SEQUENTIAL BUNDLE ADJUSTMENT
USING KALMAN FILTERING
AND OPTIMAL SMOOTHING

A. Mohamed[1], A. Benjamin, B. Wilkinson

Geomatics Program, University of Florida,
304 Reed Lab, Gainesville, FL 32611-0565 amohamed@ufl.edu

KEY WORDS: Bundle Adjustment, Kalman Filtering, Vision Aiding

ABSTRACT:

Vision aiding of the navigation solution has become an integral component of low-cost IMU/GPS sub-systems providing direct georeferencing to remote sensing systems. The data workflow to recover the orientation parameters rigorously requires the simultaneous handling of large amount of imagery and navigation data. In some situations, with small unmanned aerial vehicles for example, a flight block of thousands of images is the norm. The normal matrix of the blended imagery and navigation data can be very large in size for regular computers to handle efficiently. We use a Kalman filtering approach to sequentially process the blended navigation and imagery data; georeferencing parameters are then computed for every exposure station. In overlapping areas of the imagery, the exposure stations and the overlapping object are coplanar; this forms the general update equation of the Kalman filter. To rigorously account for the simultaneous optimal solution of the state parameters, we backward smooth the filtered estimates using the stored covariance information. We solve the problem in a form of overlapping strips in two directions to account for the whole block of imagery. We also account for the hybrid nature of the observation equation formulation which has mixed observations and parameters through creating equivalent condition equations and use the general least-squares approach. We use this technique on imagery collected by a small unmanned aerial vehicle used in environmental research. The small format of the imagery resulting from the low flying altitude produces large amount of images per flight mission. Because of the limitation on the vehicles payload, a lightweight MEMS-based inertial unit augmented by low-cost precise GPS is used to directly geo-reference the acquired imagery. The benchmarked accuracy of the attitude information from the inertial unit is in the order of half a degree root-mean-squared error. Simulation results show the possibility of improving the results by at least a factor of two through using image aiding. Besides, the condition of coplanar exposure stations and overlapped objects creates tighter relative models between the different images and between the different strips, resulting in a tighter adjustment of the whole block. The proposed technique should, not only improve the accuracy of the image block, but also improve the algorithm computational efficiency drastically.

[1]Corresponding author

1. INTRODUCTION

1.1 Background

Georeferencing is the process of acquiring knowledge concerning the origin of an event in space and time. Depending on the sensor type, this origin is defined by a number of parameters such as time, position/location, and attitude/orientation. When this origin information is determined directly by on-board vehicle sensors, the term direct georeferencing is utilized. This has been a rapidly developing area of concern for geospatial experts including navigation, remote sensing, and photogrammetry specialists.

To properly direct geo-reference imagery, the position and orientation of the sensor at exposure times must be known. These elements are commonly referred to as exterior orientation parameters (EOP). The position of the sensor is often represented by the incident nodal point of the lens (X_L, Y_L, Z_L) [Wolf and Dewitt 2000]. These parameters are commonly obtained by calculating offsets between an on-board Global Positioning System (GPS) unit and the camera lens. The attitude parameters of the imaging sensor are often derived from an on-board inertial measurement unit (IMU) mounted rigidly to the same frame as the imaging sensor. The attitude angles can be represented by omega/phi/kappa (OPK), azimuth/tilt/swing (ATS), or heading/pitch/roll (HPR). The OPK convention will be implemented herein. OPK involves sequential rotations about the x, y, and z axes. Specifically, a positive omega (ω) rotation is a counterclockwise rotation about the x-axis; a positive phi (φ) rotation is a counterclockwise rotation about the once rotated y-axis; a positive kappa (κ) rotation is a counterclockwise rotation about the twice rotated z-axis.

1.2 Aerial Triangulation (AT)

In in-direct geo-referencing, aerial triangulation is used to obtain the exterior information parameters of individual exposure station using known ground point coordinates and its counterparts in the images. The other interest of AT is obtaining the object space coordinates of other points imaged in the aerial images.

Bundle Adjustment (BA) is a technique used in analytical AT to obtain EOPs. BA uses a least squares approach to minimize the errors of a bundle of rays connecting the photo coordinate measurements with the control coordinates on the ground. This is usually done through incorporating either the co-linearity equations to determine the EOPs through determining the object space coordinates of the imaged points or directly through the co-planarity condition [Kersten and Baltsavias 1994; Tommaselli and Tozzi 1996; Haala et al 1998; Wolf and Dewitt 2000; Mikhail et al 2001; Wang and Clarke 2001].

Co-linearity is the condition that the incident nodal point of the lens (exposure station), any object point, and its corresponding image point all lie on the same line in three-dimensional space [ibid]. Co-planarity, on the other hand, is the condition that two exposure stations of a stereo-pair, any object point, and the corresponding image points of the two photos all lie in the same plane [ibid]. Since no object coordinates are involved in the co-planarity condition, initial approximations of the object space coordinates are not necessary. This proves to be a valuable asset of the algorithm developed herein.

Bundle adjustments can be performed either sequentially or simultaneously. Each method has three main components: relative orientation of each stereo-model, connection of adjacent models to form continuous strips and/or blocks, and simultaneous adjustment of the photos from the strip/block to ground control. Relative orientation is used to determine the relative angular attitude and positional difference between two photographs when the images were captured [ibid]. Absolute orientation takes the relatively oriented stereo-models and transforms them to the ground using three-dimensional conformal coordinate transformations. The unknown

quantities of BA in either method are the object space coordinates (XYZ) of object points and the exterior orientation parameters (georeferencing parameters) of each photograph.

Advantages and disadvantages are prevalent for the use of either sequential or simultaneous bundle adjustments. The greatest disadvantage for sequential bundle adjustments is the nonlinear accumulation of random error along an image strip as more stereo-models are added to the adjustment. Simultaneous BA avoids this error accumulation by processing all measurements at once; thus providing a more robust method for determining the optimal solution. Simultaneous BA, however, comes with a computational burden in the form of huge matrix operations imposed by the large amounts of imagery in the strips/blocks, specifically with small format imagery. If a method existed for reducing error accumulation in sequential BA, the computational time incurred by doing the adjustment sequentially as opposed to simultaneously could be a great advantage especially in near real-time applications. This has made sequential estimation an ongoing research topic of interest for many experts especially in the field of navigation.

Many approaches to sequential estimation through bundle adjustments are in the field of robot vision or vision metrology [Kersten and Baltsavias 1994; Edmundson and Fraser 1998; Di et al 2008]. To accurately obtain the orientation and position of the robot, a simultaneous bundle adjustment of all the previous geospatial data would not be feasible as the robot/UAV needs the geospatial information in near-real time to continue navigation. Thus, sequential estimation theory is used. Almost all approaches incorporate image-matching techniques utilizing co-linearity equations [Kersten and Baltsavias 1994; Tommaselli and Tozzi 1996; Haala et al 1998; Wang and Clarke 2001]. The algorithm developed in [Webb 2007], however, applies co-planarity as the observation model for Kalman filtering to the sequential aero-

triangulation problem with success. The premise of that research was navigation not georeferencing the acquired imagery. Thus, our research analyzes Kalman filtering with an optimal smoother in sequential BA as a method for reducing the sequential accumulation of error normally associated with AT to provide accurate and precise geo-referencing parameters.

1.3 **Kalman Filtering**

The Kalman Filter (KF) is an algorithm that implements a predictor-corrector estimator to minimize the estimated error covariance of the state [Gelb 1974]. KF achieves that by utilizing: knowledge of system and measurement dynamics, assumed statistics noises and measurement errors, and initial condition information [ibid]. Kalman filtering is the most common technique for estimating the state of a linear system and is widely used in many applications such as navigation with INS-GPS systems [Nassar et al 2007; Webb 2007], satellite orbit prediction [Xiong et al 2009], and in many other fields.

The objective of the Kalman filter is to obtain the system state estimate $\hat{x}_k$ (*a posteriori* state estimate) as a linear combination of a predicted estimate $\hat{x}_k^-$ (*a priori* estimate) and a weighted difference between an observation z_k and a measurement prediction $(H\hat{x}_k^-)$. In equation form,

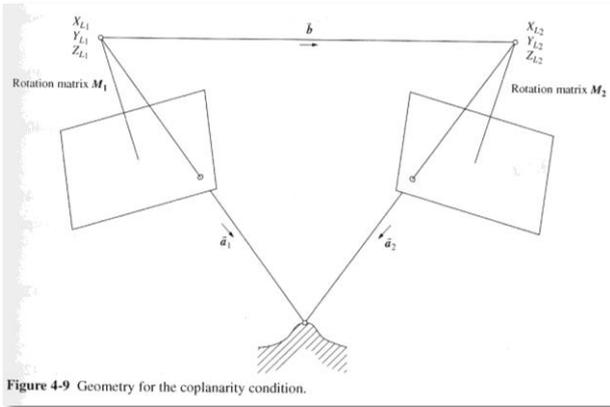$$\hat{x}_k = \hat{x}_k^- + K(z_k - H\hat{x}_k^-) \qquad (1)$$

Where

K is the Kalman gain that minimizes the a posteriori error covariance and

$(z_k - H\hat{x}_k^-)$ is the measurement innovation or residual.

Interested readers are referred to Kalman Filtering text such as [Kalman 1960; Rauch et al 1965; Gelb, 1974; Welch and Bishop 2001] for details of the algorithm.

1.4 Co-planarity Condition

We use of the co-planarity condition (CC) as the observation model in the Kalman Filter algorithm. If two photographs are relatively oriented with respect to each other, then the object space rays defined by a pair of conjugate image points and their respective exposure stations will intersect at exactly one point [Wolf and Dewitt 2000; Mikhail et al 2001]. The object space position of that point occurs at the intersection. The two object space rays in combination with the position vector connecting the two exposure stations form the three sides of a triangle. This triangle defines the plane for which this condition satisfies; see a depiction of the condition in figure (4-9) from [Mikhail et al 2001] below.

Figure 4-9 Geometry for the coplanarity condition.

The co-planarity condition is based on the fact that the volume (V) of the parallelepiped (a polyhedron consisting of all parallelogram faces) of three co-planar vectors will be 0 as shown in equation 2 below.

$$V = \left| \vec{b} \cdot (\vec{a}_1 \times \vec{a}_2) \right| = \begin{vmatrix} b_x & b_y & b_z \\ u_1 & v_1 & w_1 \\ u_2 & v_2 & w_2 \end{vmatrix} = 0 \quad (2)$$

As illustrated in the figure, the base vector between the two exposure stations is $(\vec{b})$. Vectors, $(\vec{a}_1)$ and $(\vec{a}_2)$, are the object space rays from the exposure stations through their respective conjugate image points to the common object space point.

To further develop the observation model, the co-planarity condition takes the observation equation form [Mikhail et al 2001]

$$Z_e = \begin{bmatrix} x_1 & y_1 & z_1 \end{bmatrix} F \begin{bmatrix} x_2 \\ y_2 \\ z_2 \end{bmatrix} = 0 \quad (3)$$

Where

Z_e is the equivalent observation, and

F is the fundamental matrix that contains the exterior orientation parameters (EOPs) and the interior orientation parameters (IOPs) of both images in the stereo-model, which take the form

$$F = C_1^T M_1 K_b M_2^T C_2 \quad (4)$$

Where

C is the calibration matrix that contains the interior orientation parameters in the form

$$C = \begin{bmatrix} 1 & 0 & -x_0 \\ 0 & 1 & -y_0 \\ 0 & 0 & -f \end{bmatrix} \quad (5)$$

And

K_b is the skew-symmetric matrix that contains the base vector information between the two exposure stations

$$K_b = \begin{bmatrix} 0 & -b_z & b_y \\ b_z & 0 & -b_x \\ -b_y & b_x & 0 \end{bmatrix} \quad (6)$$

And

M_i is the rotation matrix of image i in the stereo-pair

$$M_i^T = \begin{bmatrix} m_{11} & m_{21} & m_{31} \\ m_{12} & m_{22} & m_{32} \\ m_{13} & m_{23} & m_{33} \end{bmatrix}_i \quad (7)$$

The components of these matrices as well as the details of the co-planarity equations are discussed in photogrammetry books such as [Wolf and Dewitt; Mikhail et al 2001] and will

not be repeated here.

2. METHODOLOGY

2.1 Observation Equation

The objective of this development is to test the impact the co-planarity condition can have as the observation model for a Kalman filtering approach to solving for the georeferencing parameters (EOPs). Simplification of the dynamics model was desired to isolate the impact that this particular observation model had. Thus, certain assumptions were made that affected the time update equations in the Kalman filter algorithm. First, it was assumed that the *a priori* state estimate $(\hat{x}_k^-)$ is equal to the *a posteriori* state estimate from the previous epoch $(\hat{x}_{k-1})$. Thus, the original linear discrete-time controlled process Kalman filter equation mentioned earlier is simplified from

$$\hat{x}_k^- = \hat{x}_{k-1} \qquad (8)$$

In a similar fashion,

$$P_k^- = P_{k-1} \qquad (9)$$

Furthermore, the measurement update equation for the Kalman gain is rewritten below as

$$K_k = P_k^- H^T (H P_k^- H^T + R_e)^{-1} \qquad (10)$$

Where
R_e is the equivalent measurement noise covariance matrix formed by the measurement noise covariance matrix R and the Jacobian matrix of the co-planarity equation with respect to the image point observations B

$$R_e = BRB^T \qquad (11)$$

Likewise, the observation z_k is the equivalent to Z_e from the linearized form of the co-planarity equation. Substituting into the second discrete Kalman filter measurement update equation to find the a posteriori state estimate

$$\hat{x}_k = \hat{x}_k^- + K(Z_e - H\hat{x}_k^-). \qquad (12)$$

By integrating these assumptions with the observation model modifications, the Kalman filtering process can be investigated via simulation. To make the exterior orientation parameter solution more robust than a simple forward filtering process, an optimal smoothing technique is also implemented.

2.2 Optimal Smoothing

Optimal Smoothing (OS) does not require observation model; thus, reimplementation of the co-planarity condition observation equation is not necessary. Smoothing is a function of stochastic weighting only. It is a post-mission processing scheme that uses all measurements from the initial epoch to a time N to estimate a system state x_k^s at epoch k such that $0 \leq k \leq N$. In the case of this algorithm, the *a priori* and *a posteriori* system state estimates as well as the *a priori* and *a posteriori* covariance estimates from the forward Kalman filtering process for each epoch between 0 and N are stored. These estimates form the basis of the backwards smoothing operation. In theory, backward OS accounts for the shortcoming of the forward filtering algorithm to achieve optimal solution equivalent to the simultaneous batch processing of all the data.

While near real-time processing is desired, the initial implementation of this algorithm is to simplify the model as much as possible without sacrificing optimization. There are three main types of optimal smoothers: fixed-interval, fixed-point, and fixed-lag. This model implements a fixed-interval smoothing algorithm called the Rauch-Tung-Striebel (RTS) backward smoother. The RTS smoother is the least complex fixed-interval optimal backward smoother [Gelb 1974; Nassar et al 2007].

The backward sweep of the RTS commences at the culmination of the forward Kalman filter

sweep. At this point, the initial smoothed system state estimate x_{K+1}^s is equal to the *a posteriori* system state estimate $\hat{x}_N$. Likewise, the initial smoothed covariance estimate P_{K+1}^s is equal to the *a posteriori* system state estimate P_N. The smoothed system state estimate x_K^s at time k in the RTS algorithm [Rauch et al 1965]

$$x_k^s = \hat{x}_k + D_k(x_{k+1}^s - \hat{x}_{k+1}^-) \qquad (13)$$

Where

D_k is the smoothing gain matrix (similar to the Kalman gain matrix from the forward filtering process)

$$D_k = P_k I (P_{k+1}^-)^{-1} \qquad (14)$$

Note that the covariance matrix of the smoothed estimates P_k^s is not necessary for computing the smoothed system state estimate in our case. However, analysis of that covariance matrix was deemed necessary for determining the relative precision of the algorithm.

$$P_k^s = P_k - D_k(P_{k+1}^- - P_{k+1}^s)D_k^T \qquad (15)$$

3. **EXPERIMENTAL RESULTS**

To test the functionality of this algorithm, a simulation model was created to mimic imagery acquisition over a predetermined flight path. While the end goal is the utilization of this algorithm with thousands of images per flight, the initial simulation model was a strip of 25 photographs. Over this minimal period, both the forward filter and the backwards optimal smoother could be seen converging on a steady-state solution. Thus, enlarging the test strip at this point would only contribute additional redundancy to the steady-state solution.

Standardization of the simulations was necessary for comparison across trials. To do this, a seeded random number generator was utilized to perturb the original inputs by altering their respective observation standard deviations (precision). A simple structure was designed to simulate and standardize an image matching algorithm based upon the desired number of tie points between the overlapping images. Recall that tie points are crucial for the proper implementation of the co-planarity condition as the observation model.

The focal point of this research is to determine how well this Kalman filtering and optimal smoothing algorithm can handle different parameters that affect the accuracy and precision of georeferencing simulated UAV imagery. Numerous trials were run altering the number of tie points and the initial GPS positional precision. The trials standardized the input standard deviation of the exterior orientation parameters to 1°. This is an overly pessimistic estimate as a MEMS-based IMU typically found on a UAV can achieve a RMS of half a degree or better. The following is a discussion of the initial findings from implementing this filtering and smoothing algorithm to sequential aero-triangulation.

In this set of simulations, the flying height was set to 200m and four conjugate tie point pairs were used to satisfy the co-planarity condition in each stereo-pair. The positional precision was evaluated with horizontal precision twice as good as vertical precision. For example, the trial with a horizontal precision of 10m was given an initial vertical precision of 20m. Five separate trials were investigated for the given input parameters. This was done to see the effect input GPS position precision had on the output orientation angle precision. The following graph shows the convergence of the forward filter to a steady state solution for the precision of the exterior orientation angles; similar results are obtained for the remaining angles.
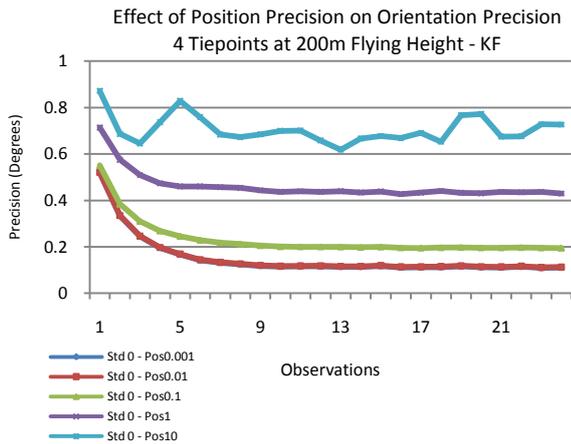
Fig. (1): Improved orientation precision due to improved position precision – Forward Kalman Filtering

Fig. (2): Improved orientation precision due to improved position precision – Backward Optimal Smoothing

It is apparent from Figure 1 that the filter is increasing the precision of the orientation angles as the sequential orientation proceeds down the strip of photos. It should be noted that although the GPS position precision is improving by a factor of ten for each of the five trials, the improvement in the orientation precision does not improve by a factor of ten. The initial findings suggest that improving GPS positional precision from consumer grade (±10m σ) to survey grade (±0.1m σ) result in substantial orientation precision gains. Return on investment for advanced geodetic grade receivers with positional accuracies of ±0.01m σ to ±0.001m σ do not offer results that substantially increase orientation precision. Therefore, investment in these expensive receivers may be unnecessary for the desired georeferencing precision.

Figure 2 is from the same set of data showing the results from the backward smoothing algorithm. The orientation precision results illustrate that the smoother is further refining the forward filter output. This marked improvement in orientation precision from ±1° σ to ±0.15° σ offers encouraging initial results for the implementation of this algorithm. Figure 2 also shows that there is a rebound effect at the end of the steady state; the reason is not yet known and believed to be programming error.
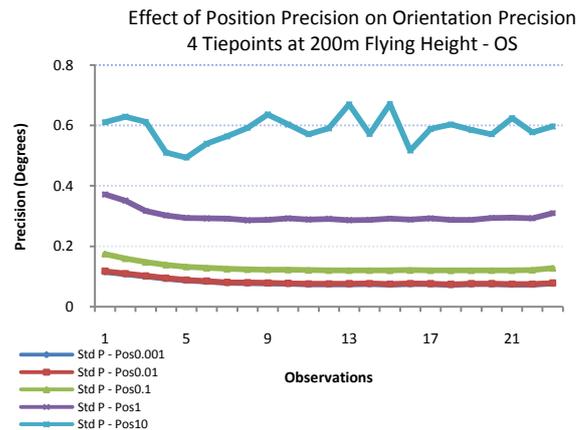
In the simulation model, the actual or 'true' values are known. This allowed us to analyze the algorithm accuracy by comparing the true values with the filtered/smoothed results. The term 'residual' will be used to define the difference between the true value and the estimated (filtered/smoothed) value, whilst the term 'error' will be used to define the injected simulation error. Using the data from the previous set of trials, Figure 3 shows the errors in the orientation angles.
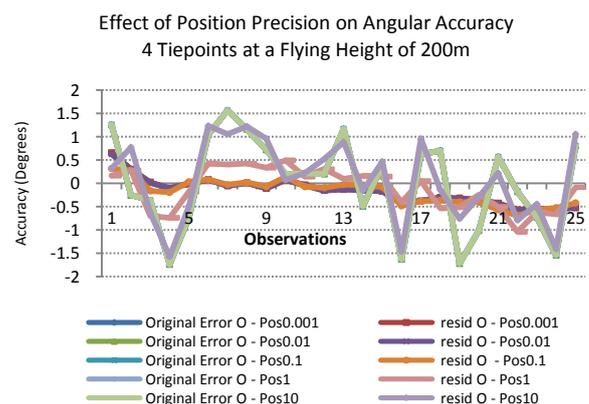

Fig. (3): Orientation accuracy at different position precision – Forward Kalman Filtering

Ideally, the residuals should be converging to 0° with oscillations of ±0.2° based on

algorithm design and the precision results. Noticeably the orientation angles, omega and phi, are not converging to the optimal residual. This bias is definitely of concern. Before reaching a definitive conclusion, an analysis of the results from the smoothing operation is shown in Figure 4.
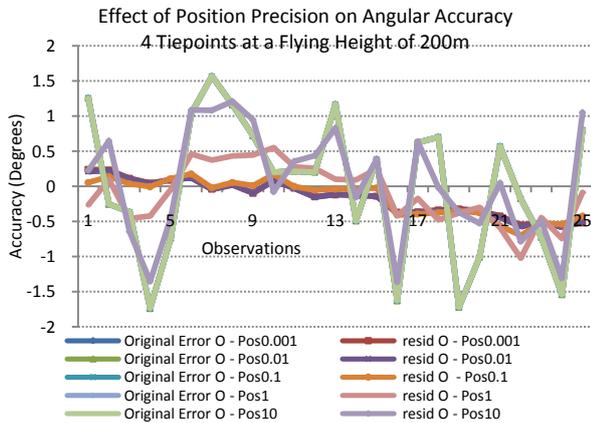

Fig. (4): Orientation accuracy at different position precision – Backward Optimal Smoothing

Figure 4 shows that the angular accuracy improves only for decent to high accuracy GPS positional precision; the optimal solution converged to the optimal solution of 0° towards the second half of the strip. The kappa angle residual (not shown), however, did not show the same behavior. More analysis of the algorithm is necessary to troubleshoot this problem. It should be noted that the difference in orientation angular accuracy between very fine GPS position precision (±0.001m σ) and decent GPS position precision (±0.1m σ) is negligible. This further reinforces the notion that a less expensive geodetic GPS unit may be just as suitable for georeferencing accuracy and precision as a top of the line geodetic GPS unit when this algorithm is implemented.

While some uncertainty exists as to the overall accuracy of this sequential orientation algorithm, we are quite confident in the improvement of the angular orientation precision as a result of algorithm implementation. Therefore, further simulation trials were run to analyze the impact of

increasing the number of tie points between the images. More tie points would lead to improved geometric stereo-model strength. This leads to the expectation that even greater gains in orientation angular precision can be expected with an increase in the number of tie points used. The results of increasing the number of tie points to nine supported the hypothesis that increasing the number of tie points leads to a subsequent increase in the angular precision. Relative to the four tie point session of simulation trials, the increase in orientation angle precision was significant. For higher orders of GPS positional precision (±10m σ), the relative increase in the smoothed angular precision between the two sessions was more than 100%. For smaller orders of GPS positional precision (±0.1m σ), the relative increase in the smoothed angular precision between the two sessions was still significant at approximately 30%. The orientation angle kappa had the most to gain from using more tie points. With four tie points, kappa had the weakest precision of the three angles. After using nine tie points, the three orientation angles had greater angular precision of equal magnitude.

A session of simulations was run using sixteen tie points to see if the trend continued. While improvements in angular precision were made, the difference in steady state precision was not significant. The significant change was a quicker convergence for 16 tie points relative to 9 tie points. Quicker convergence would seem to be the primary motive for utilizing this greater number of tie points. The improvement in angular precision is subtle and insignificant relative to using 9 tie points.

4. SUMARY AND CONCLUSION

We presented a method to process aerial imagery sequentially using an algorithm based on Forward Kalman Filtering and Backward RTS Optimal Smoothing. The approximately 90% increase in angular precision achieved with this algorithm relative to the original input data is a testament to the gains that can be made utilizing this algorithm for

extrapolating quality georeferencing parameters from sequential aero-triangulation. However, further investigation into the accuracy anomalies is necessary before an all inclusive endorsement is given to the full blown implementation of this algorithm.

When utilizing this algorithm, it is suggested that images be taken prior to the subject area to allow the system to initialize and reach steady state. Likewise, obtaining a few images beyond the subject area is advised to avoid end of the strip errors from affecting the beginning of the smoothing algorithm.

Further development of the simulation model to incorporate different flying heights is in progress. After the accuracy anomalies have been addressed, analysis of this algorithm with block sequential aero-triangulation will be the primary focus for further development. This is a natural progression since imagery acquisition in the real world is commonly collected in overlapping strips that form blocks of aerial photographs. Real data is available and will be utilized once the algorithm is tuned.

5. REFERENCES

Di, K. C., F. L. Xu, et al 2008. Photogrammetric processing of rover imagery of the 2003 Mars Exploration Rover mission. ISPRS Journal of Photogrammetry and Remote Sensing 63(2): 181-201.

Edmundson, K. and C. S. Fraser 1998. A practical evaluation of sequential estimation for vision metrology. ISPRS Journal of Photogrammetry and Remote Sensing 53(5): 272-285.

Gelb, A., 1974. Applied optimal estimation. Cambridge, Mass., M.I.T. Press.

Haala, N., D. Stallmann, et al 1998. Calibration of Directly Measured Position and Attitude by Aerotriangulation of Three-Line Airborne Imagery. ISPRS Journal of Photogrammetry and Remote Sensing 32(3): 8.

Kalman, R. E. 1960. A New Approach to Linear Filtering and Prediction Problems. Transactions of the ASME–Journal of Basic Engineering 82D: 35-45.

Kersten, T. P. and E. P. Baltsavias 1994. Sequential Estimation of Sensor Orientation for Stereo Images Sequences. International Archives of Photogrammetry and Remote Sensing 30(5): 8.

Mikhail, E. M., J. S. Bethel, C. McGlone 2001. Introduction to Modern Photogrammetry. New York, Wiley.

Nassar, S., X. J. Niu, et al 2007. Land-vehicle INS/GPS accurate positioning during GPS signal blockage periods. Journal of Surveying Engineering-ASCE 133(3): 134-143.

Rauch, H. E., F. Tung, et al 1965. Maximum Likelihood Estimates of Linear Dynamic Systems. AIAA Journal 3(8): 1445-&.

Tommaselli, A. M. G. and C. L. Tozzi 1996. A Recursive Approach to Space Resection using Straight Lines. Photogrammetric Engineering and Remote Sensing 62(1): 57-66.

Wang, X. and T. A. Clarke 2001. Separate Adjustment in Close-range Photogrammetry. ISPRS Journal of Photogrammetry and Remote Sensing 55(5-6): 289-298.

Webb, T. P. 2007. Vision-based State Estimation for Uninhabited Aerial Vehicles using the Co-planarity Constraint. PhD Dissertation - University of Florida, Gainesville, FL.

Welch, G. and G. Bishop 2001. An Introduction to the Kalman Filter. SIGGRAPH 2001. Chapel Hill, NC, University of North Carolina at Chapel Hill: 45.

Wolf, P. R. and B. A. Dewitt 2000. Elements of Photogrammetry : with Applications in GIS. Boston, McGraw-Hill.

Xiong, K., L. D. Liu, et al. 2009. Modified Unscented Kalman Filtering and its Application in Autonomous Satellite Navigation. Aerospace Science and

Technology 13(4-5): 238-246.