

A stereo vision system for support of planetary surface exploration

Maarten Vergauwen, Marc Pollefeys* and Luc Van Gool

K.U.Leuven ESAT-PSI, Kasteelpark Arenberg 10
B-3001 Leuven, Belgium
firstname.lastname@esat.kuleuven.ac.be

Abstract. In this paper a system will be presented that was developed for ESA for the support of planetary exploration. The system that is sent to the planetary surface consists of a rover and a lander. The lander contains a stereo head equipped with a pan-tilt mechanism. This vision system is used both for modeling of the terrain and for localization of the rover. Both tasks are necessary for the navigation of the rover. Due to the stress that occurs during the flight a recalibration of the stereo vision system is required once it is deployed on the planet. Due to practical limitations it is infeasible to use a known calibration pattern for this purpose and therefore a new calibration procedure had to be developed that can work on images of the planetary environment. This automatic procedure recovers the relative orientation of the cameras and the pan- and tilt-axis, besides the exterior orientation for all the images. The same images are subsequently used to recover the 3D structure of the terrain. For this purpose a dense stereo matching algorithm is used that -after rectification- computes a disparity map. Finally, all the disparity maps are merged into a single digital terrain model. In this paper a simple and elegant procedure is proposed that achieves that goal. The fact that the same images can be used for both calibration and 3D reconstruction is important since in general the communication bandwidth is very limited. In addition to the use for navigation and path planning, the 3D model of the terrain is also used for Virtual Reality simulation of the mission, in which case the model is texture mapped with the original images. The system has been implemented and the first tests on the ESA planetary terrain testbed were successful.

Keywords: Active Vision systems, Prototype systems, 3D reconstruction, Stereo calibration.

1 Introduction

The work described in this paper was performed in the scope of the ROBUST¹ project of the European Space Agency (ESA). In this project an end-to-end system is developed for a planetary exploration mission.

The ROBUST system consists of three important parts.

* Postdoctoral Fellow of the Fund for Scientific Research - Flanders (Belgium) (F.W.O. - Vlaanderen).

¹ The ROBUST consortium consists of the Belgian companies SAS and OptiDrive, the K.U.Leuven departments PMA and ESAT-PSI and the German companies DLR and vH&S.

- the planetary rover: the Nanokhod, a small and simple rover, designed to carry instruments in the immediate surroundings of a lander. It is equipped with a tether cable, providing the rover with power and data connection to the lander which allows a very high ratio instrument-mass/rover-mass [10]. Figure 1 on the left shows an image of the Nanokhod.
- The Planetary Lander which contains the Imaging Head, an On-Board computer and the Control System for both Nanokhod and Imaging Head. The right image of figure 1 shows one of the cameras.
- The On Ground Control System

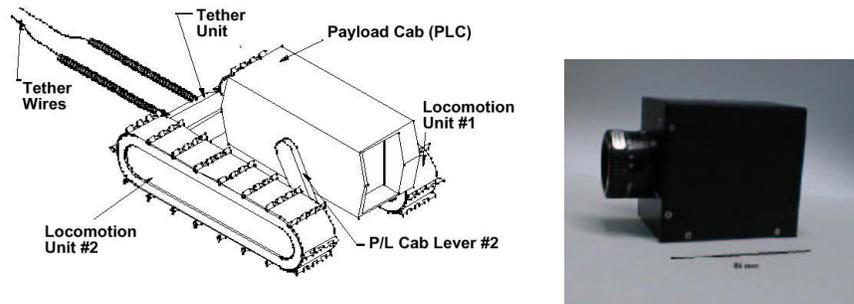


Fig. 1. left: the Nanokhod; right: one of the cameras of the stereo rig

The Imaging Head is both used for recording images from which a reconstruction of the planetary terrain is computed and for controlling the motion of the rover, using Light Emitting Diodes on the payload cab of the rover for the latter. It consists of a stereo head, mounted on a unit which allows for pan and tilt motions and which is approximately 1.5 meter high. The two cameras of the stereo head are space approved 1024x1024 CCD cameras. The stereo head has a baseline of 0.5 meter.

A typical utilization scenario will deploy the Imaging Head as soon as possible after the landing of the planetary system. Because of the strain on the parts during launch and landing, the Imaging Head needs to be recalibrated. To accomplish this, it takes images of the terrain which are sent to earth where the calibration is performed using these images. From the same images a 3D reconstruction of the terrain is then computed. Since the cameras have a limited field of view (23x23 degrees) the entire environment is not recorded at once but it is segmented into rings according to the tilt angle and each ring is divided into segments according to the pan angle of the Imaging Head (see figure 2). The outermost boundary of the recorded terrain lies at twenty meters from the camera. For each of the segments a stereo image pair is recorded and sent down. The values of the actual pan and tilt angles can be read out from the encoders of the pan-tilt motors and are sent down together with the corresponding images.

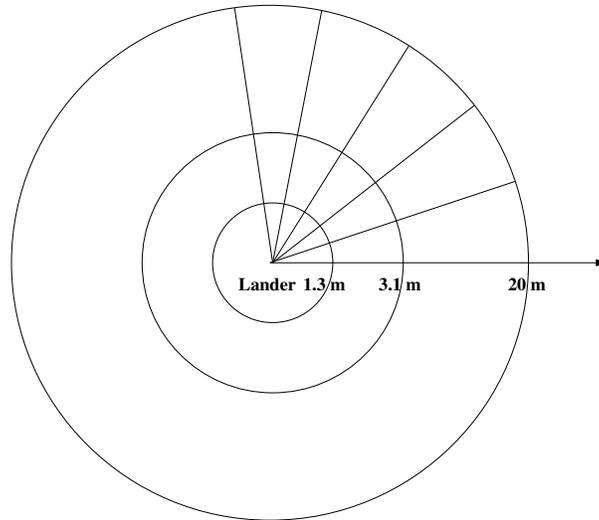


Fig. 2. Segmentation of the terrain into segments

2 Calibration

Every planetary mission is a high-risk operation. During launch and landing, the lander and its contents are subject to extreme forces. The mechanical properties of the Imaging Head are likely to have been affected by mechanical and thermal effects. For high accuracy equipment, such as the Imaging Head, a small change in these mechanical properties results in large degradation of the results, unless the new properties can be estimated. The cameras themselves are built so that the intrinsic parameters during the mission can be assumed identical to the parameters obtained through calibration on ground. If the camera housing were not so rigidly built and the camera intrinsics were likely to change during launch or landing. Algorithms exist that can retrieve these intrinsic parameters from images too [12, 7].

Traditional calibration algorithms rely on known calibration objects with well-defined optical characteristics in the scene. If cameras take images of these artificial objects, the pose of the cameras can be computed, yielding the extrinsic (mechanical) calibration of the stereo rig. There are two reasons why this scheme is not suited in our case where the Imaging Head is deployed on a distant planet. First there is the problem of where to place the calibration objects. One needs to be absolutely sure of the pose of these objects for the calibration to have any meaningful result. It is of course impossible to add objects to the terrain, so one has to think of placing calibration "markers" on the lander itself. A typical lander consist of a "cocoon" which opens after landing, comparable to an opening flower. The markers could be applied to the opening "petals". However, one is never sure of the exact position of these petals which makes the markers much harder to use. Even if one did dispose of accurate markers on the lander, a second problem arises. During the design of the Imaging Head, robustness was a very important issue and therefore the number of moving items was minimized. Therefore, a fixed focus lens

was chosen. Since the accuracy of the stereo matching decreases with the square of the distance, the cameras are focussed on the far range to gain as much accuracy in the far regions as possible. As a consequence, the images of near regions are blurred. Since the markers would be on the lander, images of the markers would always be blurred, reducing the accuracy of the calibration up to the point where the markers are useless. It is clear that standard calibration algorithms can not be used in our system. A new strategy had to be developed that only uses images of the terrain to calibrate the Imaging Head.

The calibration procedure that was implemented for the ROBUST project is able to calibrate the Imaging Head using images of the terrain only. This means that the images which are sent down from the planet to earth to reconstruct the terrain, can also be used for calibrating the Imaging Head. Therefore, the terrain based calibration causes no important overhead on transmission. The calibration of the extrinsic (mechanical) properties of the Imaging Head is split into two parts which are executed consecutively. First the relative transformation between the two cameras is computed. Once this relative calibration is performed, a procedure can be performed which computes the relative transformations between the cameras and the lander. This boils down to computing the pan and tilt axes of the pan-tilt unit.

2.1 Relative calibration

The relative transformation between the two cameras of the Imaging Head can be computed from images of the terrain only. The algorithm to do this uses the concept of the essential matrix. This matrix represents the epipolar geometry between two views, including the internal parameters of the cameras as extra information. We make use of the fact that the relative transformation between the cameras does not change when the different segments of the terrain are recorded, which allows for different measurements of the epipolar geometry to be combined to yield one accurate solution.

If the essential matrix between the two views is computed, the relative transformation (position and orientation) between the two cameras can be calculated up to the baseline (i.e. the distance between the two cameras).

Computing epipolar geometry The first step in obtaining the relative calibration is the computation of the epipolar geometry of the stereo head. The epipolar geometry constraint limits the search for the correspondence of a point in one image to points on a line in the second image. Figure 3 illustrates this.

To find back the epipolar geometry between two images automatically, a feature detector, called the Harris Corner Detector [3] is applied to the images. Next, the corners are matched automatically between pairs of images using cross correlation. This process yields a set of possible matches which is typically contaminated with an important number of wrong matches or outliers. Therefore a robust matching scheme, called RANSAC[2], is used to compute and update epipolar geometry and matches iteratively.

In the case of the ROBUST Imaging Head the data of the different segments of the terrain can be combined to compute the epipolar geometry much more robustly because the relative transformation between the cameras does not change. Figure 4 illustrates this. Stereo images of different rings are obtained by tilting the Imaging Head. How-

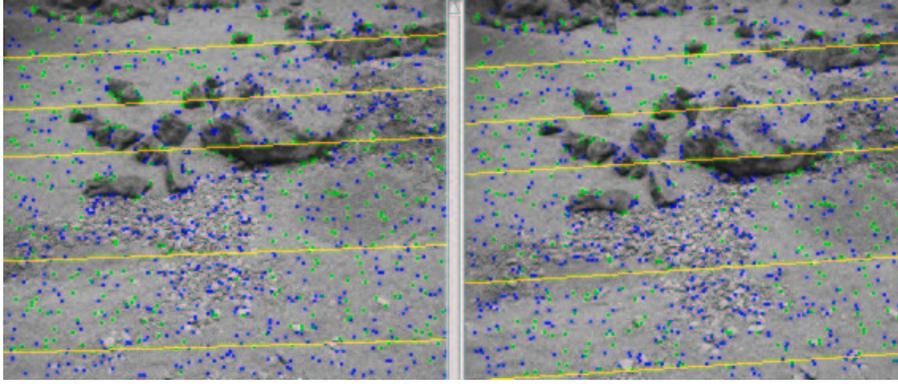


Fig. 3. Epipolar geometry of an image pair

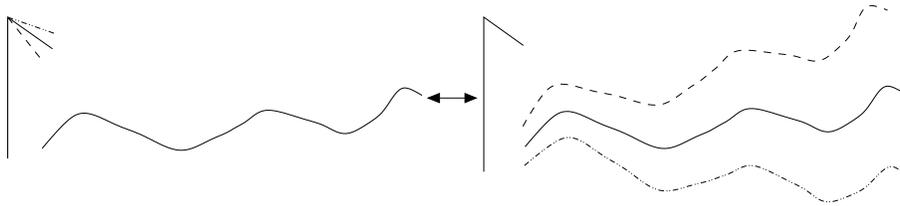


Fig. 4. Combining different segments

ever, one could imagine the camera to be kept steady and the terrain to be tilted. This would result in the same stereo images. That is why the possible correspondences of the different rings and segments can be combined to compute the epipolar geometry more accurately.

It is even the case that a specific degenerate case for the computation of the epipolar geometry is solved by the combination scheme we described. Computing the epipolar geometry of a pair of images of a planar scene is impossible from correspondences only. If the planetary terrain is planar or close to it, computing the epipolar geometry for one pair of images becomes an ill-posed problem. By combining correspondences from different segments, this problem is solved.

Computing relative transformation Once the epipolar geometry is computed in the form of the fundamental matrix F , the relative transformation between the two cameras of the Imaging Head can be calculated. First the essential matrix is obtained as $\mathbf{E} = \mathbf{K}^T \mathbf{F} \mathbf{K}$ with \mathbf{K} the 3x3 matrix with the intrinsic calibration of the cameras. The relative transformation (\mathbf{R}, \mathbf{t}) (up to scale) can easily be obtained from it [6] since $\mathbf{E} = [\mathbf{t}]_{\times} \mathbf{R}$.

As mentioned, the absolute value of the distance between the two cameras, i.e. the baseline, can not be computed from the images. It is unlikely, however, that this value will deviate a lot from the original value. Since all measurements for 3D reconstructions and localization will consistently be carried out using the same stereo head, this has only

very little influence. The only case where absolute measurements are really needed is for estimating tether consumption and the absolute size of obstacles to overcome. In these cases even a deviation of a few percent would be smaller than the uncertainty induced by other factors.

The computed values for \mathbf{R} and \mathbf{t} are used as an initialization for a non-linear Levenberg-Marquardt minimization which finds back the values of \mathbf{R} and \mathbf{t} that minimize sum of all distances between points and their corresponding epipolar lines. The result is a very accurate estimation of the relative transformation between the two images.

2.2 Pan-tilt calibration

To be able to bring all the measurements in a single frame, it is not sufficient to have the relative calibration between the cameras, but also the pan- and the tilt-axis are needed. Since for the same reasons as for the relative calibration these values can change due to the stress that occurs during launch and landing.

The evaluation of the pan- and the tilt-axis is more complicated, but can also be achieved from the image data, at least if some overlap exists between neighboring image pairs. This is guaranteed in the pan direction due to the fixed vergence set-up that does not allow 100% overlap between two views of the same stereo pair (and thus yields overlap with neighboring pairs). For the tilt axis care should be taken to foresee a sufficient overlap between the rings of Figure 2, at least for some image pairs.

To compute the different axes, the relative motion between stereo pairs is required. Matches can be obtained in a way similar to the one described in Section 2.1, however, in this case most features have already been reconstructed in 3D for one of the stereo pairs and a robust pose estimation algorithm (i.e. also based on RANSAC) can be used to determine the relative transformation between the stereo pairs. Once this has been done for all the overlapping pairs, the results can be combined to yield one consistent estimate of both pan and tilt axes. In this case available constraints such as the rotation angles (read from the encoders) are also enforced.

2.3 Synthetic Calibration experiment

The calibration algorithm has been tested on artificial data. A planar scene with texture from a real image from Mars was constructed and pairs of images were generated with a visualization toolkit. First the relative calibration between the two cameras was computed. During calibration, data of 9 image pairs was combined and 2591 corners were matched to calculate the relative transformation. We could compare the result with the ground-truth value of the relative transformation. For comparison, the rotational part of the relative transformation is represented as a rotation around an axis (l, m, n) with a certain angle θ . The ground truth was

$$(l_0, m_0, n_0) = (0, 1, 0) \text{ and } \theta_0 = 10^\circ$$

The computed values were

$$(l, m, n) = (0.000397, 0.99951, 0.00018) \text{ and } \theta = 10.093^\circ$$

The angle between (l, m, n) and (l_0, m_0, n_0) was 0.0252 degrees. The difference between θ and θ_0 was 0.093 degrees. Both values are small, meaning that the rotation was estimated accurately because of the combination of data.

The pan and tilt axes were calibrated from the same data. The angle between the computed tilt axis and the ground truth was 0.138 degrees. The angle between the computed pan axis and the ground truth was 0.618 degrees. The larger error for the pan axis can be explained from the fact that only correspondences between three images are used to estimate it while correspondences from four images can be exploited to compute the tilt axis. During calibration of the real system, better results can be expected because much more image pairs of a non planar scene will be used.

3 3D Terrain modeling

After the calibration of the IH is performed, the process of generating a 3D model or models of the planetary terrain can start. This modeling is vital to accomplish the goal of planetary exploration. Its input are all images of the terrain and the calibration of the Imaging Head. The output of the terrain modeling can have different forms but the most important is the Digital Elevation Map (DEM). In the following sections we will describe the different steps that are performed to obtain such a DEM.

3.1 Generation of disparity maps

On each pair of images recorded by the Imaging Head, a stereo algorithm is applied to compute the disparity maps from the left image to the right and from the right image to the left. Disparity maps are an elegant way to describe correspondences between two images if the images are **rectified** first. The process of rectification re-maps the image pair to standard geometry with the epipolar lines coinciding with the image scan lines [11, 8]. The correspondence search is then reduced to a matching of the image points along each image scan-line. The result (the disparity maps) is an image where the value of each pixel corresponds with the number of pixels one has to move to left or right to find the corresponding pixel in the other image.

In addition to the epipolar geometry other constraints like preserving the order of neighboring pixels, bidirectional uniqueness of the match and detection of occlusions can be exploited. The dense correspondence scheme we employ to construct the disparity maps is the one described in [5]. This scheme is based on the dynamic programming scheme of Cox [1]. It operates on rectified image pairs and incorporates the above mentioned constraints. The matcher searches at each pixel in the left image for the maximum normalized cross correlation in the right image by shifting a small measurement window along the corresponding scan line. Matching ambiguities are resolved by exploiting the ordering constrain in the dynamic programming approach.

3.2 Digital Elevation Maps

A digital elevation map or DEM can be seen as a collection of points in a “top view” of the 3D terrain where each point has its own height or “elevation”. The algorithm

proposed for generating regular DEMs in the ROBUST project fills in a “top view” image of the terrain completely, i.e. a height value can be computed for every pixel in the top view image, except for pixels that are not visible in the IH because of occlusions. These occlusions are found in a very simple way.

The terrain is divided into cells: the pixels of the DEM. For each cell the stereo pair image is selected in which the cell would be visible if it had a height of zero. A vertical line is drawn and the projection of this line in the left and right image of the stereo pair is computed. Figure 5 illustrates the algorithm that is used to determine the height of the terrain on that line.

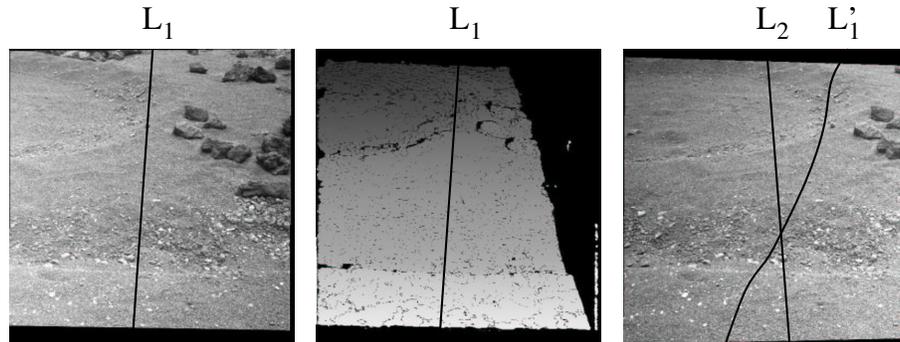


Fig. 5. Digital Elevation Map generation in detail. The figure contains the left rectified image (left), the corresponding disparity map (middle) and the right rectified image (right). The 3D line L corresponding to a point of the DEM projects to L_1 resp. L_2 . Through the disparity map the shadow of L_1 on the right image can also be computed. The intersection point of L_2 and L_1 corresponds to the point where L intersects the surface.

- It is possible to detect occluded regions easily. This is the case for cells that are not visible in both stereo images. The height value of these cells can not be computed and these cells get a certain predefined value in the DEM which marks them as unseen.
- This particular scheme makes it possible to generate regular digital elevation maps at any desired resolution, interpolating automatically if needed.
- For the parts of the terrain close to the boundary of a ring, different parts of the vertical line will be projected in different stereo views. Therefore it is possible that data of two different stereo views has to be combined. This is not a problem because the transformation between the views can easily be computed since the calibration has been calculated.

4 Path Planning

Once the planetary terrain has been reconstructed, scientists can indicate sites of interest (Points Of Reference or PORs) which the rover should visit. Paths have to be

computed between successive PORs. For this purpose, an path planning module has been developed. Given a terrain map, an initial rover position and heading and a desired rover position and heading, the path planner will find a path, using an A* framework, which takes the rover from the initial state to the goal state and which is optimal in terms of a series of parameters, like energy consumption, risk of tip-over, use of tether etc.

4.1 Travel Cost Map

The Travel Cost Map (TCM) provides a measure for the cost of traversal based on metrics inherent to the terrain. In the current implementation, a simple metric based on the gradient of the Digital Elevation Map (DEM) is used. Another metric characterizes the uncertainty of the terrain data, the farther from the lander camera the higher the uncertainty. Areas occluded by rocks also have high uncertainty.

4.2 The Hierarchical Approach

The rover can move according to a set of available operators (also called rover movements), which take the rover from one position and heading (this pair is also known as a state) to another position/heading. Each operator has an associated cost. The main term of this cost is computed from the above mentioned TCM. Given that A* is computationally very complex, finding a path in a reasonably large terrain, using complex operators for the rover movements, can take a long time. This has led to the choice of a hierarchical approach to the path planning problem.

Finding the Corridor At the first stage, a traverse is planned between the start and goal states using A* covering the whole terrain but with reduced resolution, the cells being somewhat larger than the size of the rover so that it can move comfortably within the corridor. A low-resolution TCM is used for this. The transition operators are simple forward, backward, left and right, allowing to apply a highly optimized and fast version of A*. The result is a corridor (see figure 6) in which the rover may safely move.

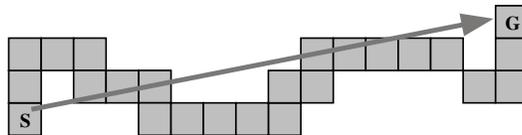


Fig. 6. Corridor

Refinement of the Path At the second stage the path is refined using the high-resolution TCM. By restricting the search to cells marked in the Restrained Grid constructed in the

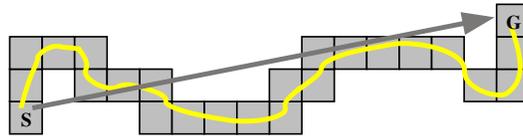


Fig. 7. Refined path

previous stage more complex operators and full available resolution can be used within reasonable time constraints.

The representation of the operators to take the rover from one state to another is kept very general, i.e. a rotation followed by a translation. The cost of applying an operator is determined by using a number of cost evaluation points. The cost is calculated as a weighted sum of the costs at these points, evaluated at the resulting rover pose. The evaluation points represent parts of the 'virtual' rover during an after the completion of the corresponding move. The position of the evaluation points are calculated based on the rover dimensions, the parameters of the rover movement, the desired safety margin and the resolution of the TCM.

The result of the hierarchical A* algorithm is a high-resolution path (see figure 7), represented by an ordered list of rover poses, bringing the rover from its start pose to the desired destination.

5 System test

A first test of the complete system was performed at the ESA-ESTEC test facilities in Noordwijk (The Netherlands) where access to a planetary testbed of about 7 by 7 meters was available. The Imaging Head was set up next to the testbed. Its first task was the recording of the terrain according to the setup shown in figure 2. A mosaic of the pictures taken by this process can be seen in figure 8.



Fig. 8. Mosaic of images of the testbed taken by the stereo head.

The autonomous calibration procedure was launched and it computed the extrinsic calibration of the cameras based on the images. Once the calibration had been computed

the system rectified the images and computed dense disparity maps. Based on these, a Digital Elevation Map was constructed. The result can be seen in figure 9. Because

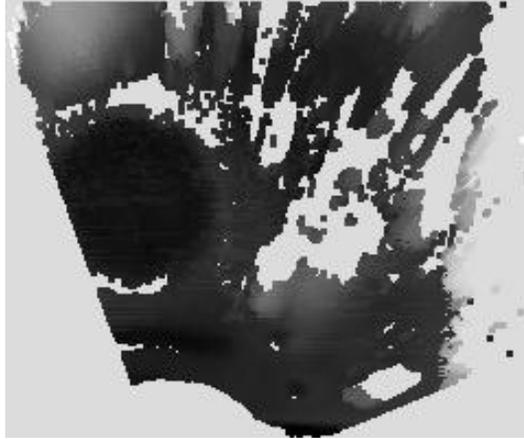


Fig. 9. Digital Elevation Map of the ESTEC planetary testbed. A significant amount of cells is not filled in because they are located in occluded areas.

of the relatively low height of the Imaging Head (approximately 1.5 meters above the testbed) and the big rocks in the testbed, a large portion of the Digital Elevation Map could not be filled in because of occlusions. The Digital Elevation Map of figure 9 was then used to construct a textured triangulated mesh model. Some views of this model can be seen in figure 10. A striping effect is visible in the texture of the model. This is due to a hardware problem of the cameras which caused the images to suffer from an intensity gradient. Since the 3D model combines data from all images, the boundaries are visible.

Once the reconstruction was complete, the model was used to plan a trajectory for the Nanokhod rover. At that time, another component of the system came into play: the **Simulator**. This component allowed for the operator to simulate low-level commands on the Nanokhod and Imaging Head in a Virtual Reality environment. A simple gravitation model was used to predict the pose of the Nanokhod after a motion. An overview of the GUI of this component is shown in figure 11. The GUI is divided into 4 windows. The upper two show what the left and right camera of the stereo head are seeing. The bottom right window is a 3D interaction window and the bottom left window shows a picture of what a camera, mounted in the cab of the Nanokhod would see.

When the trajectory was simulated, the commands were uploaded to the lander system which executed them autonomously. The pose of the rover was observed using the Imaging Head. The lander needed the calibration parameters, computed on the ground station for this and made use of the LED's that are present on the Nanokhod. The telemetry, computed by the lander could then (after downlink to the ground station) be played back by the simulator and compared with the simulated pose of the Nanokhod.

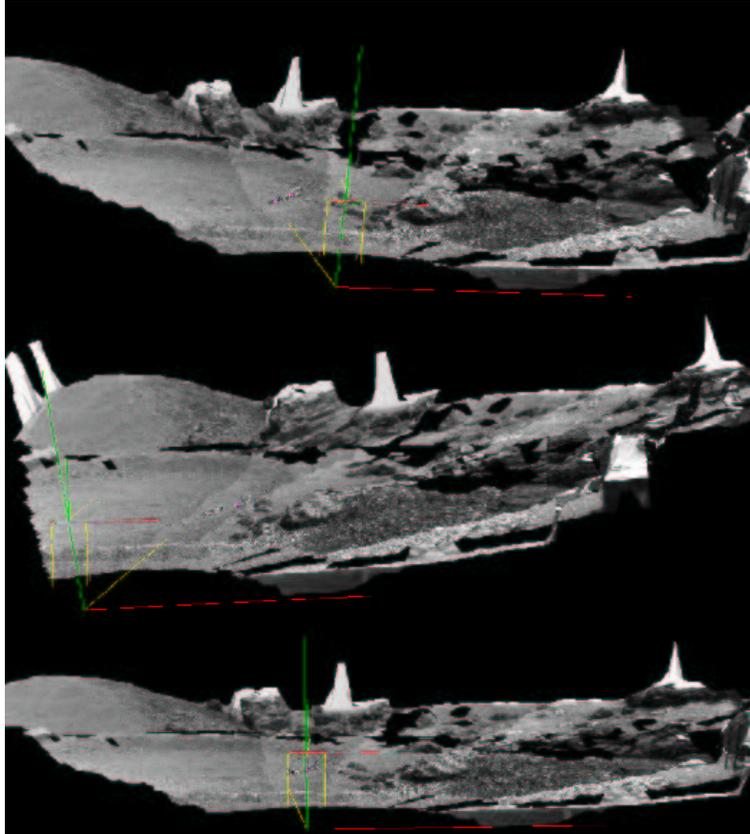


Fig. 10. Three views of the reconstructed model of the ESTEC planetary testbed.

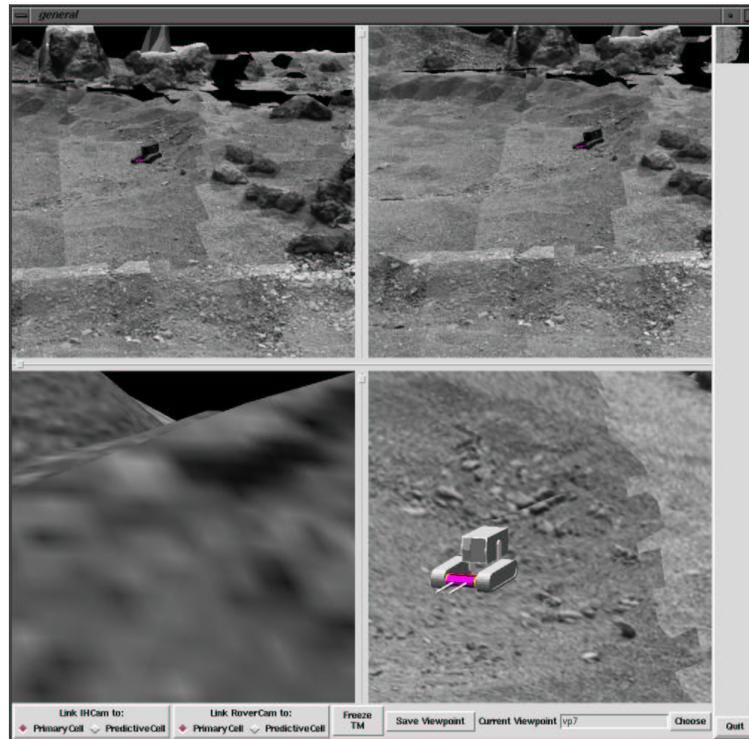


Fig. 11. Overview of the simulator. The upper two windows simulate the view of the two cameras. The bottom right window is a 3D interaction window and the bottom left window gives a view from the cab of the rover.

6 Future development

In the case of the ROBUST project, calibration of the Imaging Head is a critical issue. Because of the known specifics of the Imaging Head, the calibration algorithms can be targeted to this particular setup and take all known information on the mechanics and intrinsics of the system into account. Situations can be imagined where such information is not available. While the algorithms, described in this paper, can no longer be used, it is still possible to retrieve some calibration and 3D reconstruction.

6.1 Uncalibrated 3D reconstruction

In [9] it is described how structure and motion can be retrieved from an uncalibrated image sequence. The 3D modelling task is decomposed into a number of steps. First successive images of the sequence are matched and epipolar geometry is computed, using the same approach as described in section 2.1. The projection matrices of all views are computed by triangulation and pose estimation algorithms. The ambiguity on the reconstruction is then reduced from pure projective to metric (Euclidean up to scale) by self calibration algorithms (see also [7]). Finally textured triangular mesh models of the scene are constructed.

6.2 Experiment

The described technique could be of use during a planetary exploration mission. The most important instrument of the payload cab of the rover is most probably a camera. This camera will make images of samples and rocks on the terrain but could also be used to perform close-range reconstruction of the terrain, helping the rover to “dock” precisely onto the desired position. The camera would take images of the area of interest during the approach. These images could then be used as input for the reconstruction algorithm described before to generate a 3D reconstruction. The resolution of this reconstruction would be far superior to the one of the DEM, obtained from the Imaging Head.

During testing of the *robust* system on the planetary testbed at ESTEC, a preliminary test was performed. Images of some rocks on the testbed were taken by hand with a semi-professional digital camera. The images were processed by the 3D reconstruction system. The resulting 3D reconstruction of the rock can be seen in figure 12. The results are very good and show that this strategy could be an interesting extension of the *robust* system.

7 Conclusion

In this paper an approach for calibration and 3D measurement from planetary terrain images was proposed which allowed for an important simplification of the design of the imaging system of the lander. The components described in this paper are part of an end-to-end system which can reconstruct an unknown planetary terrain and guide a rover autonomously on the planetary surface. The system has succeeded a first test in a planetary testbed at ESTEC.

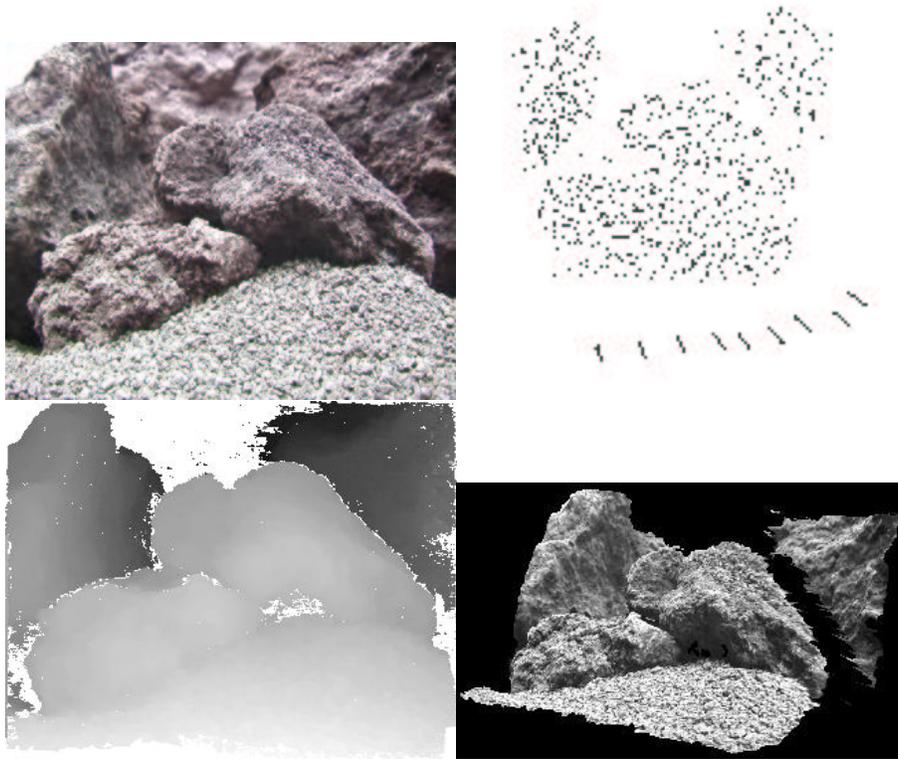


Fig. 12. Results of the reconstruction process of a rock on the planetary testbed. The upper left view shows one of the original images of the sequence. The upper right view displays reconstructed points and cameras. In the lower left view a dense depth map is shown. The lower right view shows an arbitrary view of the 3D reconstruction.

Acknowledgments

We acknowledge support from the Belgian IUAP4/24 'IMechS' project. We also wish to thank all partners of the ROBUST project for the collaboration. Special thanks go to Ronny Moreas of K.U.Leuven-PMA for his contribution on the path planning algorithm.

References

1. I. Cox, S. Hingorani and S. Rao, "A Maximum Likelihood Stereo Algorithm", In *Computer Vision and Image Understanding*, Vol. 63, No. 3, May 1996.
2. M. Fischler and R. Bolles: "RANDOM SAMPLING CONSENSUS: a paradigm for model fitting with application to image analysis and automated cartography", In *Commun. Assoc. Comp. Mach.*, 24:381-95, 1981.
3. C. Harris and M. Stephens: "A combined corner and edge detector", In *Fourth Alvey Vision Conference*, pp. 147-151, 1988.
4. J. Knight and I. Reid, "Self-calibration of a stereo-rig in a planar scene by data combination", In *Proceedings International Conference on Pattern Recognition (ICPR 2000)*, pp. 411-414, Barcelona, 2000.
5. R. Koch, "Automatische Oberflächenmodellierung starrer dreidimensionaler Objekte aus stereoskopischen Rundum-Ansichten", *PhD thesis*, University of Hannover, Germany, 1996 also published as *Fortschritte-Berichte VDI*, Reihe 10, Nr.499, VDI Verlag, 1997.
6. S. Maybank, "Theory of reconstruction from image motion", *Springer Verlag*, 1992.
7. M. Pollefeys, R. Koch and L. Van Gool. "Self-Calibration and Metric Reconstruction in spite of Varying and Unknown Internal Camera Parameters", *International Journal of Computer Vision*.
8. M. Pollefeys, R. Koch and L. Van Gool, "A simple and efficient rectification method for general motion", *Proc. ICCV'99 (international Conference on Computer Vision)*, pp.496-501, Corfu (Greece), 1999.
9. M. Pollefeys, R. Koch, M. Vergauwen and L. Van Gool, "Metric 3D Surface Reconstruction from Uncalibrated Image Sequences", *Proc. SMILE Workshop (post-ECCV'98)*, LNCS 1506, pp.138-153, Springer-Verlag, 1998.
10. R. Rieder, H. Wanke, H. v. Hoerner. e.a., "Nanokhod, a miniature deployment device with instrumentation for chemical, mineralogical and geological analysis of planetary surfaces, for use in connection with fixed planetary surface stations", In *Lunar and Planetary Science*, XXVI, pp. 1261-1262, 1995.
11. C. Loop and Z. Zhang. "Computing Rectifying Homographies for Stereo Vision". IEEE Conf. Computer Vision and Pattern Recognition (CVPR'99), Colorado, June 1999.
12. A. Zisserman, P. Beardsley and I. Reid, "Metric calibration of a stereo rig", *Proceedings IEEE Workshop on Representation of Visual Scenes*, Cambridge, pp. 93-100, 1995.