

Article

An Algorithm for Affordable Vision-Based GNSS-Denied Strapdown Celestial Navigation

Samuel Teague ^{1,2,*}  and Javaan Chahl ^{1,2} ¹ STEM Academic Unit, University of South Australia, Mawson Lakes, Adelaide, SA 5095, Australia² Platforms Division, Defence Science and Technology Group, Edinburgh, Adelaide, SA 5111, Australia

* Correspondence: samuel.teague@unisa.edu.au

Abstract: Celestial navigation is rarely seen in modern Uncrewed Aerial Vehicles (UAVs). The size and weight of a stabilized imaging system, and the lack of precision, tend to be at odds with the operational requirements of the aircraft. Nonetheless, celestial navigation is one of the few non-emissive modalities that enables global navigation over the ocean at night in Global Navigation Satellite System (GNSS) denied environments. This study demonstrates a modular, low cost, lightweight strapdown celestial navigation solution that is utilized in conjunction with Ardupilot running on a Cube Orange to produce position estimates to within 4 km. By performing an orbit through a full rotation of compass heading and averaging the position output, we demonstrate that the biases present in a strapdown imaging system can be nullified to drastically improve the position estimate. Furthermore, an iterative method is presented which enables the geometric alignment of the camera with the Attitude and Heading Reference System (AHRS) in-flight without an external position input. The algorithm is tested using real flight data captured from a fixed wing aircraft. The results from this study offer promise for the application of low cost celestial navigation as a redundant navigation modality in affordable, lightweight drones.

Keywords: UAV; navigation; celestial; strapdown; stellar; orbital; modular



Citation: Teague, S.; Chahl, J. An Algorithm for Affordable Vision-Based GNSS-Denied Strapdown Celestial Navigation. *Drones* **2024**, *8*, 652. <https://doi.org/10.3390/drones8110652>

Academic Editor: Pablo Rodríguez-González

Received: 6 September 2024

Revised: 31 October 2024

Accepted: 5 November 2024

Published: 7 November 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Celestial navigation is among the oldest forms of navigation in aviation [1]. The abundance of salient stars, known to high levels of precision, make them a useful cue for navigators when operating in clear conditions. The elevation of a star above the horizon would be measured, yielding a ‘line of position’. This process was repeated with different stars to fully determine the navigator’s position. The advancement of imaging and computing hardware saw the integration of autonomous star trackers into manned aircraft, as seen, for example, in Lockheed’s SR-71. These autonomous star trackers consisted of a mechanically stabilized telescope and inertial sensors, whose observations, when combined, could produce position estimates to within 0.3 nautical miles for up to 10 h of operation. While accurate, these sensors tended to be both heavy and voluminous, making them undesirable for more modern Size, Weight and Power Constrained (SWAP-C) applications.

By the advent of the 21st century, Global Positioning System (GPS) had become ubiquitous in avionics navigation. The introduction of GPS caused the interest in celestial navigation to wither due to its relative inaccuracy. Consequently, celestial navigation is primarily seen only in space-based systems, whose orientation must be known to high levels of precision. Nonetheless, celestial navigation was identified as a desirable alternative to GPS [2], primarily due its robustness against potential jamming. Critically, few GPS-denied alternatives exist that are capable of using passive sensors to estimate global position at night or over the ocean. For this reason, celestial navigation remains an important topic of research. The methodology presented in this paper demonstrates how celestial navigation may be utilized on low cost Uncrewed Aerial Vehicles (UAVs) that lack the precision of an expensive Attitude and Heading Reference System (AHRS).

Existing works on strapdown celestial navigation tend to be focused on simulation. One such example derives the equations for a celestial camera mounted directly to a strapdown inertial system [3]. Similarly, a more recent study identified that a strapdown celestial system is contingent on the initial conditions and will tend to diverge over time [4]. While some promise has been shown in the utility of such systems [5], the primary problem with the navigation method is that the strapdown system must estimate the camera biases and delineate these from true position error. This is difficult to achieve, as camera bias is perceptually identical to motion over the Earth. Utilizing the celestial system as a highly accurate attitude reference tends to produce useful results in high altitude aircraft [6] and spacecraft [7–11]; however, as a modular solution, the celestial system must have a direct input to the AHRS to improve the inertial position estimate. Alternatively, it has been theoretically shown that atmospheric refraction can be observed to resolve positions [12]. This technique relies on the observation of the atmospheric refraction of starlight, which occurs most significantly closest to the horizon. The angle of refraction tends to be minimal and difficult to observe, requiring a highly stabilized viewing platform.

There are a number of reasons why celestial navigation has not become ubiquitous in aviation and particularly in UAV applications. Firstly, there is significant complexity in designing and configuring the hardware. A celestial system must typically be gimballed to within arcminutes of precision, such that a narrow field of a view sensor may lock onto a single star. This could be achieved through gyroscopic stabilization; however, the inclusion of multiple kilograms of stabilization hardware is rarely justified to achieve a navigation outcome with significantly worse performance than GPS. Secondly, celestial systems require a clear view of the night sky, which places practical limitations on their use. Thirdly, alongside the mechanical complexity, there exists a significant computational complexity required to integrate a celestial payload into an existing system. A star database must be maintained, image processing algorithms must be implemented to identify and track stars, and the system may or may not require an interface for the existing AHRS to time and orientate data. These reasons have overshadowed the benefit of having an additional GPS-denied modality for navigation available to an aircraft.

Previous research has addressed the latter of these points, demonstrating that modern computer vision libraries and embedded hardware are capable of handling the computational requirements of celestial navigation [13,14]. We address the first of these points in this study by demonstrating the use of a strapdown celestial system for navigation. Contrary to the stabilized alternative, a strapdown celestial system contains low mechanical complexity, is lightweight, and can be implemented at a low cost. The primary trade-off is in the accuracy of the navigation system. While high levels of accuracy are theoretically achievable (depending on the quality of the optical system), the limiting factor with strapdown celestial navigation is the accuracy of the AHRS. As a rough guide, an attitude error of 1° correlates with a position error of approximately 100 km. It is not uncommon for low cost autopilots (such as the Cube Orange running ArduPlane v4.5 firmware) to produce attitude errors in the vicinity of multiple degrees. Such biases would lead to positional offsets that are far too significant for use in any real application.

We address this shortcoming by demonstrating how a simple orbital motion can significantly improve a celestial position estimate. This technique has been used on the ground to correct for boresight errors [15], and we demonstrate through experimentation that a similar technique may be used on aircraft to correct for attitude errors. The principle behind this maneuver is that a misaligned boresight will trace a circle about the true boresight if a full azimuthal revolution is performed. The misalignment presents itself in the navigation frame as a circular error in latitude and longitude, enabling the averaging of the results to attain an improved position estimate. We demonstrate that this technique is effective even when the camera is misaligned with the AHRS, offering a reliable method for estimating the position to within 4 km from an unknown state. While lacking precision, this method of localization is absolute, cheap to implement, and relatively lightweight.

To our knowledge, no such technique currently exists in the literature. The methods proposed here are intended to demonstrate the application of strapdown celestial navigation on a fixed-wing drone platform, while addressing the key practical difficulties in doing so. Existing works are iterative, relying on the integration of celestial measurements with inertial measurements using a filter such as an Extended Kalman Filter (EKF) or Unscented Kalman Filter (UKF). This research demonstrates that a strapdown system can be treated as a stand-alone, modular addition to an inertial system, such that the measurements do not need to be integrated into the filter. Independence from an inertial filter enables the celestial system to produce true global position measurements that are not affected by initial conditions. Provided the use of an accurate clock, the results presented in this paper will not degrade over time.

The remainder of this paper addresses the theory, implementation, and results of the proposed method. We outline the equations for observing stars and estimating global positions in Section 2, we describe the methodology used in star detection/tracking and the experimental configuration, we explicitly define the equations for computing the mean position in Section 3, we present the results from the flight trial and simulation in Section 4, we discuss the results in Section 5, and we conclude in Section 6.

2. Theory

2.1. Star Observation

An observer may estimate their position with knowledge of the current sidereal time and direction vectors to three or more stars in the local North East Down (NED) frame. A camera system is mounted relative to the body frame of the aircraft, with Direction Cosine Matrix (DCM) $\mathbf{C}_{b/c}$ describing the orientation of the camera with respect to the aircraft. Given an aircraft with the orientation described by the DCM $\mathbf{C}_{l/b}$, an observation \mathbf{X}^c in the camera frame may be transformed to the local NED frame by the following equation:

$$\mathbf{X}^l = \mathbf{C}_{l/b}\mathbf{C}_{b/c}\mathbf{X}^c \quad (1)$$

Multiple stars are observed simultaneously through the imaging system. A calibrated camera may be described by the pinhole projection model, such that

$$\alpha\mathbf{x}^p = \mathbf{K}\mathbf{X}^c \quad (2)$$

where \mathbf{x}^p contains the homogeneous pixel coordinates of the projected star, \mathbf{K} is the camera intrinsic matrix, and \mathbf{X}^c is the direction vector of the star in the camera coordinate frame. Consequently, with knowledge of the camera intrinsic matrix, given an observation \mathbf{x}^p , the vector \mathbf{X}^c may be computed as follows:

$$\mathbf{X}^c = \alpha\mathbf{K}^{-1}\mathbf{x}^p \quad (3)$$

where the scale factor α may be computed with the knowledge that \mathbf{X}^c is unitary. Given that $|\mathbf{X}^c| = 1$, it follows that $\alpha|\mathbf{K}^{-1}\mathbf{x}^p| = 1$. Therefore, α can be calculated as follows:

$$\alpha = \frac{1}{|\mathbf{K}^{-1}\mathbf{x}^p|} \quad (4)$$

2.2. Position Estimation

The global position may be calculated using the zenith-angle, ζ , of stars. Following the methodology presented by [16], each measurement ζ generates a plane that intersects with the terrestrial sphere, yielding a circle on which the observer may be located. Two observations generate a line, of which only two points on this line intersect with the terrestrial sphere. This case is depicted in Figure 1. Three observations precisely define the location of the observer, and more than three observations can create an over-constrained system. We

may utilize redundant measurements by estimating the least-squares approximation from three or more star observations.

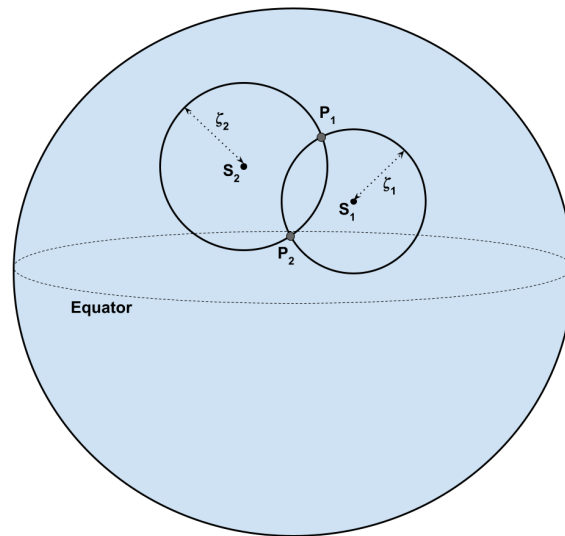


Figure 1. The circles around points S_1 and S_2 show the intersection of their respective planes with the terrestrial sphere, forming circles about which the observation could have been made. Points S_1 and S_2 are taken from the observed stars, located such that their respective star is observed at the zenith. The zenith-angles ζ_1 and ζ_2 represent the angle at which the stars were actually observed. In this case, two stars were observed, resulting in two potential points (P_1 and P_2) at which this observation could have been made. Additional observations serve to reduce this ambiguity.

Given n stars, with right ascension α and declination δ , observed at a zenith angle ζ , we may define n planes within the geographic coordinate system. If we define λ as the longitude of an observer whose zenith is directed towards the star, which can be expressed as the difference between the star’s right ascension and the current Greenwich hour angle of Aries, $\lambda = \alpha - GHAT$, then it can be seen that the equation of a plane whose normal vector is directed towards the star is given by

$$ax + by + cz - p = 0 \tag{5}$$

where

$$\begin{aligned} a &= \cos \lambda \cos \delta \\ b &= \sin \lambda \cos \delta \\ c &= \sin \delta \\ p &= \cos \zeta \end{aligned} \tag{6}$$

Therefore, given a minimum of three star observations, we can build matrices \mathbf{A} and \mathbf{p} to formulate a least-squares approximation of the intersection of these planes. Vertically stacking each equation yields

$$\mathbf{Ax} = \mathbf{p} \tag{7}$$

where

$$\mathbf{A} = \begin{bmatrix} a_1 & b_1 & c_1 \\ a_2 & b_2 & c_2 \\ \vdots & \vdots & \vdots \\ a_n & b_n & c_n \end{bmatrix}, \mathbf{x} = \begin{bmatrix} x \\ y \\ z \end{bmatrix}, \mathbf{p} = \begin{bmatrix} p_1 \\ p_2 \\ \vdots \\ p_n \end{bmatrix}$$

We apply the standard least-squares solution to find the point of intersection of the planes, \mathbf{x} , as follows:

$$\mathbf{x} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{p} \quad (8)$$

Putting \mathbf{x} back into geographical coordinates, we find the latitude and longitude which best fits the star observations:

$$\phi = \tan^{-1} \left(\frac{z}{\sqrt{x^2 + y^2}} \right) \quad (9)$$

$$\lambda = \tan^{-1} \left(\frac{y}{x} \right) \quad (10)$$

This value for the latitude and longitude minimizes the squared error in the position vectors. This will be used in Section 3 for converting many observations into a single position estimate.

3. Methods

A modular celestial navigation system should be mounted rigidly with respect to the AHRS. The orientation of the aircraft, $\mathbf{C}_{l/b}$, is provided by the autopilot. The orientation of the camera, $\mathbf{C}_{b/c}$, is unknown and must be calibrated to obtain accurate positional information. For the purposes of this study, we assume that the precise orientation of the camera remains unknown, and we consequently accept that the individual position estimates will be erroneous. It will be shown later that this misalignment may be overcome through the use of averaging. This is particularly useful given that the camera itself may be mounted separately from the autopilot, thus being subjected to factors such as vibration, aerodynamic loading, or changes in AHRS biases, which cause the relative orientation between the AHRS and the camera to change over time.

3.1. Star Detection and Tracking

Stars are detected within an image using a basic binary thresholding operation. Given a mean pixel value within the frame, μ , and standard deviation, σ , a binary threshold is set at $\mu + 5\sigma$, as advised in [17]. The contours are extracted, and an agglomerative clustering algorithm is applied to cluster redundant detections.

Once detected, a standard Kalman filter is used to track individual stars between frames. Each star is defined by its state vector:

$$\mathbf{x}_k = \begin{bmatrix} u \\ v \\ \dot{u} \\ \dot{v} \end{bmatrix} \quad (11)$$

and a constant sized bounding box. The states u, v, \dot{u} , and \dot{v} describe the x position, y position, x velocity, and y velocity of the star on the image plane, respectively. The position of the star is taken to be the subpixel maxima, computed as a weighted average of the 3×3 Region of Interest (ROI) centered on the peak pixel value, as described by [18]

$$(u, v) = \left(\frac{\sum_{i,j} I_{i,j} u_i}{\sum_{i,j} I_{i,j}}, \frac{\sum_{i,j} I_{i,j} v_j}{\sum_{i,j} I_{i,j}} \right) \quad (12)$$

The state transition matrix \mathbf{F} is defined assuming a constant velocity as follows:

$$\mathbf{F} = \begin{bmatrix} 1 & 0 & \Delta t & 0 \\ 0 & 1 & 0 & \Delta t \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (13)$$

The position and velocity of the stars are observed directly; therefore, the observation matrix \mathbf{H} is defined as the 4×4 identity, such that the observation equation is simply given by

$$\mathbf{z}_k = \mathbf{x}_k \quad (14)$$

Following the standard Kalman filter equations, the a priori state and covariance estimates are computed:

$$\hat{\mathbf{x}}_{k+1} = \mathbf{F}\mathbf{x}_k \quad (15)$$

$$\hat{\mathbf{P}}_{k+1} = \mathbf{F}\mathbf{P}_k\mathbf{F}^T + \mathbf{Q} \quad (16)$$

for the covariance matrix \mathbf{P}_k with process noise \mathbf{Q} . The innovation is calculated as follows:

$$\mathbf{y}_k = \mathbf{z}_k - \hat{\mathbf{x}}_{k+1} \quad (17)$$

with covariance, \mathbf{S}_k , given by

$$\mathbf{S}_k = \hat{\mathbf{P}}_{k+1} + \mathbf{R}_k \quad (18)$$

where \mathbf{R}_k is the measurement noise. The Kalman gain is computed as follows:

$$\mathbf{K}_k = \hat{\mathbf{P}}_{k+1}\mathbf{S}_k^{-1} \quad (19)$$

and the posterior updates are applied to the a priori estimates:

$$\mathbf{x}_{k+1} = \hat{\mathbf{x}}_{k+1} + \mathbf{K}_k\mathbf{y}_k \quad (20)$$

$$\mathbf{P}_{k+1} = (\mathbf{I} - \mathbf{K}_k)\hat{\mathbf{P}}_{k+1} \quad (21)$$

The measurement noise, \mathbf{R}_k , is defined as a diagonal matrix, with elements equal to $[0.5 \ 0.5 \ 1.0 \ 1.0]$, and the process noise is also defined as a diagonal matrix, with elements equal to $[4.0 \ 4.0 \ 2.0 \ 2.0]$. The measurement noise was selected based on the calibration accuracy of the camera, and the process noise was experimentally tuned to minimize the occurrence of lost tracks under high motion conditions. The a priori state estimate is used to propagate the bounding box containing the region of interest, within which the peak is detected using Equation (12). The bounding box is centered on $[u, v]$, rounded to the nearest pixel. The width and height of the bounding box were fixed at 21 pixels for this imaging system, given an image resolution of 1936×1216 and a field of view of 53.5 deg. A visual snapshot of the tracking system can be seen in Figure 2.




Figure 2. Star tracker operating on video footage captured in-flight. Image intensity is amplified $10\times$.

3.2. Experimental Configuration

Celestial imagery was captured in-flight from a 4 m flying wing UAV. The selected autopilot (Cube Orange) was mounted at the center of mass, and the celestial payload was mounted in the shoulder of the aircraft (see Figure 3). The celestial payload collected imagery at a rate of 10 Hz, as well as the attitude and position data from the autopilot, enabling post-analysis. A Raspberry Pi 5 (Raspberry Pi Ltd., Cambridge, UK) was used as the companion computer, interfacing with the celestial camera and the autopilot. An Alviium 1800 U-240 (Allied Vision, Stradtroda, Germany) monochrome sensor fitted with a f/1.4 6 mm wide angle lens was chosen for the imaging system. A serial Universal Asynchronous Receiver Transmitter (UART) link between the Raspberry Pi and the Cube Orange facilitated the transport of MAVLink v2.0 messages. The AHRS data from ArduPlane’s EKF3 was recorded at a rate of 30 Hz, and the ground truth GPS position data were recorded at a rate of 10 Hz. The Raspberry Pi clock was synchronized with GPS time prior to takeoff. The Raspberry Pi and the camera were mounted on a PLA 3D-printed structure (see Figure 4) for integration into the airframe.

The test flight was conducted on a moonless night. The wind was modest, typically remaining below 5 m/s. The flight plan consisted of both straight legs and orbital trajectories with varying radii. An overview of the flight plan can be seen in Table 1.

Table 1. Flight plan for the test flight.

Description	Direction	Repeats	Radius (m)	Altitude (m)
Takeoff				30
Orbit to Altitude	CCW	4	300	800
Straight Legs	N/S	6	2500	800
Orbit	CW	1	1200	800
Orbit	CW	2	600	800
Orbit	CCW	1	600	800
Orbit	CCW	1	1200	800
Orbit to Altitude	CCW	1	300	100
Land				0

The total flight lasted 72 min. Astronomical twilight ceased at 19:32, and takeoff was conducted at 19:53. The GPS receiver in the aircraft was not allowed to be switched off to enable emergency failsafes and prevent fence breaches. We address the consequences of this in Section 4.4.

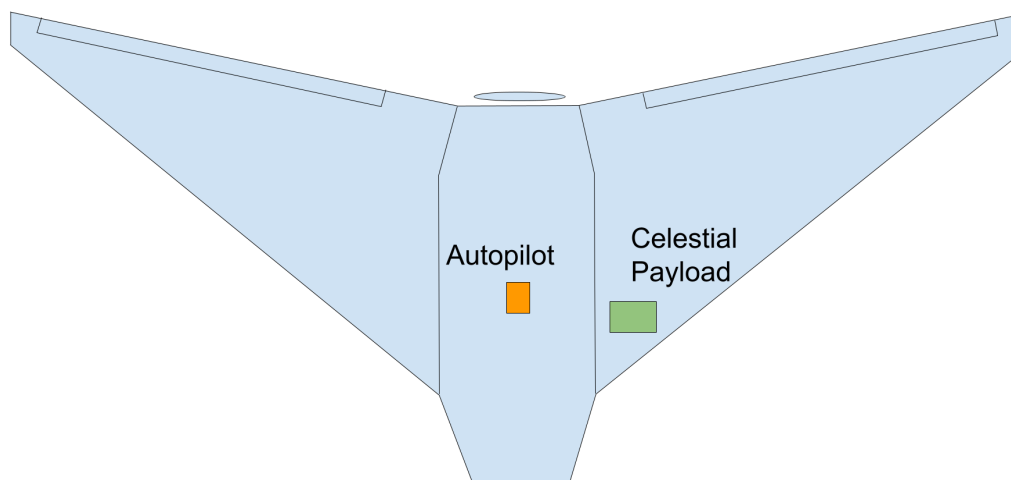


Figure 3. Platform layout of the autopilot and celestial payload within the airframe.



Figure 4. The celestial payload, consisting of a Raspberry Pi 5 and an Alvium 1800 U-240 monochrome sensor fitted with a 6 mm f/1.4 wide angle lens.

3.3. Position Estimation

The star tracker in Section 3.1 operates on distorted images. The pixel location from each star tracker is extracted, rectified, and subsequently converted to a unit vector in NED coordinates according to Section 2.1. Given an observation $\mathbf{X}^l = [x \ y \ z]^T$, the elevation of a star above the horizon is calculated as

$$el = -\tan^{-1} \frac{y}{\sqrt{x^2 + y^2}} \quad (22)$$

and, subsequently, the zenith-angle ζ is calculated as $90 - el$.

In the initial case, the camera orientation $\mathbf{C}_{l/b}$ is not known to a high level of precision. Assuming that an attempt is made to mount the camera in alignment with the autopilot, an initial DCM may be formulated from some combination of 90° Euler rotations. It will be seen that as long as this orientation is accurate to within a hemisphere of tolerance, the exact value does not matter, as it will be recalculated in flight.

If a minimum of six stars are visible within the frame, a Random Sample Consensus (RANSAC) approach to position estimation may be used to remove bias from false detections. RANSAC is a technique used to identify outliers in sample data and omit them from the estimation process. This is useful for detecting misidentified stars or for removing non-star detections (such as satellites or overhead planes). In the context of position estimation, the RANSAC algorithm randomly selects three stars to generate a position estimate. Subsequently, the algorithm compares the location of the remaining stars against the position estimate. If the remaining stars are within some heuristic tolerance, they are considered inliers. If they are outside of tolerance, they are considered outliers. The position estimate which minimizes the mean squared error of the inliers is chosen as the estimate. A pseudocode implementation of this RANSAC position estimation is outlined in Algorithm 1.

Algorithm 1 Position-RANSAC

```

 $C_{l/b}, C_{b/c}$ 
GHAF  $\leftarrow$  getGHAAries()
As  $\leftarrow$  []
Ps  $\leftarrow$  []
for obs in observations do
    az, el = computeAzEl(obs,  $C_{l/b}, C_{b/c}$ )
    a, b, c, p = computePositionCoefficients(GHAF, az, el, obs.ra, obs.dec)
    As.append([a, b, c])
    Ps.append([p])
end for

inliers  $\leftarrow$  0
outliers  $\leftarrow$  0
best error  $\leftarrow$  inf
x  $\leftarrow$  [0, 0, 0]
for 0 < i < iterations do
    j  $\leftarrow$  3 randomly selected indexes
    build A and p using As[j] and Ps[j]
     $\hat{x} \leftarrow (A^T A)^{-1} A^T p$ 
    inlier error  $\leftarrow$  0
    for all remaining indexes k do
         $\mathcal{E} \leftarrow Ps[k] - As[k]\hat{x}$ 
        if  $\mathcal{E} < tolerance$  then
            inliers += 1
            inlier error +=  $\mathcal{E}$ 
        else
            outliers += 1
        end if
    end for
    if inliers : outliers > acceptable ratio then
        mean error = inlier error / inliers
        if mean error < best error then
            best error = mean error
            x  $\leftarrow$   $\hat{x}$ 
        end if
    end if
end for
 $\phi \leftarrow \tan^{-1} \left( \frac{x[2]}{\sqrt{x[0]^2 + x[1]^2}} \right)$ 
 $\lambda \leftarrow \tan^{-1} \left( \frac{x[1]}{x[0]} \right)$ 
return  $\phi, \lambda$ 

```

We first demonstrate that, in the general case, the DCM $C_{l/b}$ is not fixed during flight. In an integrated solution, the celestial system would provide attitude measurements to the EKF, and the offset from this attitude would be estimated. As a modular solution, however, this is infeasible. Taking the output from the EKF to be the true orientation of the aircraft, we can use the true GPS position to compute $C_{l/b}$. We use the Kabsch-RANSAC methodology presented in [19] to calculate the ideal rotation between a star's theoretical location and its actual location. In conjunction with the autopilot output, this rotation yields $C_{l/b}$. Testing over an 89 s section of flat and level flight, it can be seen in Figure 5 that the camera roll angle shifts by approximately 0.2° , the pitch angle shifts by approximately 0.05° , and the yaw angle shifts by approximately 0.3° . In the context of celestial navigation, these are significant perturbances which, had the offsets been attributed to position as opposed to camera orientation, would have been interpreted as around 30 km of positional

offset. Indeed, by assuming that $C_{l/b}$ is fixed, we can see in Figure 6 that the latitude drifts by approximately 0.2° over the same short window of flat and level flight. Over an extended period (hours), the orientation may shift such that positional errors are in the range of hundreds of kilometers.

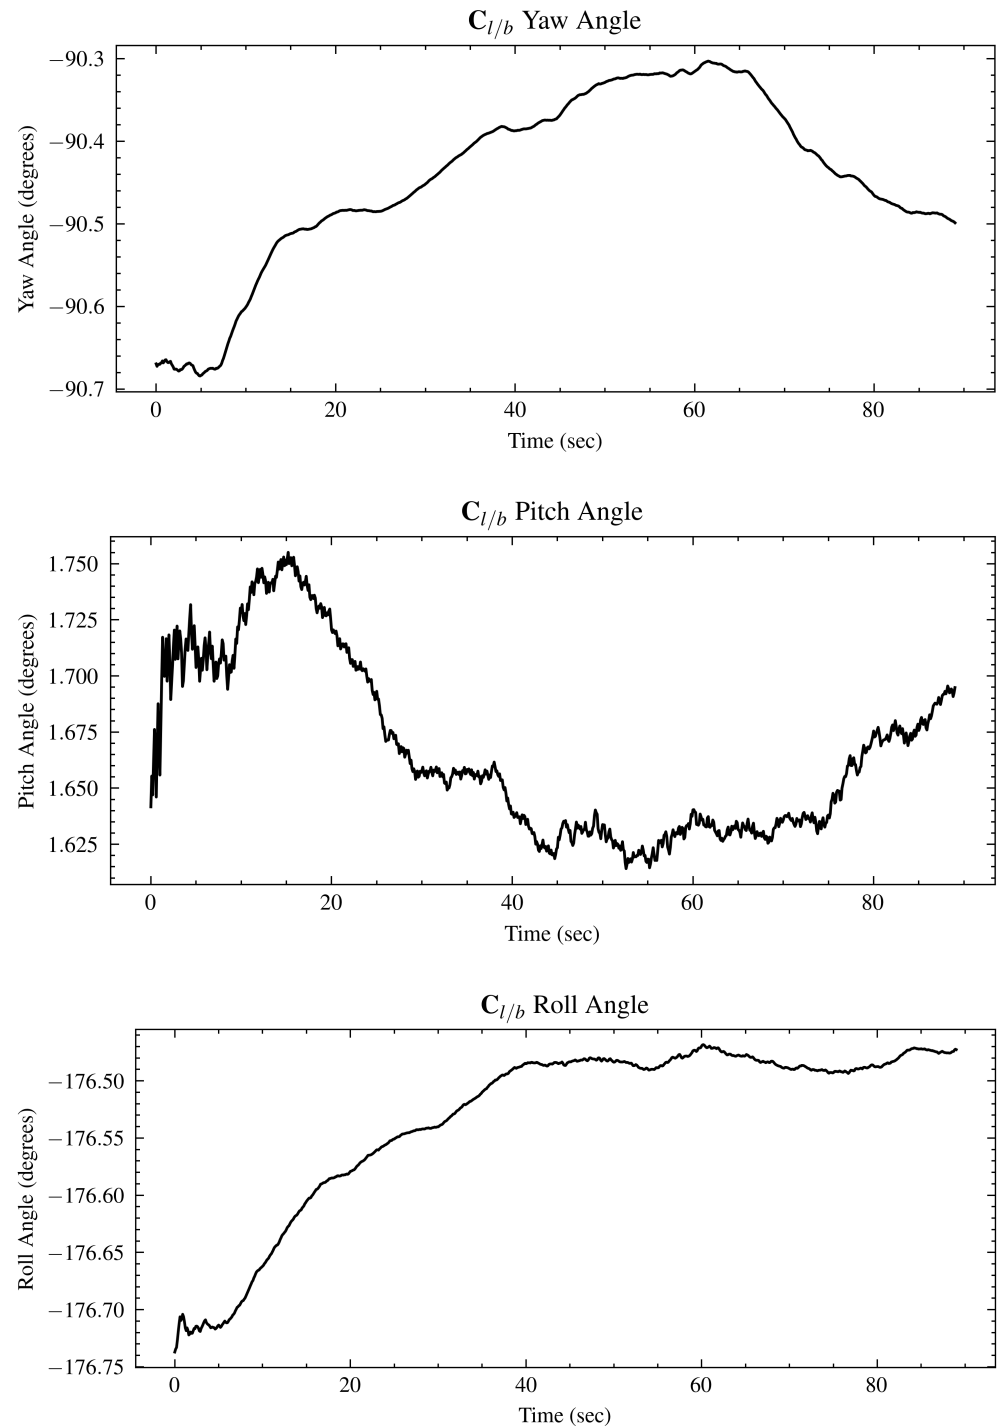


Figure 5. True camera orientation in the aircraft body frame, as calculated using a combination of stellar observations, aircraft attitude, and GPS data. Data are captured from a straight segment of flight over 89.5 s (895 video frames). **Top:** Calculated camera yaw angle in aircraft body frame. **Middle:** Calculated camera pitch angle in aircraft body frame. **Bottom:** Calculated camera roll angle in aircraft body frame.

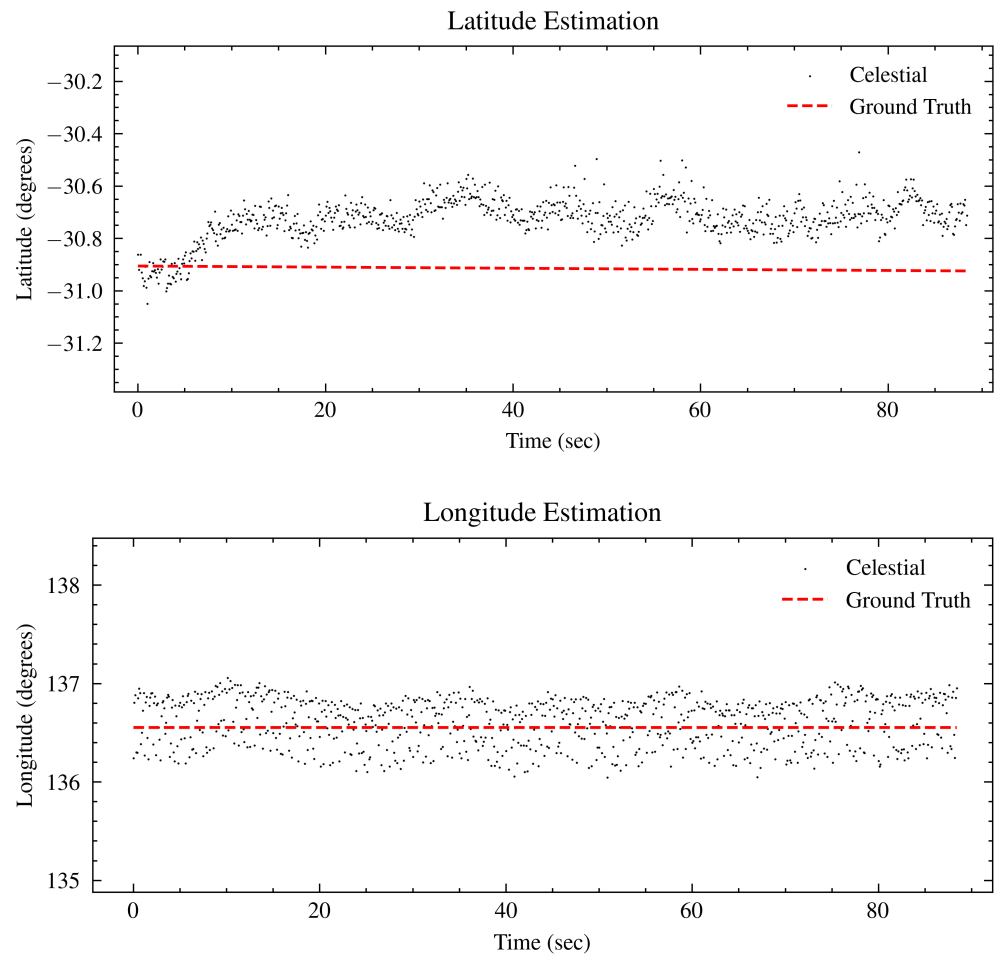


Figure 6. Position estimation over an 89.5 s straight segment of flight. It can be seen that, by assuming the camera orientation is fixed, significant positional error makes its way into the estimate. **Top:** Latitude estimated using celestial imagery with Ardupilot AHRS. **Bottom:** Longitude estimated using celestial imagery with Ardupilot AHRS.

It is clear that a stand-alone celestial module cannot function in the conventional manner with consumer-grade hardware, unless the autopilot is capable of integrating the attitude output from the celestial system into its own filter and estimating the camera orientation. We will demonstrate now a method which, given a very rough initial estimate of the camera orientation, is capable of returning position estimates to within 4 km. This is significant in the context of long flights, where alternative dead-reckoning solutions would experience positional drift that is either linear (assuming velocity measurements are available) or quadratic (assuming only acceleration measurements) as a function of time.

By performing a rotation through 360° of compass heading, at an approximately constant yaw rate, the position estimates can be averaged to find an approximation of the orbital center. Provided n images throughout an orbit are available, there exists n latitude and longitude estimates, which may be expressed as unit vectors in the Earth Centred Earth Fixed (ECEF) frame \mathbf{p}_i^{ecef} . The mean position is taken to be the mean of these vectors:

$$\bar{\mathbf{p}}^{ecef} = \frac{1}{n} \sum_{i=1}^n \mathbf{p}_i^{ecef} \quad (23)$$

The latitude and longitude are simply calculated as

$$\phi = \sin^{-1}(z) \quad (24)$$

$$\lambda = \tan^{-1}\left(\frac{y}{x}\right) \quad (25)$$

where x , y , and z are the elements of the mean position vector.

Once an initial position estimate has been formulated, it may be used to calculate a more precise estimate of the camera orientation $\mathbf{C}_{b/l}$. The aim of orientation estimation is to find the rotation matrix \mathbf{R} that minimizes the residual error between a set of observations \mathbf{u}^c represented in the camera frame of reference and a set of theoretical unit vectors \mathbf{v}^l in the NED frame of reference:

$$\mathcal{E} = \sum_{i=1}^n \|\mathbf{v}_i^l - \mathbf{R}\mathbf{u}_i^c\|^2, \quad (26)$$

In this case, the rotation matrix \mathbf{R} is composed of the aircraft DCM, $\mathbf{C}_{l/b}$, and the camera DCM, $\mathbf{C}_{b/c}$:

$$\mathbf{R} = \mathbf{C}_{l/b}\mathbf{C}_{b/c} \quad (27)$$

where we accept $\mathbf{C}_{l/b}$ as a deterministic output from the autopilot, and so the camera orientation may be found given \mathbf{R} and $\mathbf{C}_{l/b}$:

$$\mathbf{C}_{b/c} = \mathbf{C}_{l/b}^T \mathbf{R} \quad (28)$$

Each observation \mathbf{u}_i^c is correlated with a database star. This correlation is obtained through star identification. During the instantiation of the star tracker, a lost-in-space log-polar star identification algorithm is used to determine the IDs of each star in the frame [20]. The theoretical position vectors \mathbf{v}_i^l in the NED coordinates are generated from the star's right ascension and declination, the time of observation, and the estimated latitude and longitude [13]. This allows for the use of the Kabsch algorithm [21] to find the rotation \mathbf{R} between the observed stars and the database stars. Following the implementation in [19], the algorithm is provided in Algorithm 2, where \mathbf{A} is the matrix of n observation vectors with dimension $[n \times 3]$ and \mathbf{B} is the matrix of n theoretical vectors, also with dimension $[n \times 3]$. The resulting rotation \mathbf{R} is used in conjunction with the most recent autopilot attitude data to find the camera orientation $\mathbf{C}_{b/c}$.

Algorithm 2 Kabsch


```

Translate vectors in  $\mathbf{A}$  and  $\mathbf{B}$  such that centroid is at origin
Compute the matrix  $\mathbf{C} = \mathbf{A}^T \mathbf{B}$ 
Compute the singular value decomposition  $\mathbf{U}, \mathbf{\Sigma}, \mathbf{V} = \text{SVD}(\mathbf{C})$  such that  $\mathbf{C} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ 
Set diagonal elements  $\Sigma_1$  through to  $\Sigma_{n-1} = 1$ .
if  $\det(\mathbf{UV}) > 0$  then
     $\Sigma_n = 1$ 
else
     $\Sigma_n = -1$ 
end if
Compute the rotation matrix  $\mathbf{R} = \mathbf{V}\mathbf{\Sigma}\mathbf{U}^T$ 


```

Once the camera orientation has been found, two options are presented, depending on the computational power of the flight computer. If the computer is capable of parallel processing, then the set of theoretical star locations \mathbf{v}^l may be re-calculated using the updated latitude and longitude, and the estimated positions may be re-calculated using the updated camera DCM $\mathbf{C}_{b/c}$. This enables the system to recursively converge on a position estimate from a single set of orbital data. Alternatively, it is possible to repeat an orbit and calculate a new set of latitudes and longitudes. This method does not require post hoc processing and may be better suited to real-time applications. For the purposes of this


study, we use the former method, as this allows us to treat each orbit as an independent sample, yielding an independent position estimate, thus giving us deeper insight into how the characteristics of a given orbit affect the position estimate. An example of this process can be seen in Figure 7.




(a) Initial position estimate with an uncalibrated camera. It can be seen that the initial position error is large, at 44.50 km.



(b) Second iteration of position estimation, using the position from the previous iteration to calibrate the camera. It can be seen that the position error from the first estimate has been interpreted as a camera orientation offset, and the radius of the position estimates is now approximately equal to the positional error in the output of the first iteration. The position error is now 2.39 km.



(c) Third iteration of position estimation. The improved position estimate from the previous output has been used to calibrate the camera. The primary source of error now lies in the asymmetry of the orbital dynamics. The position error has increased slightly to 3.03 km.



(d) Final iteration of position estimation. The final positional error is 2.54 km.

Figure 7. Visualization of the position estimates as the algorithm converges on the location of the aircraft. Each blue point represents an independent position that was calculated from an image frame, the red dot represent the estimated position taken from the mean of the individual estimates, and the yellow dot represents the true center of the aircraft's orbit.

4. Results

4.1. Position Estimation

Following the experimental configuration in Section 3, the position of the aircraft is estimated for each orbit. The camera was initially aligned using a rotation matrix built from Euler angles of -90° , 0° , and 180° in yaw, pitch, and roll, respectively, for a yaw–pitch–roll rotation sequence. This is a coarse approximation based on the orientation with which the camera was mounted to the airframe. Of course, the true orientation differs from this; as can be seen in Figure 5, there is approximately a 5° error in this initial orientation.

Each orbit is treated as an independent observation. This is achieved by resetting the camera orientation at the initialization of every orbit, thus requiring the algorithm to recalculate the camera DCM. The output from each of the orbital motions listed in Table 1 is shown in Table 2. The position error is measured as the distance between the final celestial position estimate and the center of the orbit, calculated using the Haversine function.

Table 2. Celestial positioning results.

Description	Direction	Radius (m)	Pos. Error (km)	Iterations
Orbit to Altitude (1)	CCW	300	3.56	5
Orbit to Altitude (2)	CCW	300	7.18	4
Orbit to Altitude (3)	CCW	300	9.67	4
Orbit to Altitude (4)	CCW	300	9.89	4
Orbit (5)	CW	1200	2.21	4
Orbit (6)	CW	600	3.48	5
Orbit (7)	CW	600	2.54	5
Orbit (8)	CCW	600	1.73	5
Orbit (9)	CCW	1200	2.90	4
Orbit to Altitude (10)	CCW	300	7.16	6

4.2. Estimation Accuracy

It can be seen that the orbits which are neither climbing nor descending produced more accurate positional data. These orbits had larger radii, which consequently produced greater numbers of position estimates per orbit. Additionally, the variance in the pitch and roll axis during ascent/descent tends to be greater, as the total energy control system utilizes the pitch axis to throttle the climb/descent rate. By taking these two factors into account, we calculate the standard error as a function of the covariance in pitch and roll, as well as the number of samples. The generalized variance is calculated by taking the determinant of the covariance matrix in pitch and roll, where this matrix is given by

$$\Sigma_{\theta,\phi} = \begin{bmatrix} \text{COV}(\theta, \theta) & \text{COV}(\theta, \phi) \\ \text{COV}(\phi, \theta) & \text{COV}(\phi, \phi) \end{bmatrix} \quad (29)$$

where the function $\text{COV}(x, y)$ describes the covariance between sets x and y :

$$\text{COV}(x, y) = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{n} \quad (30)$$

given n samples. The standard deviation is found given the generalized variance:

$$\sigma = \sqrt{|\Sigma_{\theta,\phi}|} \quad (31)$$

and thus the standard error in roll and pitch is calculated given the standard deviation and the number of samples as follows:

$$SE = \frac{\sigma}{\sqrt{n}} \quad (32)$$

It can be seen in Figure 8 that the relationship between the estimated standard error and the true error is approximately linear. The accuracy of the first orbit is likely due to chance; this is a good indicator that the factors dictating an accurate positional estimate are indeed the variance in pitch and roll throughout the orbit and the total number of samples. The inverse of the trend-line may be used as an estimate for the Circular Error Probable (CEP), where

$$CEP = 1205 \times SE - 0.567 \tag{33}$$

and CEP is estimated in kilometers.

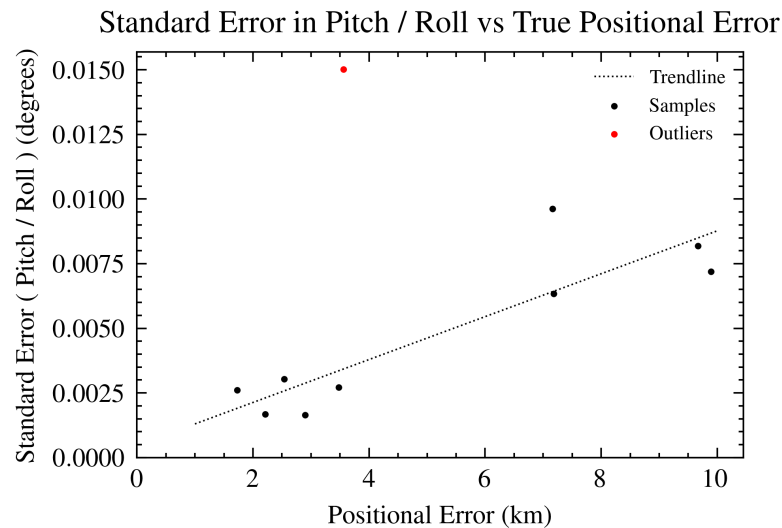


Figure 8. Plot of the standard error in pitch and variance versus the true error. It can be seen that the relationship is approximately linear.

4.3. Initial Conditions

We demonstrate here that the initial value of the camera rotation matrix has little effect on the final position estimate. This is significant because, as shown in Section 3, many factors may cause changes in the relative orientation between the camera and the AHRS. Choosing orbit 5 as an example, we show that as long as the estimated boresight of the camera is accurate to within a hemisphere of tolerance (that is, the estimated boresight is within 90° of the true boresight), the algorithm will converge near the true location. A graphical representation of the first iteration given various camera calibrations is shown in Figure 9. In each case, with a position error less than 90°, the position estimate converged to within the estimated CEP. It can be seen in Table 3 that an initial calibration error beyond 90° results in a position estimate on the opposite side of the Earth.

Table 3. Sensitivity to initial camera calibration.

Initial Orientation Error	Final Pos. Error (km)	Iterations
Nil (calibrated)	2.47	2
Initial Guess (≈5°)	2.49	4
45°	2.46	5
60°	2.51	5
85°	2.49	6
120°	19,987.66	5




Figure 9. Visualization of position estimates for a given camera misalignment. It can be seen that as long as the camera is aligned to within 90°, the mean estimated position provides a good approximation of the true position.

4.4. Simulation of Wind Effects

The flight trial was conducted with modest amounts of wind, utilizing GPS to conduct an orbit. We postulate that in the presence of wind, the position is better attained by fixing the pitch and roll of the aircraft, such that a constant yaw rate is achieved in the local NED frame. This approach does not require GPS to perform an orbit and only relies on compass heading and inertial attitude sensors to follow a trajectory. Simulation results strongly indicate that GPS is not required but is in fact detrimental to the orbital method of position estimation. As a consequence of not using GPS, the aircraft is subjected to lateral drift in the local NED frame. The amount of drift experienced by the aircraft during an orbit is, however, minimal in comparison to the scale to which position is being estimated.

We used JSBSim to simulate the aircraft's flight dynamics, in conjunction with Ardupilot's software-in-the-loop simulator to generate synthesized motion-blurred imagery. The images were generated following the methodology in [13], factoring in the motion blur experienced by the camera. A strong southerly wind of 54 km/h was added to the simulation. An orbit was performed using GPS, and another was performed using a fixed pitch and roll angle with GPS disabled. During the fixed-attitude orbit, the aircraft drifted 1.05 km. The simulated trajectory of the GPS-denied orbit can be seen in Figure 10.





Figure 10. Simulated trajectory of the aircraft during a GPS-denied orbit in high wind conditions.

Applying the same methodology, the position was calculated for both the GPS-guided orbit and the GPS-denied orbit. The control protocol of the GPS-guided orbit was to maintain a ground track at a fixed radius about a point in an identical manner to what was performed in the real flight test. This method adjusts the pitch and roll of the vehicle to compensate for wind and altitude. Up-wind sections of the orbit are prolonged in time, and down-wind sections are contracted in time. The resulting position estimates are heavily biased towards the prolonged portion of the orbit, which consequently skews the mean calculation. Under high wind conditions, it can be seen that following a fixed-radius ground track is not optimal for position calculation. The resulting position error was 18.27 km, as seen in Figure 11. By contrast, the fixed-attitude orbit achieved a position error of 2.29 km, despite being subjected to over 1 km of drift. This highlights the importance of fixing the aircraft attitude rather than following a constant-radius orbit.



(a) Final iteration of position estimation from a GPS-guided orbit. The position error is 18.27 km. **(b)** Final iteration of position estimation from GPS-denied orbit. The position error is 2.29 km.

Figure 11. Comparison of position estimation in high wind conditions. It can be seen that the course followed by a GPS-guided orbit introduces errors in the estimation process. The GPS-denied orbit follows a fixed-attitude orbit, resulting in drift, but ultimately yielding a more accurate position.

5. Discussion

The results presented in Section 4 demonstrate that strapdown celestial navigation as a modular solution has potential use in Global Navigation Satellite System (GNSS)-denied UAV navigation. The celestial camera was mounted to the airframe, separate from the autopilot, but still utilizing the AHRS orientation information coming from the autopilot. This has significant implications for integration into SWAP-C airframes, in which the inclusion of stabilization hardware adds unwanted mass. With a modern GPS receiver weighing only a few grams and producing estimates to within 1m accuracy, it is

understandable that, given the choice, an alternative celestial navigation solution would not be included. Yet, GNSS denial is increasing in prevalence, and alternative navigation solutions must be explored if a UAV is to operate under such circumstances.

The most profound outcome from this study is the celestial system's independence from the initial conditions. Provided a functional AHRS and navigation system which is capable of performing an orbit, the celestial system can produce a position estimate from any unknown position on the globe, with a camera that is aligned to within 90° . The remaining source of error lies in the clock drift, which is typically equal to 3 ppm (0.3 s per 24 h) for a modern real-time clock. This has significant implications for long endurance aircraft, which may need to operate for many hours in RF contested environments. It has been shown that the rate of divergence in dead reckoning navigation is substantial without a velocity estimate [22], particularly in consumer grade systems. Even with velocity measurements, in the presence of wind, a tactical grade aircraft may drift by 10 km per hour. In such cases, over the course of multiple hours, the precision offered by the proposed method is a vast improvement.

For loitering aircraft, this method of localization is convenient, as the flight plan need not change. For aircraft travelling significant distances, it may be sufficient to intermittently perform a single orbit, and utilize the more erroneous position estimate for course correction. The act of performing an orbit through a full compass rotation effectively nullifies the biases and offsets between the AHRS and the celestial camera. The main source of error for each independent position estimate during an orbit, is the misalignment of the estimated zenith with the true zenith. This may be caused by a number of factors, such as aerodynamic loading causing minor perturbations in the camera's orientation relative to the autopilot, improper alignment or calibration of the inertial system, or simply due to estimation errors caused by improperly estimated centrifugal acceleration. In each of these cases, it is expected that the error remains approximately constant throughout an orbit. Consequently, zenith errors at one particular azimuth are cancelled out by the zenith errors at the opposing (180° offset) azimuth. This is the reason why such a method is capable of working with almost arbitrary levels of initial error.

This flight was conducted using GPS to maintain the orbital position. We recognize that, in a true GNSS denied environment, wind errors will cause the aircraft to drift during an orbit. While this may introduce some error into the position estimate, we show in Section 4 that the primary factor governing the accuracy of a position estimate is the variance in the zenith-angle throughout the orbit. The circular position errors seen for example in Figure 7, are not caused by the change in aircraft position. They are caused by a misalignment of the optical system, and consequently may be reproduced by simply maintaining constant roll and pitch throughout a full compass rotation. In Section 4.4 we show that even in high wind conditions, the appropriate strategy is to maintain roll rate and accept that there will be positional drift throughout the orbit. This strongly suggests that the presence of GPS during the flight test offered no advantages to the algorithm.

An obvious limitation of this method is its dependence on sky visibility. Some research has shown that short-wave infrared cameras offer a daylight visible alternative to visible spectrum cameras [23,24]. These tend to have far lower signal-to-noise ratios, resulting in more erroneous measurements. It may be the case that the proposed positioning method is capable of nullifying the increased observational error from short-wave infrared observations. This may be a potential topic for future research.

6. Conclusions

This study proposed a method for obtaining more accurate positional estimates from a modular strapdown celestial navigation system. We hypothesized that by flying in an orbital motion, the errors in estimated zenith angles would cancel one another at opposing azimuths. The methodology was tested in a real flight, demonstrating that the position can be routinely estimated to within 4 km by performing orbits at a fixed altitude and airspeed. Throughout each orbit, positional estimates are generated from the individual

celestial images, and these positions are averaged at the conclusion of the orbit. We show that by recursively estimating the mean position and using this position to recalibrate the orientation of the camera, the algorithm is capable of converging near the true location. Testing found that the algorithm is robust against initial conditions, requiring no knowledge of the prior position and only requiring the camera to be aligned to within a hemisphere of tolerance.

Author Contributions: Conceptualization, S.T. and J.C.; methodology, S.T.; software, S.T.; validation, S.T.; formal analysis, S.T.; investigation, S.T.; resources, J.C.; data curation, S.T. and J.C.; writing—original draft preparation, S.T.; writing—review and editing, J.C.; visualization, S.T.; supervision, J.C.; project administration, J.C. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Acknowledgments: This work was supported by Scope Global Pty Ltd. under the Commonwealth Scholarships Program, and the Commonwealth of South Australia under the Australian Government Research Training Program.

Conflicts of Interest: The authors declare no conflict of interest.

References

- Gatty, H. Aerial Navigation—Methods and Equipment. *SAE Trans.* **1932**, *27*, 153–170. [[CrossRef](#)]
- Pappalardi, F.; Dunham, S.; LeBlang, M.; Jones, T.; Bangert, J.; Kaplan, G. Alternatives to GPS. In Proceedings of the MTS/IEEE Oceans 2001. An Ocean Odyssey. Conference Proceedings (IEEE Cat. No. 01CH37295), Honolulu, HI, USA, 5–8 November 2001; IEEE: New York, NY, USA, 2001; Volume 3, pp. 1452–1459. [[CrossRef](#)]
- Ali, J.; Zhang, C.; Fang, J. An algorithm for astro-inertial navigation using CCD star sensors. *Aerosp. Sci. Technol.* **2006**, *10*, 449–454. [[CrossRef](#)]
- Ting, F.; Xiaoming, H. Inertial/celestial integrated navigation algorithm for long endurance unmanned aerial vehicle. *Acta Tech. CSAV (Ceskoslov. Akad. Ved.)* **2017**, *62*, 205–217.
- Levine, S.; Dennis, R.; Bachman, K.L. Strapdown Astro-Inertial Navigation Utilizing the Optical Wide-angle Lens Startracker. *Navigation* **1990**, *37*, 347–362. [[CrossRef](#)]
- Chen, H.; Gao, H.; Zhang, H. Integrated navigation approaches of vehicle aided by the strapdown celestial angles. *Int. J. Adv. Robot. Syst.* **2020**, *17*, 1729881420932008. [[CrossRef](#)]
- Gai, E.; Daly, K.; Harrison, J.; Lemos, L. Star-sensor-based satellite attitude/attitude rate estimator. *J. Guid. Control Dyn.* **1985**, *8*, 560–565. [[CrossRef](#)]
- Wang, J.; Chun, J. Attitude determination using a single star sensor and a star density table. *J. Guid. Control Dyn.* **2006**, *29*, 1329–1338. [[CrossRef](#)]
- Mao, X.; Du, X.; Fang, H. Precise attitude determination strategy for spacecraft based on information fusion of attitude sensors: Gyros/GPS/Star-sensor. *Int. J. Aeronaut. Space Sci.* **2013**, *14*, 91–98. [[CrossRef](#)]
- Guo, C.; Tong, X.; Liu, S.; Lu, X.; Chen, P.; Jin, Y.; Xie, H. High-precision attitude estimation method of star sensors and gyro based on complementary filter and unscented Kalman filter. *Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci.* **2017**, *42*, 49–53. [[CrossRef](#)]
- De Almeida Martins, F.; Carrara, V.; d’Amore, R. Positionless Attitude Estimation with Integrated Star and Horizon Sensors. *IEEE Access* **2023**, *12*, 2340–2348. [[CrossRef](#)]
- Gao, Z.; Wang, H.; Wang, W.; Xu, Y. SIMU/Triple star sensors integrated navigation method of HALE UAV based on atmospheric refraction correction. *J. Navig.* **2022**, *75*, 704–726. [[CrossRef](#)]
- Teague, S.; Chahl, J. Imagery synthesis for drone celestial navigation simulation. *Drones* **2022**, *6*, 207. [[CrossRef](#)]
- Teague, S.; Chahl, J. Strapdown Celestial Attitude Estimation from Long Exposure Images for UAV Navigation. *Drones* **2023**, *7*, 52. [[CrossRef](#)]
- Zhang, H.; Zhang, C.; Tong, S.; Wang, R.; Li, C.; Tian, Y.; He, D.; Jiang, D.; Pu, J. Celestial Navigation and Positioning Method Based on Super-Large Field of View Star Sensors. In Proceedings of the China Satellite Navigation Conference, Jinan, China, 22–24 May 2024; Springer: Berlin/Heidelberg, Germany, 2023; pp. 475–487. [[CrossRef](#)]
- Van Allen, J.A. Basic principles of celestial navigation. *Am. J. Phys.* **2004**, *72*, 1418–1424. [[CrossRef](#)]
- Liebe, C.C. Accuracy performance of star trackers—a tutorial. *IEEE Trans. Aerosp. Electron. Syst.* **2002**, *38*, 587–599. [[CrossRef](#)]
- Delabie, T.; Schutter, J.D.; Vandenbussche, B. An accurate and efficient gaussian fit centroiding algorithm for star trackers. *J. Astronaut. Sci.* **2014**, *61*, 60–84. [[CrossRef](#)]

19. Teague, S.; Chahl, J. Bootstrap geometric ground calibration method for wide angle star sensors. *JOSA A* **2024**, *41*, 654–663. [[CrossRef](#)]
20. Wei, X.; Zhang, G.; Jiang, J. Star identification algorithm based on log-polar transform. *J. Aerosp. Comput. Inf. Commun.* **2009**, *6*, 483–490. [[CrossRef](#)]
21. Lawrence, J.; Bernal, J.; Witzgall, C. A purely algebraic justification of the Kabsch-Umeyama algorithm. *J. Res. Natl. Inst. Stand. Technol.* **2019**, *124*, 1. [[CrossRef](#)]
22. Gebre-Egziabher, D.; Powell, J.; Enge, P. Design and performance analysis of a low-cost aided dead reckoning navigation system. *Gyroscopy Navig.* **2001**, *4*, 83–92.
23. Wang, W.; Wei, X.; Li, J.; Zhang, G. Guide star catalog generation for short-wave infrared (SWIR) all-time star sensor. *Rev. Sci. Instrum.* **2018**, *89*. [[CrossRef](#)]
24. Ni, Y.; Wang, X.; Dai, D.; Tan, W.; Qin, S. Adaptive section non-uniformity correction method of short-wave infrared star images for a star tracker. *Appl. Opt.* **2022**, *61*, 6992–6999. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.