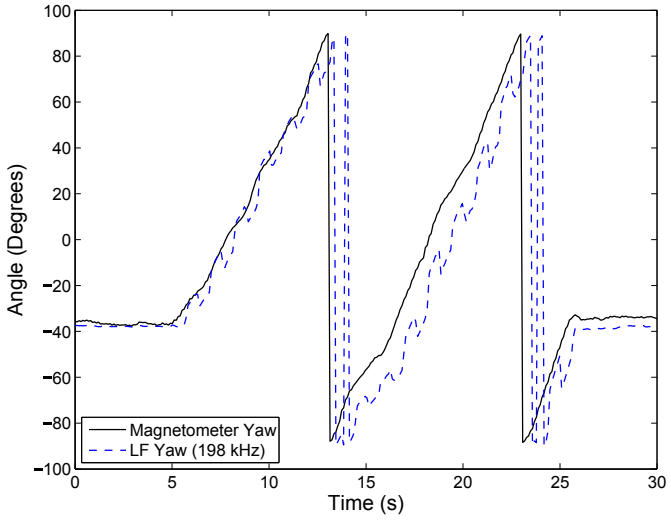

Fig. 5. Yaw angle estimated from LF polarization of 198 kHz Radio 4 carrier signal, with magnetometer-based reference yaw angle

can be seen that the LF estimate provides a reasonable yaw estimate, although with sometimes significant errors relative to the magnetometer reference, and with periodic interference present.

B. Flight Tests

Flight tests were conducted to verify the sensor's behavior in a small Remotely Piloted Aircraft (RPA). A Dynam Hawk Sky airframe was used to carry a payload including the LF radio polarization sensor, an ADXL345 three-axis MEMS accelerometer, an HMC5883L three-axis magnetometer, an ITG-3200 three-axis MEMS gyroscope, a GPS module, a data logger and a forward-looking camera. Attitude estimates were calculated offline in post-processing. The ground truth pitch and roll angles were determined using an open-source horizon-tracking algorithm implemented in MATLAB [27], operating on the data from the forward-looking 20 frame per second camera, which was first calibrated using the open-source Camera Calibration Toolbox for MATLAB [28].

Some interference from the RPA's electrical systems was observed in the LF polarization measurements. It was found that by careful layout of the aircraft's systems this effect could be minimized, such that useful polarization estimates could be obtained after filtering. However, the results of this section were obtained during glide phases of flight, when the motor was off, for clarity and simplicity of analysis. Since it is the low-bandwidth performance of the sensor that contributes to overall AHRS performance, neglecting the high frequency interference in this way was not considered a significant limitation.

Three flight segments were analyzed. Segment one includes a period of straight and level flight, followed by a pair of steep turns; segment two includes a short series of small pitch variations due to longitudinal oscillations; and segment three includes a series of small variations in attitude (primarily roll) on approach to landing. These flight segments allow the performance of the attitude estimates to be evaluated during a number of standard flight maneuvers. However, it is the performance in the steep turns that is most critical, given the significant limitations of existing systems in that case.

Fig. 6 shows a series of pictures from the forward-looking camera during flight segment one. It can be seen that the camera images contain a good horizon reference throughout most of this part of the flight, although it is not visible for occasional brief periods (such as $t \approx 17$ s). The reference angle is plotted as zero during these periods. Some views also include linear features that produce spurious horizon detections in intermittent frames, which introduce noise into the computer vision attitude estimate. However, this effect is small, and is only present during a few brief intervals.

Fig. 7 shows simple trigonometric estimates of roll during flight segment one, alongside the computer vision reference angle. A traditional accelerometer estimate is obtained using (3), and a comparable estimate based on LF measurements is obtained using (9). The accelerometer clearly provides no information about roll angle during the turns, as would be expected from the analysis of section II. However, the LF estimate agrees closely with the reference, subject to a small offset during straight and level flight and with increased wideband noise. As discussed in section IV, these are not significant limitations.

Singular Value Decomposition (SVD) was used to calculate full 3D attitude solutions, as detailed in section IV. An LF-based estimated was obtained using the synthesized LF vertical vector and magnetometer measurements, while a traditional estimate based on accelerometer and gyroscope measurements was also produced for comparison.

LF polarization measurements were made on the ground prior to the flight test to allow the true orientation of the synthesized vertical vector to be calculated. It was found to deviate from the vertical by 11.5°, and this measured value was used as the reference vector in the SVD attitude solution.

Figs. 8 and 9 show the resulting roll angle estimates during flight segment one. Fig. 10 plots the errors relative to the

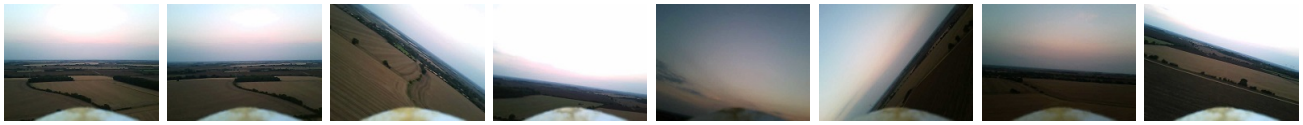

(a) $t = 1$ s (b) $t = 7$ s (c) $t = 13$ s (d) $t = 16$ s (e) $t = 17$ s (f) $t = 23$ s (g) $t = 27$ s (h) $t = 32$ s

Fig. 6. Images from forward-looking camera during flight segment one

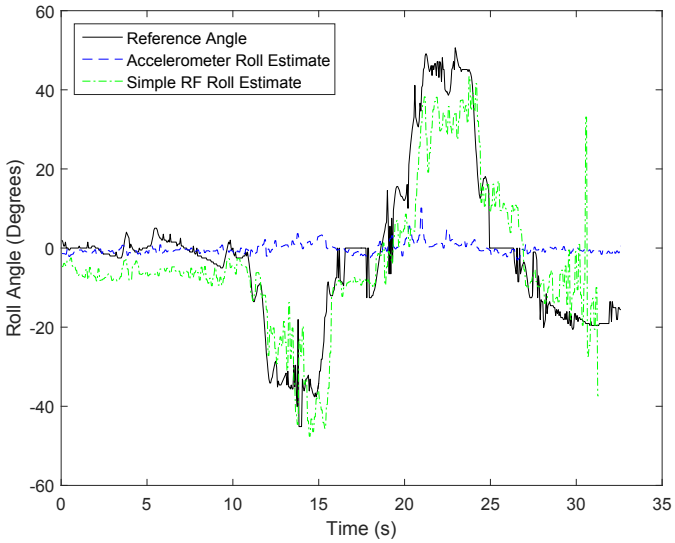

Fig. 7. RPA roll angle estimated from accelerometers alone and from LF alone, compared against computer vision reference, for flight segment one

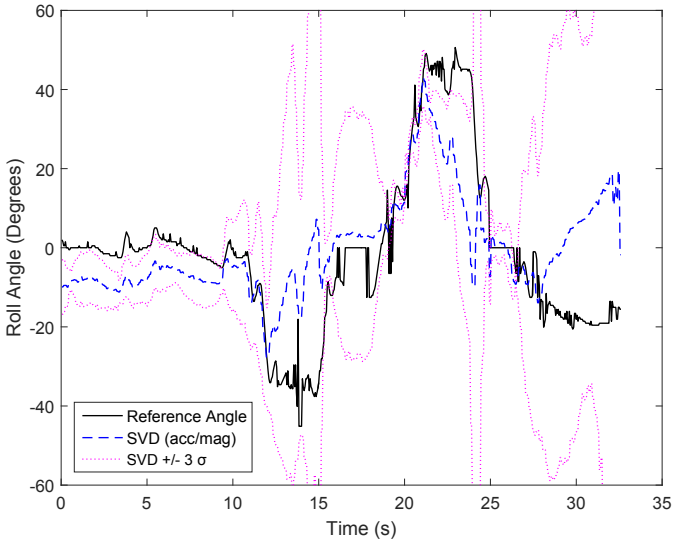

Fig. 8. Traditional SVD estimate of roll angle (based on accelerometer and magnetometer), alongside computer vision reference, for flight segment one

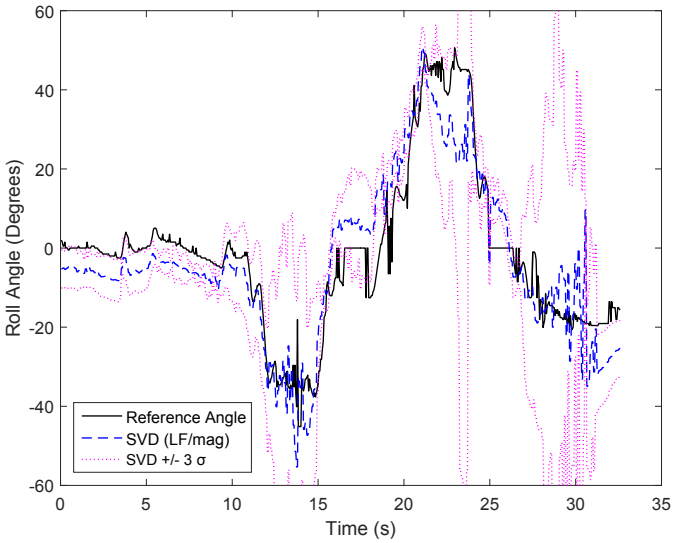

Fig. 9. LF SVD estimate of roll angle (based on LF vertical vector and magnetometer), alongside computer vision reference, for flight segment one

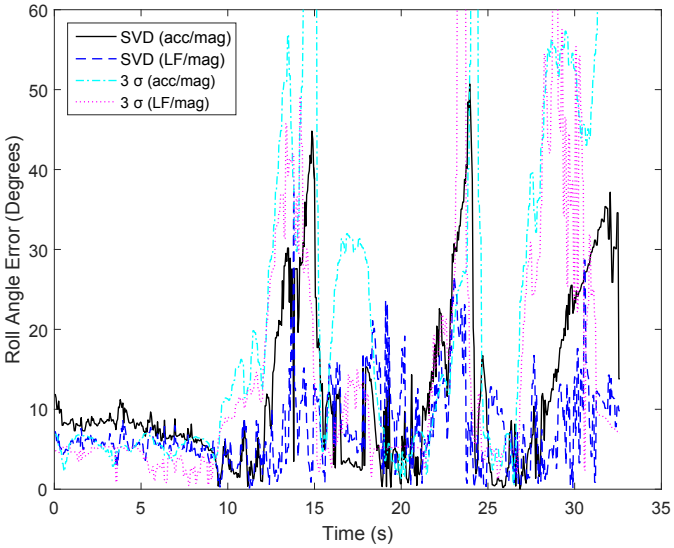

Fig. 10. Absolute errors in the traditional and LF SVD roll estimates for flight segment one

computer vision reference angle for comparison. It can be seen that the traditional accelerometer and magnetometer estimate has some ability to detect roll angle during the turns, due to the inclusion of magnetometer measurements, but that replacing the accelerometer measurements with LF polarization measurements leads to a significant improvement in accuracy.

Fig. 11 shows the errors in the traditional and LF-based SVD estimates of pitch during flight segment 2. The LF estimate is significantly more accurate than the traditional accelerometer estimate, which largely fails to measure the short-period longitudinal oscillations present during this flight segment. This is again likely to be due to the presence of dynamic accelerations. Although these errors are not as likely to be sustained over long periods of time as they are in the

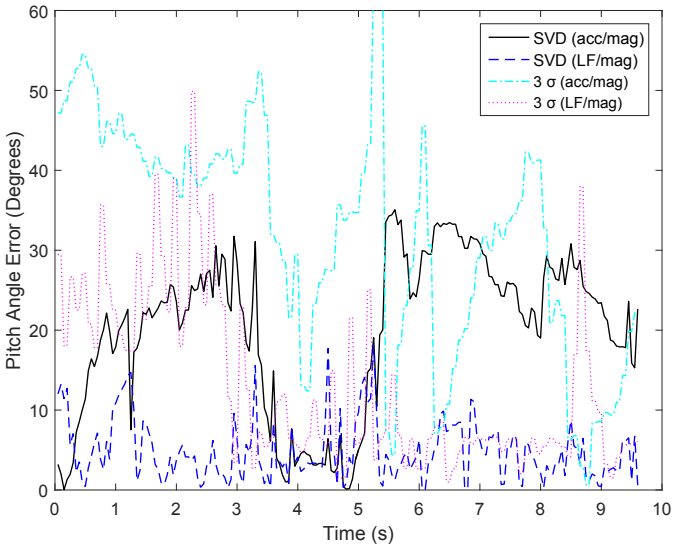

Fig. 11. Absolute errors in the traditional and LF SVD pitch estimates for flight segment two

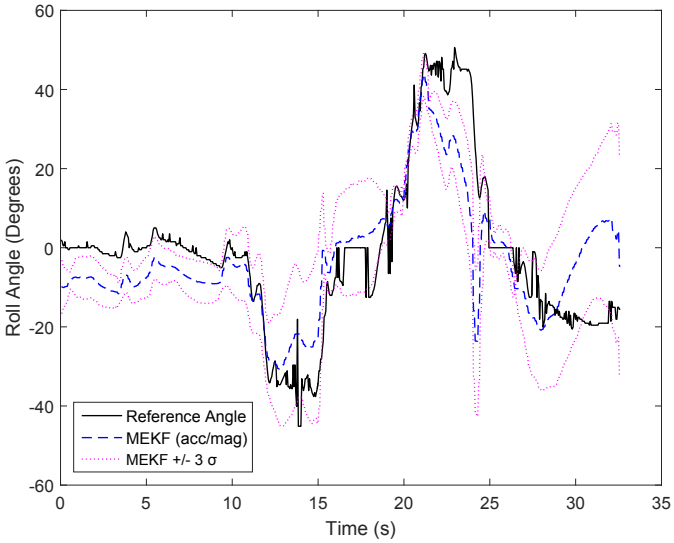

Fig. 13. Traditional MEKF estimate of roll angle (based on accelerometer, magnetometer, and gyroscope), alongside computer vision reference, for flight segment one

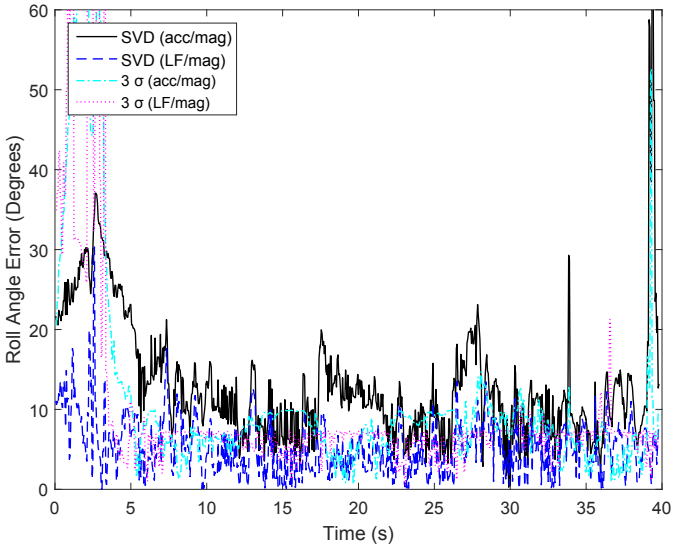

Fig. 12. Absolute errors in the traditional and LF SVD roll estimates for flight segment three

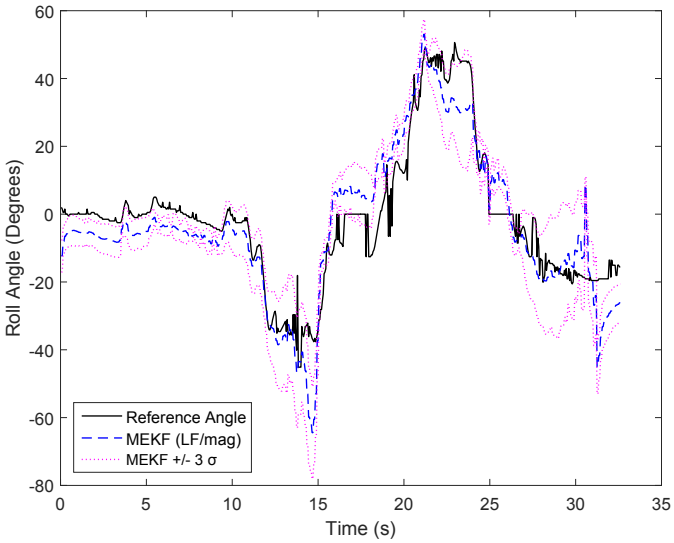

Fig. 14. LF MEKF estimate of roll angle (based on LF vertical vector, magnetometer, and gyroscope), alongside computer vision reference, for flight segment one

roll case, these results demonstrate that the use of LF-based estimates can improve pitch, as well as roll, estimates.

Fig. 12 shows the errors in the traditional and LF-based SVD estimates of roll during flight segment three. This segment covers an approach to landing, during which small changes in roll are observed. Here there is a smaller performance difference between the estimates, given that the magnitude of the dynamic accelerations is lower, but the LF-based estimate is still superior to the accelerometer-based estimate.

Figs. 13 and 14 show the roll estimates, for flight segment one, obtained by fusing the traditional and LF-based SVD estimates respectively with gyroscope data in an MEKF, as detailed in section IV. Fig. 15 shows the errors relative to the computer vision reference for each estimate. As would be expected, the use of gyroscope data significantly improves the traditional estimate in the case of steep turns, although this improvement can only be achieved over a limited bandwidth - that is, for short time periods over which the gyroscope is reliable. It also reduces the noise present in the LF-based estimate, improving its high-bandwidth performance.

Table I summarizes the RMS errors for each roll estimate considered during flight segment one. Errors are calculated relative to the computer vision reference for the whole representative flight segment being analyzed, as well as for the parts of that segment corresponding to straight and level flight ($t < 10$ s) and to steep turns (10 s $< t < 27.5$ s). Short periods for which the computer vision reference is not available are ignored. It is clear that:

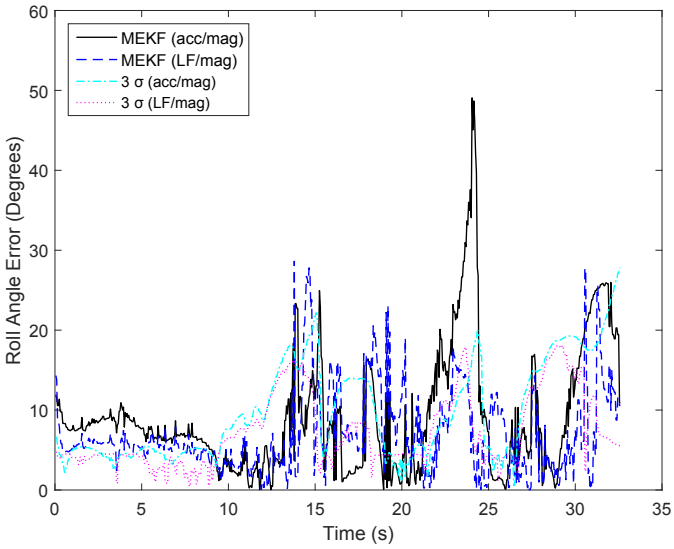

Fig. 15. Absolute errors in the traditional and LF MEKF roll estimates for flight segment one

TABLE I
RMS ROLL ERRORS FOR ATTITUDE ESTIMATES DURING RPA FLIGHT SEGMENT ONE

	Flight Segment 1		
	Straight and Level Flight	**Steep Turns**	**Overall**
Simple Accelerometer Roll	2.4°	28.0°	21.7°
Simple LF Roll	6.3°	11.3°	10.9°
SVD	7.4°	20.6°	17.2°
SVD with LF	5.5°	9.9°	9.5°
MEKF	7.4°	17.5°	12.9°
MEKF with LF	5.5°	9.4°	9.1°

- The simple accelerometer-based estimate of roll performs extremely badly in the turns, and the simple LF-based estimate offers a significant improvement.
- The traditional SVD estimate offers a slight improvement over the accelerometer-only estimate by also including the more reliable magnetometer data.
- The LF SVD estimate offers an improvement relative to the LF only estimate, primarily because the reference vector calibration used in the SVD case reduces the effect of distortion of the LF signals.
- The gyroscope-based estimates perform well over this time scale, and are consequently able to reduce the errors in the SVD estimates.

Table II summarizes the RMS errors in the traditional and LF-based SVD estimates of pitch and roll in each of the three flight segments studied. In all cases the LF estimate provides improved accuracy, and in most cases the improvement is significant. The RMS error values in both estimates depend on the performance of the sensors and the particular filter implementations used, but the general trend is clear.

TABLE II
RMS PITCH AND ROLL ERRORS FOR ATTITUDE ESTIMATES DURING RPA FLIGHT SEGMENTS ONE, TWO AND THREE

	Segment 1	**Segment 2**	**Segment 3**
SVD (roll)	17.2°	12.3°	15.4°
SVD with LF (roll)	9.5°	6.9°	6.5°
SVD (pitch)	16.5°	22.2°	17.5°
SVD with LF (pitch)	13.0°	6.2°	9.2°

These results demonstrate that a conventional low-cost AHRS, which relies heavily on the accelerometer to correct for drift in the gyroscope bias, encounters large errors when dynamic accelerations are present for a long period of time. It is the rate of drift in the gyroscope bias that determines the length of time over which a valid solution can be obtained even in the absence of a good low-bandwidth estimate. Good system performance can therefore only be obtained using a high-quality gyroscope, and the system necessarily has a maximum turn duration, above which gyroscope drift errors prevent an accurate roll estimate from being obtained. This maximum turn duration places limits on the operation of the UAV, for example potentially restricting loiter time and mission duration.

By obtaining an improved low-bandwidth, low-drift attitude estimate from LF polarization measurements, good performance can be obtained despite dynamic accelerations being sustained for an arbitrarily long time. Furthermore, good overall system performance could potentially be obtained using lower cost gyroscopes than in a traditional AHRS.

VI. CONCLUSION

In this paper, the principle of operation of low-cost Attitude and Heading Reference Systems based on an accelerometer, magnetometer and gyroscope has been reviewed, in order to highlight the limited low-bandwidth accuracy of this approach in the case of high and sustained dynamic accelerations.

The use of measurements of the H-field polarization of LF radio signals in order to determine attitude was outlined, and shown to offer the potential to avoid the dependence of the AHRS output on dynamic accelerations.

A measurement model for a suitable LF polarization sensor was presented, along with the signal processing and calibration steps used to obtain LF-based attitude estimates.

Results obtained in flight tests verified that accelerometer-based roll estimates perform poorly, particularly in turns, as would be expected from simple analysis. Conversely, roll estimates based on LF polarization measurements were found to perform well in a small RPA carrying out turns and other maneuvers, offering a significantly more accurate solution.

It was further shown that good performance can be obtained by combining a traditional low-bandwidth roll estimate with a high-bandwidth gyroscope estimate using an MEKF. However, this improvement is the result of including the gyroscope data,

and therefore can only be obtained over a limited bandwidth. Applying the same approach to the LF estimates resulted in slightly improved performance.

Replacing the accelerometer used in a traditional AHRS with an LF polarization sensor resulted in a significantly improved low-bandwidth roll estimate, eliminating errors due to the presence of high and sustained dynamic accelerations. It may therefore also permit a good attitude estimate to be obtained despite the use of a low-cost gyroscope with poor bias stability. Significant cost savings could therefore be possible.

Given that suitable LF signals are already broadcast in many regions, LF-based attitude determination systems can often be realized with no construction of additional infrastructure. Although the reliability of such a system may be a limitation, the use of multiple LF signals would add a significant degree of redundancy. This should result in an acceptable system, given the lower reliability requirements inherent in small low-cost UAV applications.

Further work is needed to develop a real-time AHRS incorporating an LF radio polarization sensor, and to perform flight tests in a fully autonomous UAV. The system's accuracy could also be improved by implementing additional calibration steps, and extending the attitude filter to track changes in the LF reference vectors over time.

REFERENCES

[1] S. Siebert and J. Teizer, "Mobile 3D mapping for surveying earthwork projects using an unmanned aerial vehicle (UAV) system," *Automation in Construction*, vol. 41, pp. 1 – 14, May 2014.

[2] H. Bendea, P. Boccardo, S. Dequal, F. Giulio-Tonolo, D. Marenchino, and M. Piras, "Low cost UAV for post-disaster assessment," in *The Int. Archives of the Photogrammetry, Remote Sensing and Spatial Information Sci.*, vol. 37, no. B8, Beijing, China, July 2008, pp. 1373–1379.

[3] M. A. Goodrich, B. S. Morse, D. Gerhardt, J. L. Cooper, M. Quigley, J. A. Adams, and C. Humphrey, "Supporting wilderness search and rescue using a camera-equipped mini UAV," *J. Field Robotics*, vol. 25, no. 1-2, pp. 89–110, January-February 2008.

[4] J. Berni, P. J. Zarco-Tejada, L. Suarez, and E. Fereres, "Thermal and narrowband multispectral remote sensing for vegetation monitoring from an unmanned aerial vehicle," *IEEE Trans. Geosci. Remote Sens.*, vol. 47, no. 3, pp. 722 –738, March 2009.

[5] W. T. Higgins, "A comparison of complementary and kalman filtering," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 11, no. 3, pp. 321 –325, May 1975.

[6] R. Mahony, T. Hamel, and J.-M. Pflimlin, "Nonlinear complementary filters on the special orthogonal group," *IEEE Trans. Autom. Control*, vol. 53, no. 5, pp. 1203–1218, 2008.

[7] K. Jensen, "Generalized nonlinear complementary attitude filter," *J. Guid. Control Dyn.*, vol. 34, no. 5, pp. 1588–1593, September-October 2011.

[8] E. J. Lefferts, F. L. Markley, and M. D. Shuster, "Kalman filtering for spacecraft attitude estimation," *J. Guid. Control Dyn.*, vol. 5, no. 5, pp. 417–429, 1982.

[9] F. L. Markley, "Attitude error representations for kalman filtering," *J. Guid. Control Dyn.*, vol. 26, no. 2, pp. 311–317, March-April 2003.

[10] J. L. Crassidis, F. L. Markley, and Y. Cheng, "Survey of nonlinear attitude estimation methods," *J. Guid. Control Dyn.*, vol. 30, no. 1, pp. 12–28, January-February 2007.

[11] F. L. Markley, "Attitude determination using vector observations and the singular value decomposition," *J. Astronaut. Sci.*, vol. 36, no. 3, pp. 245–258, July-September 1988.

[12] M. D. Shuster and S. D. Oh, "3-axis attitude determination from vector observations," *J. Guid. Control Dyn.*, vol. 4, no. 1, pp. 70–77, January-February 1981.

[13] F. L. Markley and D. Mortari, "Quaternion attitude estimation using vector observations," *J. Astronaut. Sci.*, vol. 48, no. 2-3, pp. 359–380, April-September 2000.

[14] M. Euston, P. Coote, R. Mahony, J. Kim, and T. Hamel, "A complementary filter for attitude estimation of a fixed-wing UAV," in *Proc. 2008 IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, Nice, France, September 2008, pp. 340–345.

[15] S. Maguire and P. Robertson, "Low frequency radio signal polarisation sensor with applications in attitude estimation," in *Proc. IEEE Sens. Conf. 2014*, Valencia, Spain, Nov 2014, pp. 1587–1590.

[16] ——, "Low frequency radio polarization sensor with applications in attitude estimation," *IEEE Sens. J.*, vol. 15, no. 12, pp. 7304–7311, December 2015.

[17] W. F. Phillips, C. E. Hailey, and G. A. Gebert, "Review of attitude representations used for aircraft kinematics," *J. Aircraft*, vol. 38, no. 4, pp. 718–737, July-August 2001.

[18] D. Gebre-Egziabher, R. C. Hayward, and J. D. Powell, "Design of multi-sensor attitude determination systems," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 40, no. 2, pp. 627–649, April 2004.

[19] G. Wahba, "A least squares estimate of satellite attitude," *SIAM Review*, vol. 7, no. 3, p. 409, July 1965.

[20] J. L. Farrell, J. C. Stuelpnagel, R. H. Wessner, and J. R. Velman, "A least squares estimate of satellite attitude," *SIAM Review*, vol. 8, no. 3, pp. 384–386, July 1966.

[21] G. M. Lerner, "Spacecraft attitude determination and control," Berlin, Germany, pp. 421 – 429, 1978.

[22] R. P. Kornfeld, R. J. Hansman, and J. J. Deyst, "The impact of GPS velocity based flight control on flight instrumentation architecture," Ph.D. dissertation, MIT, June 1999.

[23] J. J. Deyst, R. P. Kornfeld, and R. J. Hansman, "Single antenna GPS information based aircraft attitude redundancy," in *Proc. 1999 American Control Conf.* San Diego, CA, USA: IEEE, June 1999, pp. 3127–31.

[24] R. P. Kornfeld, R. J. Hansman, and J. J. Deyst, "Single-antenna GPS-based aircraft attitude determination," *J. Inst. Nav.*, vol. 45, pp. 51–60, 1998.

[25] R. P. Kornfeld, R. J. Hansman, J. J. Deyst, K. Amonlirdviman, and E. M. Walker, "Applications of global positioning system velocity-based attitude information," *J. Guid. Control Dyn.*, vol. 24, pp. 998–1008, September-October 2001.

[26] S. A. Parvez, "Attitude determination using antenna polarization angles," *J. Guid. Control Dyn.*, vol. 14, no. 2, pp. 236–240, March-April 1991.

[27] M. Schwendeman and J. Thomson, "A horizon-tracking method for ship-board video stabilization and rectification," *J. Atmos. Oceanic Technol.*, 2014.

[28] J.-Y. Bouguet, "Camera calibration toolbox for matlab," 2013.

Sean Maguire received the M.Eng. degree in Electrical and Information Sciences from the University of Cambridge, Cambridge, UK in 2011, and subsequently joined the Photonics and Sensors Group at the University of Cambridge to pursue the Ph.D. degree. His research interests include Low Frequency electromagnetic sensors and low-cost attitude determination systems.

Paul Robertson received the BA degree in Electrical Sciences from the University of Cambridge, Cambridge, UK in 1984 followed by a Ph.D. on thin film electronics in 1987 from the same institution.

From 1987 to 1995 he worked at the PA Consulting Group technology center in Melbourn before returning to the faculty in the Engineering Department at the University of Cambridge. He is author of over 50 technical papers and patents. His current research interests include electromagnetic sensors for instrumentation, hybrid-electric propulsion systems for aircraft and power generating systems. He has been a Chartered Engineer and a Member of the Institution of Engineering and Technology (IET) since 1995.