Photogrammetry in Light Aircraf

Sérgio Cunha*,**, Sérgio Silva*, F. Piqueiro*

(*) University of Porto, Faculty of Engineering, FEUP
(**) Centro Integrado de Investigação Marinha e Ambiental, CIIMAR

Keywords: Photogrammetry, GPS, INS.

Abstract

Digital SLR off-the-shelf photographic cameras have achieved performance levels today that compare to professional photogrammetric equipment, at a fraction of the cost. Together with navigation equipment (GPS and inertial) and a light aircraft, they offer an attractive perspective for high resolution and local area airborne photogrammetry applications. In this type of applications, equipment and flight costs are particularly relevant. This text presents the setup, data processing requirements and results of such a system.

1. Introduction

Airborne photogrammetry has been typically targeted at surveying areas of thousands of square kilometres with terrain resolution of tens of centimetres per pixel. Although this balance between covered area and resolution is adequate for many applications, there are emerging markets for higher resolution surveys of local areas: resolutions of few centimetres per pixel of areas of less than one square kilometre. Such applications include terrain surveying for building, road and other infrastructure construction, coastal monitoring for erosion and vegetal species classification. The setup that proves most competitive for larger scales is not ideal for local area surveys. The former requires flight time efficiency with respect to the covered area. The latter is more concerned with efficiency with respect to total flight time, aircraft manoeuvrability and guidance accuracy. In particular, local area surveys are particularly suited for light aircraft.

The need to cut down flight costs implies also that equipment installation and setup should not depend on the aircraft as much as possible. This enables to use different aircraft, avoiding incurring on costs related to complex installation or simply dependence on the aircraft. The developed system was designed respecting the constraint of not requiring any structural changes in the aircraft. The only change to the aircraft was the removal of a door to allow shooting photographs from the cockpit, as explained in the next section.

2. System Architecture

The proposed system is constituted by two fundamental subsystems that are supported by the same common navigation system, running on a rugged embedded PC: the flight guidance system and the image acquisition system (figure 1).

Figure 1 – Aircraft and system functional diagram.

The navigation system receives data from three (or more) GPS antenna/receiver pairs, installed in different positions in the aircraft, logs its data and processes it to create a navigation solution (aircraft position and attitude) available for use by the remaining subsystems. These receivers must be able to output raw data, including carrier phase measurements, which by differential processing enable the computation of the aircraft attitude (heading, pitch and roll angles) and precise absolute position relative to a reference station. The purpose of computing attitude is to provide a means to migrate the computed positions of one of the receivers to the location of the inertial sensor. The receivers only need to acquire L1 measurements, although there is a significant advantage if one of them is also able to receive L2. By keeping a tight synchronization with GPS time through the GPS data messages carrying time information and, more importantly, the PPS signal [1], the navigation system also serves as a common time reference system, enabling time tagging of various events and their correspondence to the correct GPS measurements.

All time events are connected to specific hardware interrupt pins available on the embedded PC I/O ports and so their time stamps can be accurately measured through customized operative system drivers. This guaranties low latency and consequently a small time offset and low variance in the measurements.

The flight guidance system is an application developed within the framework of this project using the OpenGL graphic libraries (figure 2). It is constituted by a horizontal situation indicator (HSI) and a flight director that are displayed on a small portable monitor located above the aircraft control panel. This system uses the calculated navigation solution and previously supplied flight profiles to calculate a course error that is conveyed to the pilot through a set of visual indicators. At any instance the pilot knows his track and altitude error making possible to follow profiles within the necessary error interval. In this visualization system there is also information to control the shooting of photographs.

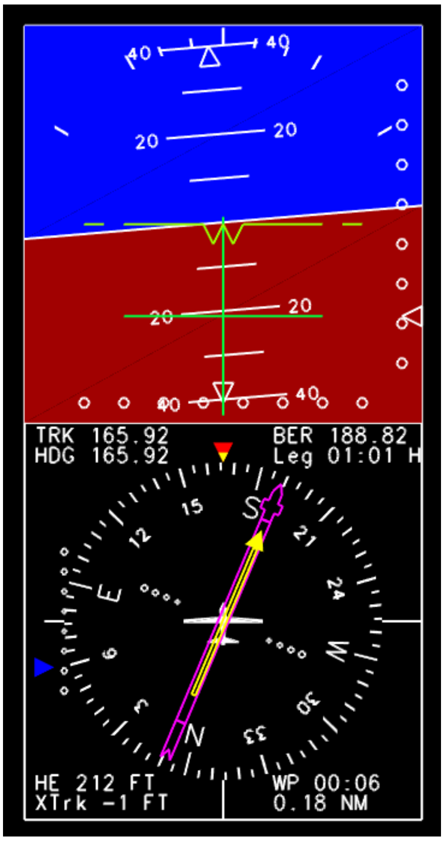

Figure 2 – Flight Guidance Display.

The pictures are taken using a high-end digital single lens reflex (D-SLR) camera that is able to acquire high resolution pictures (17 mega pixels) at high frame rates (several frames per second at continuous shooting). Each time a picture is taken a trigger signal is fed to the embedded PC and a time tag is recorded.

The IMU data related to the camera orientation is logged together with the time stamps relative to each measurement. Its main purpose is to compute the camera attitude, through integration of its acceleration and angle rate measurements with the GPS derived positions. This enables a posteriori correct positioning and attitude determination of the camera and, consequently, of the imaged ground area. It plays also a secondary role in refining the camera position estimates.

The camera is mounted on a specially designed support that is sustained on the camera operator leg (therefore only loosely attached to the aircraft body). Besides enabling a rigidly connection between the IMU and the camera, this support allows a correct and steady, but still versatile, camera positioning throughout the survey mission. This has the significant advantage of making the whole system aircraft independent. It has been observed that a human operator is more stable than a light aircraft in what regards pitch and roll angle oscillation. This approach has the drawback of being subject to positioning errors due to shifts between the GPS antennas and the inertial device. Although the absolute maximum shift is bounded (less than 10 cm typically) and has negligible impact by itself, relative motion errors between consecutive photographs may have a more sensitive effect in stereomatching in high resolution surveys. This is attenuated by the inertial device, as it provides better relative position estimates at the higher segment of the frequency spectrum.

A keyboard and a monitor constitute the flight control console which enables an operator to control the current mission flight profile and other system parameters.

As described, this system can use a small aircraft for operation. There are no hardware changes necessary to the aircraft body or equipment other than the removal of the right door. Although it might seem strange, this is a common and easy procedure in this type of aircraft. All the equipment necessary for each mission is easily assembled before flight. This enables low cost operation and maintenance. Figure 3 illustrates this description.

Figure 3 – Mission operation.

Besides the equipment installed in the aircraft a reference station is also established in the surveyed for each mission, enabling differential processing of the GPS receivers on board the aircraft relative to the ground station enhancing their precision.

3. Navigation Data Processing

Navigation data processing precedes photogram processing. The goal of this task is to obtain camera position and attitude estimates at the instants when the photograms were shot. The inertial sensor is the instrument that is responsible for computing the camera/IMU block final position and attitude values. This is obtained from integration of inertial data with GPS measurements, as explained further below. Given the accuracy level required, such GPS measurements must relate or be relatable to the IMU location within the aircraft. In order to migrate measurements from the GPS antennae location using the known arms between these devices, the aircraft body attitude must be known. It was found that the camera/IMU block, being only loosely attached to the operator's leg, rotates significantly relative to the body of the aircraft (10º is not surprising). Therefore, such migration cannot rely on the IMU derived attitude. The set of GPS antennae/receivers are processed in differential mode for this purpose, using the carrier phase measurements.

The algorithm used for ambiguity fixing of L1 measurements between antennae relies on the integer quasi-orthogonalization of the search space as proposed by Lenstra, Lenstra and Lovász [2], [3] in the early eighties, which became famous in the GPS world as the lambda method [4], to significantly reduce the search space. At the heart of fixing ambiguities is the problem of finding and classifying the admissible set of integer solutions corresponding to the set:

$$\left\{ z \in \mathbb{Z}^m : \min_{x \in \mathbb{R}^n} \|y - Ax - Bz\|_{Q^{-1}} \leq r \right\}, \tag{1}$$

where x is the vector of real unknowns (differential positions) in $\mathbb{R}^n$, z is the vector of integer unknowns in $\mathbb{Z}^m$ (carrier phase ambiguities), A and B relate x and z to the observations y, Q is the covariance of the equation error μ given by $\mu = y - Ax - Bz$ and r is a value that defines the size of the search space in terms of the total probability of the solutions that define the set (1). The value of r is typically between 3 and 5. The set of integer vectors that satisfy (1) is then sorted by ascending values of the residue. In fact, assuming the error has Gaussian distribution, a relative probability if attributed to each solution candidate. The comparison of the probabilities of the different candidates drives the choice of the proper solution as, ideally, the candidate that holds by itself a relative probability of 90% or more.

The individual equations that compose $y = Ax + Bz + \mu$ can be divided into two groups: the carrier phase observations (double differences between receivers and satellites) with real and integer unknowns and other observations with real unknowns only. Traditionally the latter are based on pseudo-range observations (again using double differences between receivers and satellites).

However, it was found that for low cost receivers the noise figure of these observables is too high for effective ambiguity resolution. In a light aircraft, a rough estimate of the attitude angles derived from the velocity vector and its evolution with time will provide a better estimate of the vectors between devices in Earth fixed coordinates, which compose the vector x. In fact, for arm lengths of less than 2 meters and attitude estimation errors of up to 30º (which will clearly an inflated upper bound), the search space for ambiguities has a radius of less than one meter, which is better than what is typical for pseudo-range measurements for low-cost OEM GPS receivers. The use of a flight model that includes wind estimates further decreases this error figure, allowing for more confident estimates.

Several epochs are stacked in the equations, sharing the same integer variables if no cycle slips are detected. This significantly reduces the set of candidates, increasing the discriminating power of the algorithm. In fact, a new epoch is added to the set of equations whenever no single solution is clearly identified as the proper choice. Upon solving ambiguities confidently, the use of triple differences of carrier phase measurements allows for immediate fixing of the ambiguities in the following epochs. The procedure is restarted on the event of total or significant loss of lock between consecutive epochs. Knowledge of the vector arms between antennas further helps assuring the right solutions have been selected.

There is a minor detail in the GPS measurements that has to taken into account for proper attitude computation: the fact that all receivers collect measurements at different instants in time for the same epoch. Differences of 1 msec are not unusual. Such gap induces position differences of several centimeters for typical light aircraft velocities. A simple approach to overcome this detail is to migrate the raw measurements closer to the exact epoch, using Doppler readings and the receiver clock offset estimate obtained from the stand-alone processing of its pseudo-range measurements. After migration, the time difference is in the order of the microsecond, inducing negligible positioning errors.

One of the receivers is also used to compute absolute positions, again through carrier phase differential processing, this time relative to a reference station installed in the flight area. The ambiguity-free measurements are now based on the pseudo-range readings, at the lack of any better alternatives. There are two main sources of errors in solving ambiguities for absolute aircraft positioning: pseudo-range measurement errors and atmospheric effects. The latter affect both pseudo-ranges and carrier-phase measurements. Therefore, it is difficult to assess if the obtained solution is correct. This is the main reason for installing a reference station in the area to be surveyed, making the baselength between reference and rover stations as small as possible. The use of one L1/L2 receiver in the aircraft and in the reference station is also a significant (but costly) improvement on the ability to successfully and reliably fix ambiguities.

With the help of the vector arms between aircraft antennae (aircraft attitude), the obtained rover antenna position estimates are migrated to the location of the center of navigation of the IMU. This location is not absolutely stable, as the IMU/Camera block is only loosely tied to the camera operator. However, it was observed that the translation relative to the measured position does not exceed few centimeters. The inertial device raw measurements (angular velocities and accelerations along three orthogonal axes) are now blended with these position estimates in a Kalman filter. The Kalman filter state variables are the position, velocity and attitude errors of the integrated navigation solution, plus the six sensor biases (total of 15 states). The observations are the GPS derived positions, together with their covariance, which contains the estimate of the GPS measurement error plus the unknown shift of the IMU/Camera block around its nominal point.

Experiments using the raw carrier phase measurements, properly compensated from the integer number of cycles already computed and migrated to the IMU position, have shown minor or negligible improvements on the obtained solution. In other words, the option for a tightly coupled Kalman filter approach was not found to be particularly rewarding in this case. The fundamental reason is the availability and roughly even distribution of the satellites in the sky, given the lack of obstructions in the airborne environment. Other reason is the option for a low cost inertial device, which masks any difference between using explicit (raw data) or aggregate measurements (positions). Finally, the primary outcome of the filter is the platform attitude, which is significantly insensitive to the choice between the two approaches.

The causality nature of the Kalman filter induces errors to propagate in time, affecting the accuracy of the lower parts of the spectrum of frequencies of the results. Such errors include the attitude estimates: the offsets of the results relative to the correct values will exhibit a slow variation. It is possible to mitigate this effect to some extent by further processing the data using the *data smoothing* paradigm instead of the Kalman filter. The estimate at each instant will be based on both past and future data. The experiments conducted in the scope of this work did not show, however, that this further step induced a visible improvement in the photogram data processing steps that follow.

The position estimates of the Kalman filter are strongly dominated by the GPS measurements. This is due to the quality of the latter as a system (especially in what regards relative motion measurement) relative to the accuracy of the employed inertial device. Differences regard mostly the higher frequency components of the spectrum. They are primarily related to sudden translation shifts of the handled platform. Naturally, the inertial device introduces significant improvement on this part of the spectrum, especially considering the much higher sampling rate.

The resulting position and attitude estimates of the Kalman filter finally suffer the obvious transformation of translation from the IMU center of navigation to the camera focal point and rotation from the IMU native orientation to the camera orientation. Alignment errors correction will be required, as addressed within the next section, as the platform has necessarily construction imperfections and might suffer torsion from flight to flight. Finally the results are interpolated at the instants when the photographs were shot (these instants are registered in the onboard computer in real time during the flight). The result is a file with the camera position and attitude estimate for each photogram to be processed.

4. Model Construction

This section addresses the problem of building and computing a model of the world as measured by the shot photograms. Such model is composed by the refined camera (focal point) position and attitude values, together with the coordinates of the elements used to link different photograms (tie-points) or photograms to terrain points (ground control points). It is the result of a minimization procedure corresponding to the overdetermined system of equations containing information regarding camera position and attitude (the file resulting from the procedures described in the previous section), tie-points and ground control points.

Before applying any other process to the photograms, these have to be compensated by the camera and lens distortion. The distortion parameters are easily computed by specific software, which is based on the comparison of a large number of targets in a series of photographs that capture these targets from different angles. The next step is to compute the alignment error between the IMU and the camera. Although this error includes both positional and rotational components, only the latter relevantly affect the following steps, requiring compensation.

Knowledge of the camera position and two separate ground control points in a photogram is enough to derive the full orientation of the camera. Further ground control points provide more accurate and, moreover, more reliable results. Comparison of the obtained angles to those obtained from the navigation data processing provides the 3D rotation compensation vector. The use of more than one photogram for this purpose is helpful to, again, increase precision and reliability.

At this point, there is a set of photograms accompanied by camera position and attitude information. Such estimates are, to the knowledge available at this point, unbiased and with error magnitudes in the order of the decimeter for position and one tenth of the degree for attitude (merely indicative figures for the accuracy levels of the employed equipment). The position error can be higher in case of failure to determine the right ambiguities, which is prone to happen in the case of high attitude flights with L1 only receivers. For a typical mission flown at 1000 meters above ground level, such errors correspond roughly to an error of a couple of meters in the terrain. This translates to less than 100 pixels of uncertainty in the definition of the radius of the search space for common points between photograms.

An a priori estimate of the digital terrain model (DTM) is also required to perform such search. The accuracy requirement for such DTM is not severe, mainly due to the use of high focal lengths. This is a secondary advantage of using a small aperture for the photographs. The primary advantage is to obtain photographs that are, from the beginning, close to vertical, requiring thus simpler transformation in the process of generating the orthophotos.

The next procedure is to compute tie-points between adjacent photograms automatically. This is a computer intensive procedure, based on the comparison of segments of pairs of photographs; however, the camera position and attitude estimates are helpful to significantly reduce the time consumed, as the search space is small. A large number of tie-points are chosen, as they will be used to compute the final DTM, as explained further below. A number of ground control points are added for both providing absolute positioning information at ground level and for quality control. These ground control points typically form a very sparse grid that covers the survey area (for instance, one in every corner and another one in the center).

As the computation of the model includes linearization of the mentioned equations, the process is iterative. During this process, it is not unusual to find incorrectly determined tie-points (false tie-points). These can be noticed by the large residues of the solution, especially in the equations related to these tie-points, and/or by the unlikely values of the height components of these tie-points on the solution. These tie-points have either to be removed of corrected. Since this step is performed by direct human intervention, it can be resource consuming. By reducing the search space significantly, the use of attitude information also causes the number of false tie-points to be small compared to the general case when only a rough estimate of the camera location is employed.

Upon cleaning the data set of all mismatch tie-points and after obtaining convergence of the model (the residues are small), new camera positions and orientations are obtained. The improvement on the positions is negligible. The equations are modestly sensitive to the camera positions values, especially when high focal lengths are employed. The improvement on the camera orientation values is, on the other hand, significant. The new positions and orientation values will be used to compute the orthophotos from the photograms, after a more precise DTM is extracted.

5. Orthophoto and DTM Generation

Stereomatching is the technique employed to obtain a more precise DTM, using the photograms and the model computed as explained in the previous section. There are two main approaches to accomplish this task. One is to generate epipolar pairs of photograms (using the position and orientation values) and then to compute the DTM from comparison of the photograms at the pixel level. Although this technique generates high resolution terrain models, matching problems relegate its use to only certain types of surveys.

The experience collected so far points towards three main sources of errors related to the motives of the surveyed area: existence of steep vertical gradients, moving objects and repeating patterns. The first one includes constructions whenever their height is perceptible from the photograms. This is very common in high resolution surveys of areas that suffered human intervention, especially urban areas. Moving objects includes cars, water or even people. Repeating patterns are likely to occur in urban areas (roofs with tiles are a typical example) and are generally troublesome in the presence of other objects with different heights. The result of these matching problems is the appearance of spikes and noise in the DTM, which induce a blurring effect on the obtained orthophotos. Figure 4 exhibits a piece of an orthophoto with an example of such phenomenon. Although it represents an effort towards achieving a pure vertical photograph, the result is, from an aesthetical point of view, not acceptable.

Figure 4 – Orthophoto with blurring effect.

The alternative is to use a smoother DTM. This is obtained from the tie-points computed as described in the previous section. They have the advantage of being correct from the point of view of matching between photographs. On the other hand, a smoother DTM can only be used under the assumption the photographs will be corrected from the distortion induced by the terrain height only to a certain level of detail. For instance, in an urban area, there will generally not be a correction to compensate for the height of the buildings. Therefore, there is the need that the original photographs are already close to vertical. Furthermore, this justifies the use of high focal lengths.

When building a DTM from tie-points, which has a low resolution, it is important that the (smooth) surface relates to the same object being characterized. For instance, since buildings cannot be separated from ground to the detail required to compensate for their heights, the DTM must be based on tie-points that correspond to points in the ground (and not in the roof of buildings). This leads to the need to select between tie-points. So far, this task has been performed manually, at the cost of time consumed in human labour. A few automatic techniques are now in the process of being essayed. These are based on the height of the tie-points relative to its neighbours. This and other similar techniques are not expected to be 100% effective.

Given the DTM and the position and orientation model of the camera for the different photograms, the next procedure if to compute the orthophotos. Finally, the orthophotos are brought together in mosaics. The separation line between orthos is generally computed automatically, taking advantage of areas where the difference between the overlapping orthos is minor or the contrast in the ortho themselves is higher. In the latter case, it is relevant that the line does not cross areas situated at an height considerably different from that of the DTM. Again, some humn intervention is generally required. Figure 5 illustrates the same area of figure 4, using the DTM obtained from tie-points. The blurring effect is not present. On the other hand, the façades of the buildings can be noticed. These have little impact for most of the applications.

Figure 5 – Orthophoto built with DEM from tie-points.

6. Conclusion

This system has been flown both in high resolution/local area (2 cm per pixel in areas of about 20 ha) and in lower resolution/regional area surveys (50 cm per pixel / 1000 sq Km). Although competitive in both applications, it has proven its particular suitability for the former. It has also proven that the extra burden of processing a larger number of photographs provided by an off-the-shelf camera is tolerable in the presence of position and attitude data, since automatic procedures are employed to minimize human intervention. The larger number of photographs has the advantage of, through the use of a larger focal distance, providing better verticality to the photographs. This is beneficial for the orthorectification process and produces better final results, especially in urban areas.

Acknowledgement

This research project has been supported by the Department of Electrical Engineering of Porto University, for which the authors express their gratitude.

Bibliography

[1] Silva, S., Chilro, R., Cunha, S., 2001. Building and Inexpensive Time Reference System using an OEM GPS Receiver and RT-Linux. ION GPS 2001, Salt Lake City, USA.

[2] Lenstra, A.K., Lenstra, H.W., Lóvasz, L., 1982. Factoring polynomials with rational coefficients. Mathematische Annalen, 261: 515-534.

[3] Hassibi, H., Boyd, S., 1996. Integer parameter estimation in linear models with GPS applications. Proceeedings of the IEEE Conference on Decision and Control, Kobe, Japan.

[4] Teunissen, P.J.G. 1995. The least-squares ambiguity decorrelation adjustment: a new method for fast GPS integer ambiguity estimation. Journal of Geodesy, 70 (1-2): 65-82.