HIGH PRECISION LANDING SITE MAPPING AND ROVER LOCALIZATION BY INTEGRATED BUNDLE ADJUSTMENT OF MPF SURFACE IMAGES

Kaichang Di [a], Fengliang Xu [a], Ron Li [a], Larry H. Matthies [b], Clark F. Olson [c]

[a] Department of Civil and Environmental Engineering and Geodetic Science, The Ohio State University, 470 Hitchcock Hall, 2070 Neil Avenue, Columbus, OH 43210-1275 - (li.282, di.2, xu.101)@osu.edu
[b] Jet Propulsion Laboratory, California Institute of Technology, Mail Stop 125-209, 4800 Oak Grove Drive, Pasadena, CA 9110 - lhm@telerobotics.jpl.nasa.gov
[c] University of Washington, Bothell, Computing and Software Systems
18115 Campus Way NE, Box 358534, Bothell, WA 98011 - cfolson@u.washington.edu

Commission IV, WG IV/9

KEY WORDS: Mars Pathfinder Mission, Landing Site Mapping, Rover Localization, Bundle Adjustment

ABSTRACT:

High-precision topographic information from all available data is crucial to many landing site geological and engineering applications. At the same time, precise navigation and localization of the rover as it traverses the Martian surface is important both for its safety and for the achievement of its engineering and scientific objectives. In this paper, we investigate a new approach to high-precision Mars landing site mapping and rover localization based on integrated bundle adjustment of an image network built by linking ground-based images with automatically or manually selected tie points. The method and software are tested using lander and rover image data obtained from the Mars Pathfinder mission. An innovative method for automatic tie point selection is also presented.

1. INTRODUCTION

High-precision global and local topographic information is crucial for support of engineering operations and achievement of scientific goals in Mars exploration missions. In particular, landing site mapping will be extremely important for landing and rover navigation in future lander and rover missions such as the 2003 Mars Exploration Rovers (MER), the European Beagle 2 lander of the 2003 Mars Express, the 2007 Smart Lander and Long-range Rover, and sample return lander missions beyond 2010.

In Mars rover missions, accurate navigation and localization of the rover relative to the landing center are needed so that the rover can traverse the Martian surface safely, return repeatedly to the same location to perform operations, and support coordinated multidisciplinary high-precision scientific experiments. In the Mars Pathfinder (MPF) mission, the rover Sojourner principally used a heading sensor and wheel counters for localization. Sojourner accumulated 1m location errors in an area of about 10m x 10m during this mission. In the MER 2003 mission, rovers will traverse an extended distance of 600m up to 1,000m where the terrain may be more challenging. In the 2007 Smart Lander and Long-range Rover mission, the rover traverse will be even longer. More accurate rover localization will be highly desirable for achievement of the scientific objectives of future missions. In order to prevent the accumulation of localization errors, descent and orbital images should be used as well to provide global constraints.

Since 1998, the Mapping and GIS Laboratory at The Ohio State University and the NASA JPL Machine Vision Group have been jointly developing a bundle adjustment method with relevant techniques for the processing of Mars descent and rover images for rover localization and landing site mapping.

We missed an opportunity to test the concept using actual MARDI (Mars Descent Imager) data because of the failure of the Mars Polar Lander Mission. In order to verify our algorithm and software, field tests were conducted at Silver Lake, CA in April 1999 and May 2000. Based on the data obtained, various rover localization experiments were carried out. Using descent and rover images along with either an integrated or an incremental adjustment, rover localization accuracy was reached of approximately 1m for a traverse length of 1km from the landing center (Li et al., 2000, 2001; Ma et al., 2001). Experiment results also showed that it is feasible to localize the rover by using rover images only if no descent images are available (as with the MER mission), and yet still achieve a similar accuracy when issues such as optimal traverse design and image network generation are considered (Di et al., 2002).

In addition to using simulated descent and rover images, we tested our methods and software with actual Mars data. We downloaded rover and lander IMP (Imager for Mars Pathfinder) images from the PDS (Planetary Data System) web site. The German Aerospace Center (DLR) also provided us with a complete panorama chosen from a vast mount of IMP images. In the following section we briefly describe the bundle adjustment models used in this investigation and our other relevant studies. In subsequent sections, the processing of IMP panoramic images and rover images are presented.

2. BUNDLE ADJUSTMENT MODELS

The basic model for the bundle adjustment is based on the well known collinearity equations (Wolf and Dewitt, 2000). The linearized observation equation is expressed in matrix form as:

$$V = AX - L \qquad (1)$$

In this bundle adjustment model, all of the unknowns (camera position, orientation of all the images, and 3D ground position of the tie points) are adjusted together after all the images are acquired. Therefore we call it the integrated bundle adjustment model.

Because there is no absolute ground control available on the Martian surface, the adjustment is a free net adjustment where the normal matrix is rank deficient. We used the Singular Value Decomposition technique to solve the normal equation in which the Minimum Norm is applied using the Least Squares principle. If certain distinctive features (such as mountain peaks or craters) can be observed from the orbital images, they are used as relative controls for the adjustment in the following form:

$$HX = W \qquad (2)$$

In order to improve computational efficiency and make the bundle adjustment applicable to a real-time operation, we developed two incremental bundle adjustment models (Li et al., 2000, 2001; Ma et al., 2001). Model I is represented as:

$$
\begin{aligned}
V_{m-1} &= A_{m-1}X_{m-1} - L_{m-1}, \\
V_m &= A_m X_m + B_m Y_m - L_m
\end{aligned}
\qquad (3)
$$

Model II is represented as:

$$
\begin{aligned}
V_{m-1} &= A_{m-1}X_{m-1} + B_{m-1}Y_{m-1} - L_{m-1}, \\
V_m &= B_m Y_m + C_m Z_m - L_m
\end{aligned}
\qquad (4)
$$

In Model I, the unknowns are added gradually into the adjustment system when the new observations are available. After the step *m* adjustment, unknowns solved by step *m-1* are also updated so that the final results are ultimately the same as those from an integrated bundle adjustment model. In Model II, with the gradual increase of new unknowns some older unknowns are no longer taken into account, thus offering more flexibility and efficiency to the process.

There are several frames of reference that are used in the MPF image pointing data, including the camera head coordinate system, lander (L) frame, Martian Local Level (M) coordinate frame, Mars Surface Fixed (MFX) coordinate frame, and Landing Site Cartographic (LSC) coordinate system. Bundle adjustment and topographic products of our models are based on the LSC system as defined by the U.S. Geological Survey.

We developed a program to convert the pointing data of the PDS images to the exterior orientation data. This is accomplished by a chain of translations and rotations through the above frames of reference. Converted exterior orientation values were then used as initial values in the bundle adjustment.

3. PROCESSING OF IMP IMAGES

We used 60 pairs of IMP stereo images provided by DLR to explore automatic tie point selection, bundle adjustment, DEM (Digital Elevation Model) and orthoimage generation. These IMP stereo images form a complete panorama (as shown in Figure 1), in which the azimuth and the tilt angles of the images are represented in the x and y axes, respectively.
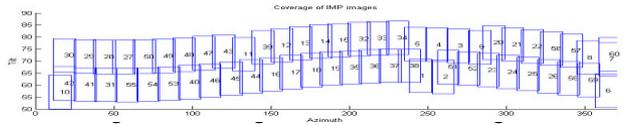
Figure 1. IMP image coverage

We developed a series of techniques to automatically select tie points within one stereo pair (intra-stereo) and between adjacent stereo pairs (inter-stereo).

For intra-stereo tie point selection, the Förstner operator is applied to extract interesting points. Then, an area-based matching is applied to match these interesting points using normalized correlation coefficients. Experiments with this data set show that in most cases over 90% of the matches can be found. Figure 2 shows the matched (white cross) and unmatched (black cross) interesting points in an intra-stereo pair. In this stereo pair, 319 and 312 interesting points were extracted from the left and right images, respectively, and 216 points were matched.
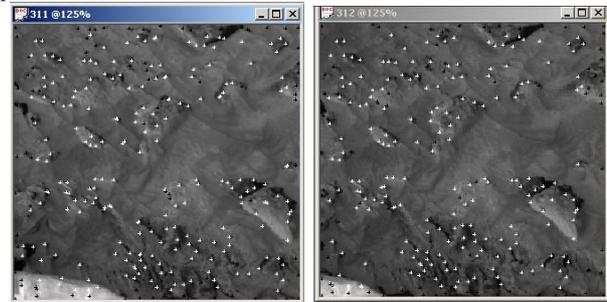
Figure 2. Interesting point matching

Within a set of matched points, some mismatched points may still exist that need further verification. Automatic verification of the matched tie points is based on the consistency of parallax. Assuming that the terrain does not vary significantly, parallaxes of the points should maintain an order throughout the image. In other works, if two points in the left image are located at the top left and bottom right corners, respectively, then their matches should preserve the same order (e.g. measured by a distance). Disordered points are considered as mismatches.

The order can be expressed simply as a function of parallaxes along the row direction, since in the column direction the parallaxes are almost a constant (close to zero). In Figure 3, we order the interesting points in the row direction and then draw their x (left) and y (right) parallax curves in the same order. Any abrupt changes of the parallaxes are assumed to represent mismatches. Mismatches are eliminated using a median filter to produce the smoother parallax curves found in Figure 4. These smooth parallax curves represent the correct order, and are used as reference parallaxes.
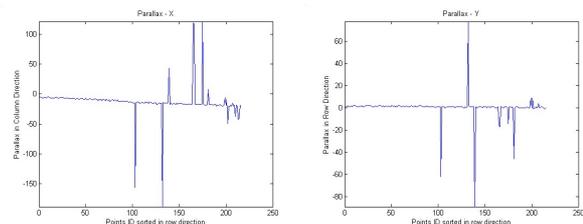
Figure 3. The x and y parallax curves
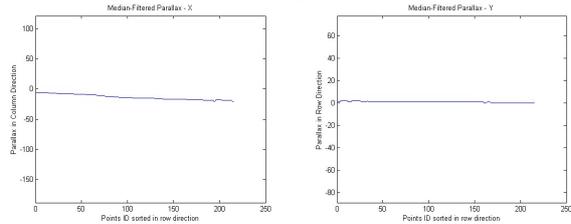
Figure 4. Smoothed x and y parallax curves

Figure 6. Distribution of all intra-stereo interesting points

Figure 7. The Coarse DEM

For each matched interesting point pair, the actual parallax is compared to the value of the reference parallax of the same row. If it is within a specified range (empirical values used are 10 pixels in x and 5 pixels in y), the pair is considered a correct match. 201 matches (93%) passed this parallax verification. A visual check was conducted to confirm the result.

To prepare an even distribution of the tie points for use in the later bundle adjustment and for improvement of computational efficiency, the final tie points are selected from the set of correct matches using a gridding scheme. Each image is evenly segmented into several patches, e.g., a 3 x 3 grid. Within each patch, the pair of correctly matched interesting points with the highest value of grayscale variance is chosen. Since this is the most distinguishable matched point within the patch, it can serve as a reliable tie point. In Figure 5, 9 pairs of final tie points have been selected for each stereo image pair.
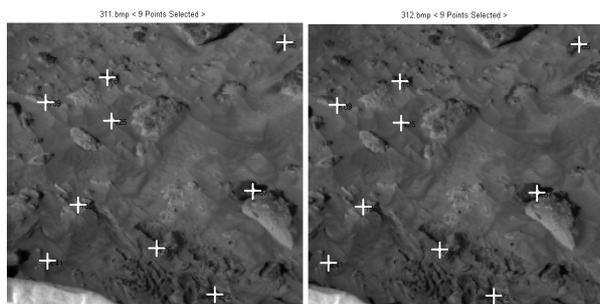
Figure 5. Nine evenly distributed final tie points are selected for each intra-stereo pair

Application of this strategy is much more difficult when selecting tie points between cross-stereo pairs or inter-stereo pairs, because disparities can often be very large. Or there may even be no overlapping across the stereo image pairs. In general, overlapping between the image margins should be very small (around 10%).

To get around this problem for inter-stereo pair images in the panorama, their relative orientations can be estimated to an acceptable level of precision by using information about the rotation of the cameras relative to the landing base. We can calculate 3D coordinates for matched tie points from the previous step and expect no large inconsistency. Figure 6 shows the distribution of 4,709 intra-stereo interesting points (matched and verified).

The 3D ground points are then used to interpolate a coarse DEM that, hopefully, will not differ very much from the final DEM. Figure 7 shows the generated coarse DEM. This is a global DEM for the entire panorama. More importantly, it can be used to bridge the stereo image pairs.
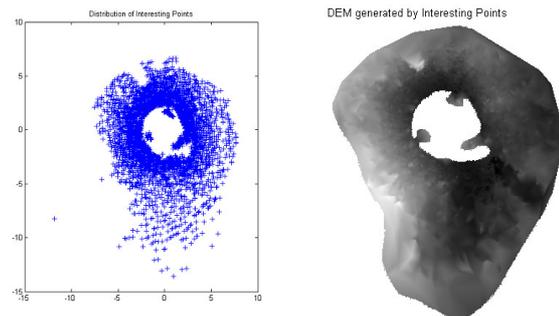
With the coarse DEM, we can predict the overlapping of inter-stereo images. In addition, we can predict for an interesting point the approximate location of its conjugate point in its corresponding inter-stereo pair image. For each interesting point, we can find its 3D coordinates in the DEM and back project it onto its inter-stereo image. The back projected location should not be too far from the actual position. A small search range can thus be used around the back projected location. In this way, most interesting points in the inter-stereo pair margins can be matched. The success rate is 90% for most cases. Verification is accomplished by again using the row-based reference parallax method. Figure 8 illustrates 4 final tie points that were selected in an inter-stereo image pair after matching, verification and gridding selection.
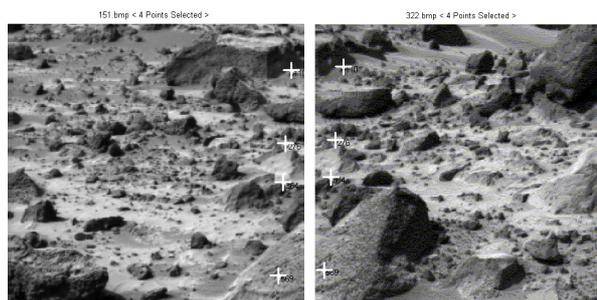

Figure 8. Selected final tie points in an inter-stereo pair

Difficulties can occur when the overlapping margins are too small and/or radiometric characteristics of the images differ greatly (e.g., when one image was taken under dim conditions while the other was taken in bright sunlight), or shadows have moved significantly from changing sun angles. However, most inter-stereo images are taken within a short time interval and the above challenge can be avoided.

Inter-stereo image pairs can be linked by three types of connections: in the horizontal, vertical, and/or diagonal direction. Because of the difficulties mentioned above, it cannot be assumed that every inter-stereo image pair will find all 3 links. If all the image pairs within a panorama can be linked by at least 1 of the 3 types, they can form a network and be bundle adjusted to estimate improved precision exterior orientation parameters. The refined parameters are then used to generate a high precision DEM.

The final links between the stereo pairs that were found through the automatically selected tie points are shown in Figure 9. Just as Figure 1, the x and y axes in Figure 9 stand for the azimuth and tilt angles of the IMP images. In total 84 pairs of inter-

stereo images were linked by the tie points. Each inter-stereo pair has 6 selected tie points. For the entire panorama, there are 5,859 matched interesting points in the overlapping margins, among which 589 have been selected as final inter-stereo pair tie points. Overall, there are 24,903 matched interesting points in both intra- and inter-stereo pair images, among which 1,128 have been selected as final tie points.
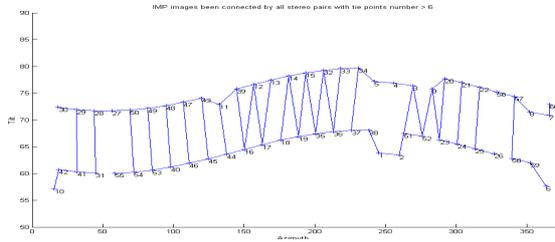

Figure 9. Connection map of the panoramic images

We are currently checking the quality of these tie points. We will then perform the bundle adjustment of the complete panorama with the automatically selected tie points and generate a precise DEM and orthoimage. Final mapping results will be presented at the symposium.

4. BUNDLE ADJUSTMENT OF IMP LANDER AND ROVER IMAGES

In order to form a rover network, the rover should "continuously" take images at designed adjacent stations so that the images can be linked together. Also the network configuration should be designed according to some network optimization principles to achieve a high accuracy in mapping and rover localization (Di et al., 2002). However, in the MPF mission, the rover did not take images continuously, so we could not test rover localization in a rover image traverse. Even for separate rover images, it is difficult to find sufficient tie points to link them to IMP images. We selected 10 stereo pairs of IMP images and 2 stereo pairs of rover images to test rover localization accuracy. Figure 10 shows the left images of 2 (out of the 10) IMP stereo images that have overlay with the rover images. Figure 11 shows the 2 rover stereo pairs used in the experiment. The rover was about 5 meters from the lander when these images were taken.
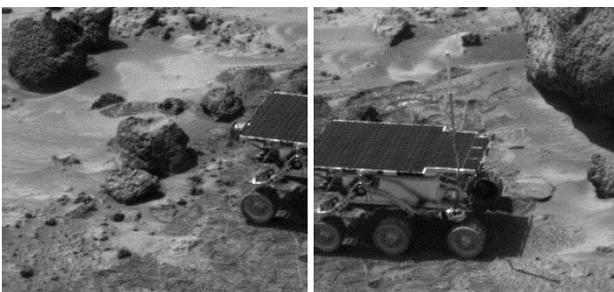

Figure 10. IMP images that have overlay with rover images

At first, we conducted a bundle adjustment by using the 10 IMP stereo pairs and no rover images. 155 tie points are selected to link the 20 images into an image network. The maximum distance from the tie point to the lander was about 25 meters, and the average distance about 6 meters. As a result of the bundle adjustment, the average of the estimated deviation of the lander camera positions were 0.013, 0.018 and 0.013 meters in the X (east), Y (north) and Z (elevation) directions respectively. The average of the estimated deviations of the tie point ground positions are 0.065m, 0.155m, and 0.028m in X, Y, and Z directions.

We manually selected 28 inter-stereo tie points that were different from those used in the bundle adjustment and used them to check the bundle adjustment accuracy. The 3D ground positions of these tie points were calculated by space intersection from one set of stereo pairs. Then we projected the computed ground points back onto the other stereo pairs. The differences between the back projected image points and those identified manually reflect the quality of the exterior orientation parameters. The space intersection and back projection were conducted using the exterior orientation parameters before and after the bundle adjustment. The average differences in x and y were 2.9 and 6.6 pixels before adjustment, and 0.4 and 0.3 pixel after adjustment. Maximum differences were 5.3 and 14.6 pixels in x and y before adjustment, and 1.0 and 0.8 pixel after adjustment. As the image coordinate differences were apparently reduced, this indicates that the accuracies of exterior orientation parameters were improved.
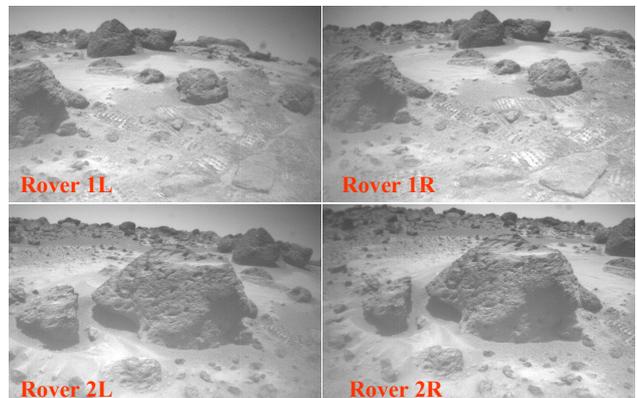

Figure 11. Rover stereo pairs that have overlay with IMP images

Next we conducted an integrated bundle adjustment by using the rover and IMP images together. An image network was built by linking the IMP and rover images: in addition to 155 tie points linked the IMP images, another 15 tie points were manually selected to link IMP and rover images and 20 tie points were selected to link the rover images.

In the image network, the distribution of the tie points between rover Stereo Pair 1 (top row in Figure 11) and IMP images is better than that of Stereo Pair 2 (lower row in Figure 11). For comparison purposes, we carried out 2 bundle adjustment experiments: one using just rover Stereo Pair 1 with all the IMP images and the other using both stereo pairs with all the IMP images.

From the results of the integrated bundle adjustment, we observed that the accuracy of IMP lander image orientation parameters is not significantly changed by adding the 4 rover images. From Experiment I, the average of the estimated deviations of the rover camera positions were 0.074, 0.085 and 0.031meters in the X, Y and Z (height) directions respectively. In Experiment II, the estimated deviations of the rover camera positions were 0.083, 0.081 and 0.035 meters in the X, Y and Z directions. We can see that the rover localization accuracy estimated from the two experiments is very similar (about 2% of the distance from the lander).

We also assessed the rover localization accuracy by comparing the ground point positions of tie points calculated from rover stereo pairs with those from IMP stereo pairs. The average difference for 10 tie points from Experiment I was 0.118, 0.051, and 0.017 meters in the X, Y and Z directions; the average difference of 15 tie points from Experiment II was 0.426, 0.178, and 0.134 meters in the X, Y and Z directions. As is apparent from above, results from the second experiment were not as good as those from the first experiment. This is due to the fact that the geometry between the second rover stereo pair and the lander images is relatively weak.

5. SUMMARY

We propose a landing site mapping and rover localization method based on bundle adjustment of an image network built by linking ground-based images. Integrated and incremental bundle adjustment models and automatic tie point selection methods have been developed. Sixty MPF IMP images of a complete panorama were used to test the effectiveness of automatic intra- and inter-stereo tie point selection, which turned out to be very promising. Ten IMP stereo pairs and two rover pairs were used to test the capabilities of the bundle adjustment. The results demonstrated that the accuracy of exterior orientation parameters of the IMP images was improved, consequently the differences between measured image points and those back-projected from adjacent stereo images were reduced from several to more than ten pixels to a sub pixel level. With this method, a rover can be localized to an accuracy of about 2% of the distance from the lander (depending on the strength of the link between rover and lander images). Mapping products from the complete panorama of IMP images, in particular precision DEM and orthoimage, are currently being processed and will be presented at the symposium.

REFERENCES

Di, K., R. Li, L.H. Matthies, C.F. Olson, 2002. A Study on Optimal Design of Image Traverse Networks for Mars Rover Localization. *2002 ASPRS-ACSM Annual Conference and FIG XXII Congress*, April 22-26, Washington DC. (CD-ROM).

Li, R., F. Ma, F. Xu, L. Matthies, C. Olson and Y. Xiong, 2000. Large Scale Mars Mapping and Rover Localization using Descent and Rover Imagery. *Proceedings of ISPRS 19th Congress,* Amsterdam, July 16-23, 2000 (CD-ROM).

Li, R., F. Ma, F. Xu, L.H. Matthies, C.F. Olson, and R.E. Arvidson, 2001. Localization of Mars Rovers using Descent and Surface-based Image Data. *Journal of Geophysical Research-Planet*, Accepted.

Ma, F., K. Di, R. Li, L. Matthies, and C. Olson, 2001. Incremental Mars Rover Localization using Decent and Rover Imagery. *ASPRS Annual Conference 2001*, April 25-27, St. Louis, MO. (CD-ROM).

Olson, C.F., L.H. Matthies, Y. Yiong, R. Li, F. Ma and F. Xu, 2001. Multi-resolution Mapping Using Surface, Descent and Orbital Imagery. *The 2001 International Symposium on Artificial Intelligence, Robotics and Automation in Space*, Montreal, June 2001.

Xu, F., F. Ma, R. Li, L. Matthies and C. Olson, 2001. Automatic Generation of a Hierarchical DEM for Mars Rover Navigation. *Proceedings of the 3rd Mobile Mapping Conference*, Cairo, Egypt, January 3-5, 2001.

Wolf, P.R., and B.A. Dewitt, 2000. *Elements of Photogrammetry with Applications in GIS*, Third Edition, McGraw-Hill, 608p.