

A Micro-Rover Navigation and Control System for Autonomous Planetary Exploration

K. Landzettel, B.-M. Steinmetz, B. Brunner, K. Arbter *
M. Pollefeys, M. Vergauwen **
R. Moreas, F. Xu ***
L. Steinicke, B. Fontaine ****

* *DLR Oberpfaffenhofen, Institute of Robotics and Mechatronics, Wessling, Germany*

** *KU Leuven, Department of Electrical Engineering, ESAT-PSI, Heverlee, Belgium*

*** *KU Leuven, Department of Mechanical Engineering, PMA, Heverlee, Belgium*

**** *Space Applications Services, Zaventem, Belgium*

Full Paper - Manuscript

Keywords: autonomous space robots, navigation, terrain modelling, path planning, ground control

ABSTRACT

This paper describes an end-to-end control system for autonomous navigation of a small vehicle at a remote place, e.g. in space for planetary exploration. Due to a realistic background of this study the proposed method has to deal with limited knowledge about the environment as well as limited system resources and operational boundary conditions, especially very large time delay in the communication between the ground control station and the space segment. To overcome these constraints the remote system has to act in a very autonomous way. Ground support minimizes the computational load of the remote system. High-level information interchange reduces the needs in communication bandwidth.

1 Introduction and Overview

In addition to manipulator-based robotics in near-Earth orbits, autonomous planetary exploration will play an important role in future space missions. After the success of the Mars Pathfinder Mission a lot of work has been started to overcome the limitations of the Sojourner Rover, especially w.r.t. the lack of local autonomy [2],[3],[4]. Also ESA has initiated the so-called Payload Support for Planetary Exploration (PSPE) project to study the feasibility of robotized planetary exploration. It includes a Lander spacecraft configuration, which should perform geo-science operations e.g. on Mars [5].

The role of the end-to-end system is to allow an investigator to remotely perform geo-science experiments on a planet, e.g. Mars or Mercury [6], using a small rover. This implies the building of a system based on the following three components:

- The Lander system (see Fig. 1) carries all of the supply engineering and the complete control system for the Rover, the scientific instrumentation and the communication with the Ground Control Station. Depending on the mission additional scientific equipment can be integrated.
- The Imaging Head (see Fig. 2), mounted on top of a vertical cantilever rod coming out of the Lander, is equipped with a space-qualified stereo camera and a two degrees of freedom pan-and-tilt unit. The cameras are optimized for both taking stereoscopic panorama images of the landing site as well as the detection of interesting objects around the Lander. The Imaging Head enables the tracking of rover motion and supports the determination of its relative position and attitude of the rover with respect to the Lander itself. It captures stereo images required for generating a map of the terrain around the Lander. The Imaging Head is - when the cantilever rod is fully extended - some 1.5m above the Lander. With the aid of the Imaging Head the On-Lander Control System performs the rover localization. The Imaging Head will be moved so that its two cameras follow the rover position and evaluates the

rover position from the stereo views. As the second task it acquires images of the environment around the Lander, to allow the Ground Station to reconstruct a Terrain Model.

The Nanokhod Rover (see Fig. 3) is a rugged, simple, reliable yet effective, track-driven carrier of scientific equipment, accommodated in a 'Payload Cab' (PLC), optimized to traverse the terrain to interesting points in the vicinity of the Lander, and to carry out in-situ measurements. Tethered to the Lander via a cable the Nanokhod makes use of the Lander's power and data handling resources rather than carrying its own. This approach results in a very high payload to total-mass ratio of the Rover, but also limits the operational range to a maximum distance of approx. 20 meters around the Lander.

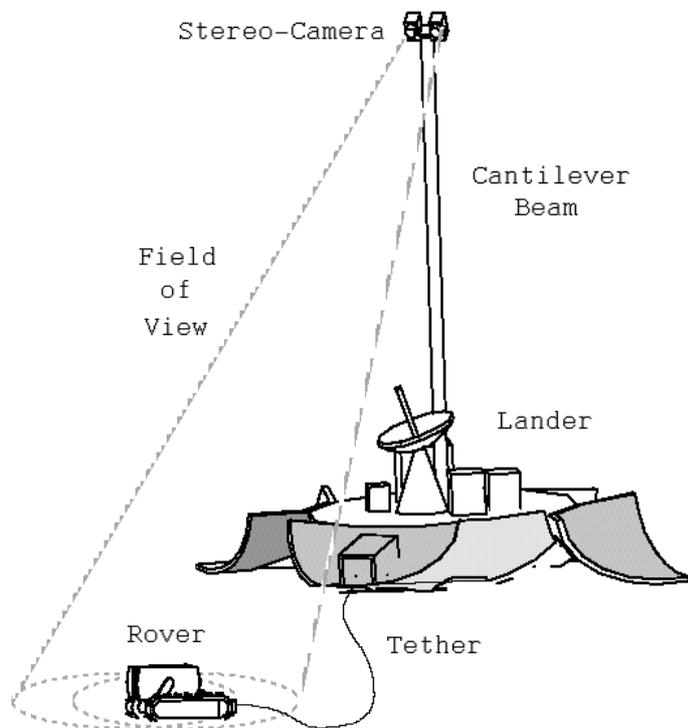


Fig. 1 Lander system with Imaging Head and Nanokhod Rover

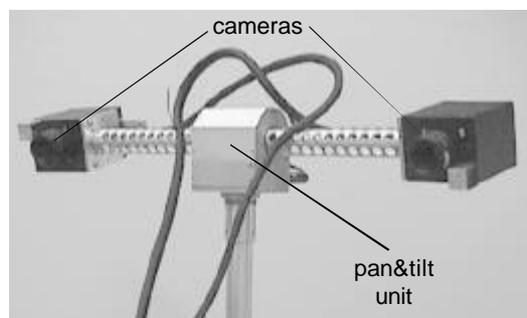


Fig. 2 Imaging Head with pan-tilt-unit

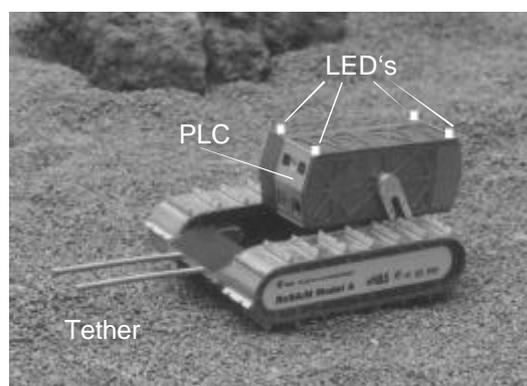


Fig. 3 The Nanokhod rover

Due to the very large time delay concerning the data link between ground and space segment (transmission times up to as much as 20 minutes) and typical limitations of communication bandwidth in space, on-line control is not feasible. That means, the Rover has to be commanded at a very abstract level: the On-Ground Control System [7] (OGCS) provides an easy-to-use command interface, optimized for intuitive scientific experiment support without the need to have specific knowledge in the field of robotics. The basis for all operations is a terrain model, reconstructed on-ground from the stereoscopic panorama images taken by the Imaging Head all around the landing site. A mission scientist selects interesting sites or objects in the terrain model, to be explored by the Nanokhod – the so called “long arm” of the Lander. In our approach a list of way-points, determined by the path planner on-ground [8], will be uploaded to the space segment for autonomous execution on site. Whereas most of the efforts has been done in the field of path planning [9], the need of an on-line navigation technique based on mission specific components has not been discussed in depth so far.

To fulfil this gap, a mission scenario has been designed for this study, which focuses on the autonomous motion control of the small Rover vehicle. The essential element to achieve this autonomy is the precise localization of the Rover without advice from ground and the capability to cope with non-nominal situations by itself. Due to unknown parameters, e.g. soil characteristics, a dead reckoning approach, based for example only on odometry data, would fail. To guide the Rover on its way, a more robust 3D-localization technique is necessary.

2 On-Ground Control System

The On-Ground Control System (OGCS) is based on the FAMOUS robotics framework leading to an actual implementation called FAMOUS/Nanokhod. It allows the control of the rover based on a layered framework, defining six levels of control abstractions. The mission scientist defines the experiments by selecting the corresponding Compound Tasks

(the highest level of control abstraction). The FAMOUS/Nanokhod then prepares, verifies and validates the Compound Tasks and sends them to the On-Lander Control System for execution.

The OGCS mainly provides an easy-to-use command interface, optimized for intuitive scientific experiment support without the need to have specific knowledge in the field of robotics. The OGCS generates the terrain model from the stereoscopic panorama-images taken by the Imaging Head. A sophisticated path planning algorithm determines the optimal path paying careful attention to given constraints (e.g. topography, estimated soil and known Rover characteristics).

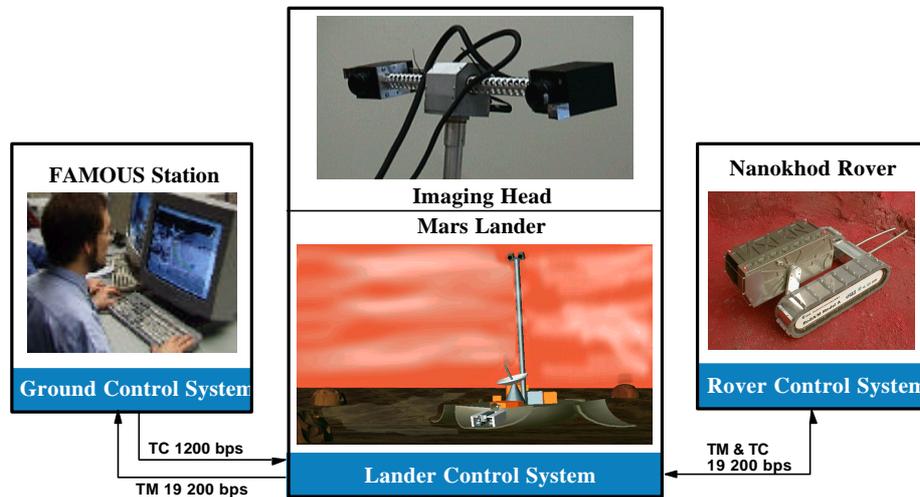


Fig. 4 Communication links between the subsystems

The On-Ground Control System is responsible for

- the definition of the initial terrain model;
- the definition of the pre-defined rover programs (including Actuators, Actions, Tasks, and Compound Tasks) and their pre-validation through simulation.

The corresponding Actuator, Actions, etc., are uploaded to the OLCS on a permanent basis, and are triggered by appropriate commands during the mission operations phase.

The OGCS is configured with an initial Digital Elevation Map, obtained e.g. from a previous survey mission, or a Test Terrain representative of the difficulties that may be encountered for traversing the kind of terrain: anticipated zones of non visibility or limited precision of localization, obstacles that cannot be overcome, etc.

During the preparation phase, the operator successively defines and validates through simulation the Actuators, Actions, Tasks and finally Compound Tasks routines. The approach is bottom-up: from the lowest level of control (Actuator) to the highest level (Compound Task).

Because, during the preparation phase, the current Terrain Model is not known, the routines of the Actuator and Action level can be checked on a test terrain model only, using a series of "worst case" (most severe) utilization scenarios.

3 Calibration

During launch and landing, the Lander and its contents are subject to extreme forces. The mechanical properties of the Imaging Head when the spacecraft arrives at the planet are hence 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.

3.1 Using markers?

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 cameras [10].

There are two reasons why this scheme is not suitable in this case where the Imaging Head is deployed on a distant planet. First, there is the problem of where to place the calibration objects. 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 consists 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 renders the markers much harder to use. Even if one did dispose of accurate markers on the Lander, a second problem arises. To maximize robustness, the cameras are not equipped with any zooming or focusing system. Since the accuracy of the stereo matching decreases with the square of the distance, the cameras are focused on infinity to gain as much accuracy in the far regions as possible. As a consequence, the images of very 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 therefore clear that standard calibration algorithms can not be used in this system. A new strategy has been developed that only uses images of the terrain to calibrate the Imaging Head.

3.2 Strategy

The calibration procedure that has been implemented is able to calibrate the Imaging Head using images of the terrain only. This means that the images which are transmitted 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 overhead on data 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. This is explained in the following section. Once this relative calibration has been computed, 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.

3.3 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).

3.3.1 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 (see Fig. 5). If one wants to find back the epipolar geometry between two images automatically, a filter, called the Harris Corner Detector [11], is applied to the images first.

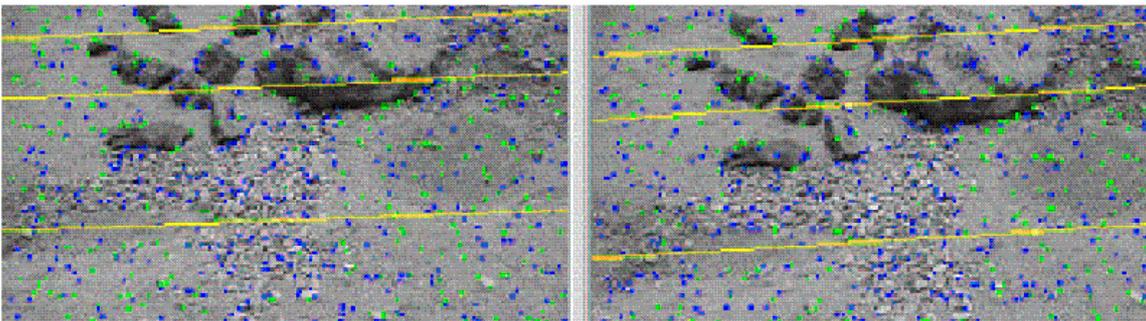


Fig. 5 Epipolar geometry of an image pair

The result consists of points or *corners* in the images determining where the image intensity changes significantly in two orthogonal directions. 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 [12], is used to compute and update epipolar geometry and matches iteratively.

In the case of our Imaging Head the images 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 and therefore the epipolar geometry does not change.

It is even the case that a specific degenerate case for the computation of the epipolar geometry is solved by this combination scheme. Computing the epipolar geometry of a pair of images of a *planar* scene is impossible from correspon-

dences 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.

3.3.2 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 constructed. This is easily done since:

$$E = K^T FK$$

with K the 3x3 matrix with the intrinsic calibration of the cameras. To derive the relative translation and rotation from the essential matrix, we refer to the work of Maybank et. al. [13].

There is one parameter that cannot be calibrated, namely the actual value of the baseline. We can, however, assume that this value will not deviate much from the mechanical specs. If there were some change in the actual value of the baseline, the consequences of fixing it to the (wrong) value of the specs are not serious because all measurements (terrain reconstruction and Rover localization) are done within the same measurement system.

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

3.4 Pan-tilt calibration

Computing the relative transformation between the cameras is an important part of the calibration but it does not suffice. For Rover localization and generation of terrain reconstructions the transformations between the cameras and the Imaging Head and between the Imaging Head and the Lander needs to be known as well.

3.4.1 The Imaging Head frame

For sake of clarity a virtual "Imaging Head frame" is defined in "the middle" of the two cameras. This means that the relative translation and rotation between the left camera and the Imaging Head frame is equal to the translation and rotation between the Imaging Head frame and the right camera.

3.4.2 From Imaging Head to Lander

Calibrating the relative transformation between the Imaging Head frame and the Lander is more complicated because it implies calibration of the pan and tilt axes. It is clear that this transformation depends on the actual angle of rotation around both the pan and tilt axes. From the world's point of view, the motion of the Imaging Head can be described as a rotation around the pan axis followed by a rotation around the tilt axis. The pan axis is never altered but the orientation of the tilt axis depends on the pan angle.

If we look from the point of view of the Imaging Head, however, it is the tilt axis that never changes and the orientation of the pan axis depends on the tilt angle. The latter view will be employed because it fits very well in the philosophy where one derives the entire chain of calibration transformations from the cameras, which are the only measurement device, to the Lander.

3.4.3 Relative transformations between views

To calibrate the pan and tilt axes, stereo images of the same ring and the same segment are used respectively. Especially the overlap between consecutive stereo images is important in the strategy.

3.4.3.1 Calibration of the Tilt axis

For the calibration of the tilt axis, a stereo image of the outer ring of a certain segment is recorded. The Imaging Head is commanded to execute a tilt motion and to record a stereo image of the second ring. One has to make sure that there is sufficient overlap between the two image-pairs.

This setup is shown on the left in Fig. 6. The area shaded from bottom left to top right is visible in the first view. The area shaded from bottom right to top left is visible in the second. A planar scene with texture from a real image from Mars was constructed and pairs of images were generated with a visualization toolkit.

Corresponding features in the images of the first image pair can be found as explained in Section 3.3.1. Because we know the relative transformation between the two cameras, we can reconstruct the features in 3D. The same is done in the second image pair. Because of the overlap, some of the features will be visible in both image pairs. We can find correspondences between these features by running the matching algorithm of Section 3.3.1 on the two images of the left or the right camera.

The corresponding features allow us to align the reconstruction of the second pair with the reconstruction of the first pair. This yields the relative transformation between the first and second IH frame.

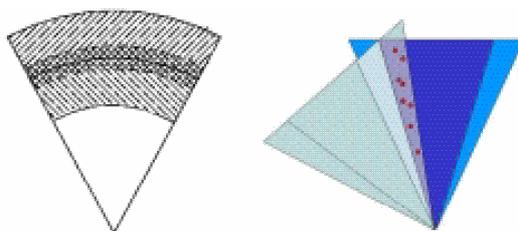


Fig. 6 Symbolic representation of the setup for the computation of a relative transformation for a tilt motion(left) and for a pan motion (right)

3.4.3.2 Calibration of the Pan axis

For the pan axis, the computation of the relative transformation between two views is slightly different. The set-up is shown on the right in Fig. 6. It is clear that in this case there are almost no features that are present in all four images of the two views. Due to the vergence of the two cameras, however, there are some features that can be seen in one stereo view and in one image of the other image pair. These features are represented by a dot in the right image of Fig. 6. Again we find back corresponding features between the left image of the first pair and the right image of the second pair with the algorithm of section 3.3.1. Because the features are visible in both images of the first stereo view, they can be reconstructed in 3D. They are also visible in one image of the second view, so one can apply a pose-estimation of the camera of the second pair in which the features are visible, yielding the pose of this camera in the frame of the first view. Using the relative transformation between camera and IH from section 3.3.1, we have found back the relative transformation between the two stereo views.

3.4.4 Actual calibration of pan and tilt axes

The previous section provides us with a set of relative transformations between Imaging Head frames. Because that each of these transformations comes from a pure rotation, the rotation axes can easily be computed. Because, from the point of view of the cameras, the pan axis changes according to the tilt angle, one first has to "undo" the influence of the tilt rotation before one can use the relative transformation to compute the pan axis.

3.4.5 Iterative procedure

During the acquisition of the data one tries not to change the pan angle if a pure tilt rotation is executed and vice versa. In any real system, however, there will be deviations from the desired angles. This means that the computation of the tilt axis will not be correct because the linear algorithm computes the real rotation axis, which is not the tilt axis if there is an - even small - pan component. But there is a solution to this problem. In the second step a good approximation of the pan axis was found, so if we account for the small deviations of pan with the current computed value of the pan axis, we can recompute the tilt axis more accurately. This in turn allows us to update the pan axis etc. We can repeat this iterative procedure until the solution for the axes has converged. In reality, three iterations are sufficient.

4 3D Terrain modeling

After the calibration of the IH has been performed, the process of generating a 3D model or models of the planetary terrain can commence. This modeling is vital to accomplishing 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 this section we will describe the different steps that are performed to obtain such a DEM.

4.1 Generation of disparity maps

On an image pair recorded by the Imaging Head, a stereo algorithm is applied to compute the disparity maps from the left image to the right and vice versa. 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 [14].

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 to 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 [15]. 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 constraints in the dynamic programming approach. The algorithm was adapted to yield sub-pixel accuracy by employing a quadratic fit of the disparities.

4.2 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". Classical approaches to generate DEMs from disparity maps or depth maps consist of two steps. First, for each stereo image pair the disparity images are used to construct depth images. These are images with the same size as the original images. The value of each pixel corresponds to its depth. Then a limited amount of points of each depth image is reconstructed in 3D. These points form the DEM. The problem of this scheme is that the resulting DEM is not regular in 3D.

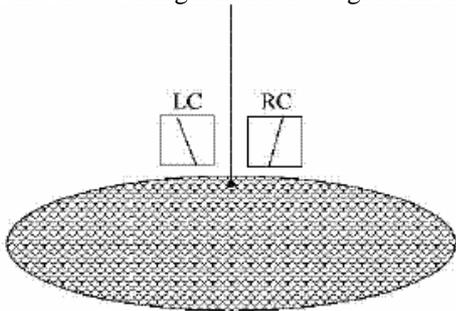


Fig. 7 Set-up of the DEM generation. The LC and RC squares represent the left and right image of an image pair, respectively.

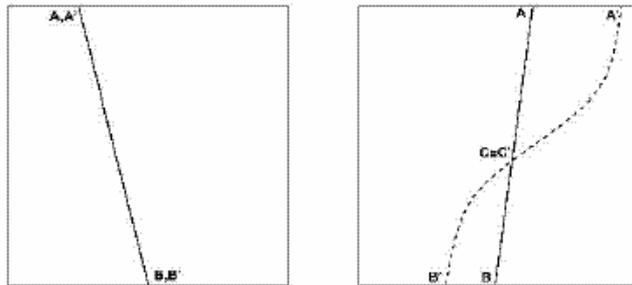


Fig. 8 DEM generation in detail: left and right disparity image with the projection of the vertical line and the "shadow"

The algorithm proposed for generating regular DEMs 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 because of occlusions. These occlusions are found in a very simple way. The principle of the algorithm is illustrated in Fig. 9.

The terrain is divided into cells. For each cell the stereo pair is selected in which the cell would be seen if it had a height of zero. A vertical line is drawn and the projection of this line in the left and right disparity image of the stereo pair is computed. Fig. 8 illustrates the algorithm that is used to find the height of the terrain on that line. The solid line $A-B$ is the projection of the vertical line in both disparity images.

Now imagine, placing a light where the left camera is. This light shines on the vertical line which throws a shadow on the terrain. In the left image this shadow of course has the same projection as the line itself. In the right image, however, this is not the case. The projection of the shadow in this image is the smooth curve from A' to B' . The part of this curve from A' to C' is the *real* part of the shadow (i.e. it would be visible on the terrain). The part from C' to B' can be seen as the *virtual* part of the shadow, coming from the part of the vertical line below the surface of the terrain. This shadow-curve can be computed using the disparity in the left disparity image of every pixel of the projected line $A-B$. The intersection point C of the vertical line and the terrain can then be found as the point where the shadow $A'-B'$ intersects the line $A-B$.

Occluded regions are easily detected since in this case no intersection point C exists. The height value of occluded cells can not be computed and these cells get a certain value in the DEM which marks them as unseen. This particular scheme also 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 calibration procedure is carried out by the OGCS after the necessary images have been transferred down to Earth. The compressed stereo image files are transmitted to Earth with a communication protocol that takes into account the long end-to-end delay and the lack of reliability of the space links. The image files are routed to the OGCS. In an actual mission, this is achieved typically via the interface services of a Mission Control: the communication protocols between the OGCS and the Mission Control are indeed usually different from the ones between the Mission Control and the Lander. Such aspects have not been considered in this project. On the basis of images of the site, the OGCS can now build a Digital Terrain Model (DTM). This model is actually the collection of a Digital Elevation Map (DEM), a Triangulated Mesh Model (TMM), a texture map and thematic maps.

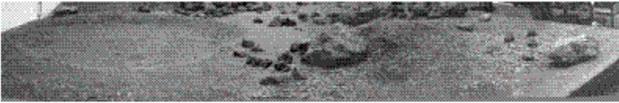


Fig. 9 Panoramic image of the ESTEC testbed

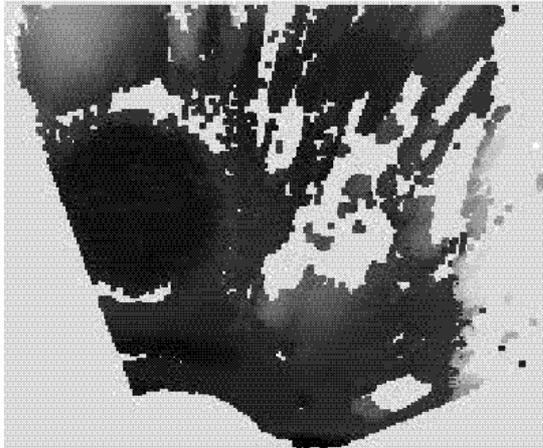


Fig. 11 Resulting DEM of the testbed

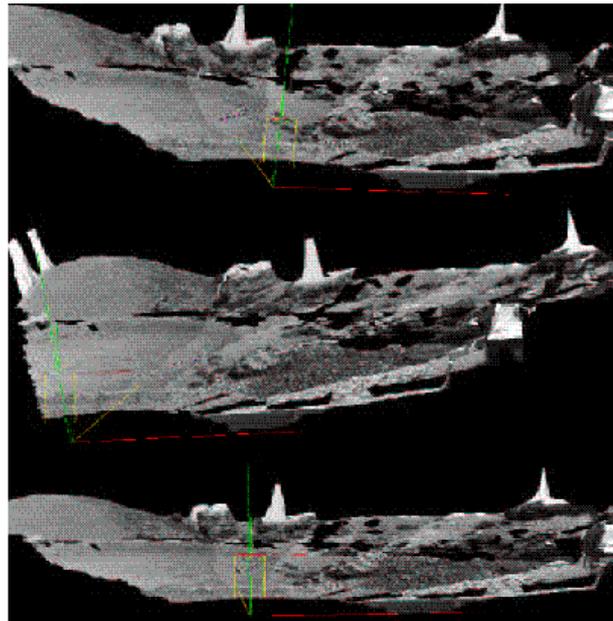


Fig. 10 Resulting views of the reconstruction

Fig. 9 shows a panoramic image of the planetary testbed at ESTEC in Noordwijk (NL), where the system was tested. In Fig. 11 the resulting digital elevation map (DEM) is shown. There are quite some (white) gaps in the map that are due to occlusions. The resulting DEM can be triangulated and textured to yield a 3D reconstruction of the terrain (Fig. 10.).

5 Path Planning

Once the Lander is on the planetary surface and the first image data are available to the scientists, a first rough planning of the overall mission is performed, identifying the sites which the scientists want to see at any rate during the mission. This is to avoid that the mission starts, expelling more and more tether, and at the end, there are some sample sites left in a sector that has not yet been explored when the tether is used up.

5.1 Definition of the Points of Reference

The scientists must therefore define at the very beginning of the mission, the broad exploration concept: what is the overall rover exploration path around the Lander to which the actual site selection will be "superimposed". This is done by defining Points Of Reference (PORs), which are "flags" to mark the most important science sites. The actual exploration path will jitter between PORs, depending on day by day science interests, as they are obtained science results, and resource availability.

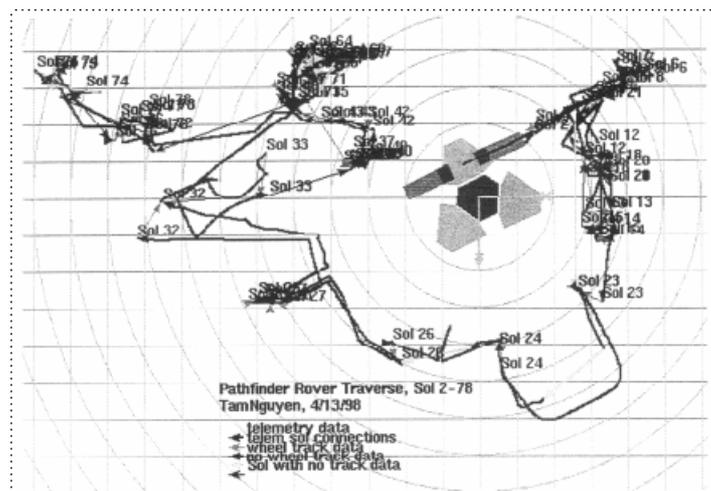


Fig. 12 Sojourner trajectories on Mars, following a more or less circular movement

The primary scientific objective behind usage of the Nanokhod rover is to achieve a geoscience mission by determining the composition of a series of rocks in particular areas. As a consequence, the corresponding broad concept will most likely be a kind of circular movement in an annular segment around the planetary Lander, with local radial excursions. Similar to the approach in the Sojourner/Pathfinder mission (see Fig. 12), the rover follows trajectories composed of many straight lines, which roughly will follow the chosen circular segment, with some local excursions.

Around this reference circular segment, the Nanokhod is expected to visit about 20 sites of scientific interest.

For the definition of the broad exploration path, the OGCS provides the Mission Interface. This is a graphical view, in which the terrain is visualized but not necessarily with the highest fidelity. The texturing, however, should be as detailed as possible to allow for a good human-oriented presentation.

Upon request, a "Red-Orange-Green" (ROG) concept is applied to clearly indicate the safe zones (green), the risky zones (orange), and the forbidden zones (red). A scientist chooses the sites (s)he wishes to visit, preferably in a green zone, and by associated information to this point (e.g. a priority and a comment).

To prepare the rover movement on Mars, a path planner on the FAMOUS/Nanokhod station is used. The path planner is capable of finding paths which minimize risk of tipping over, risk of entangling the tether, length of tether used, time to traverse, risk of sliding due to slopes and poor soil contact, and risk of getting stuck in loose soil.

The desired optimization criteria can be selected by the OGCS operator. If more than one criterion is selected, the weighted sum of the corresponding cost functions is minimized. To minimize the risk of collision with obstacles, the path planner takes into account localization errors and avoids areas that are invisible to the Imaging Head cameras.

The path planner computes a route between the current position and the desired site. This route consists of a sequence of path segments. Given the 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* algorithm, which takes the rover from the initial state to the goal state and which is optimal in terms of the mentioned parameters.

The general principle for the algorithm of the path planner is based on the construction of a corridor by using A* on a grid map covering the whole terrain but with reduced resolution, the cells being somewhat larger than the rover. With this first corridor, a refinement of the path is performed with a higher resolution map.

5.2 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 used 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. Soil characteristics are taken into account by the ROG Map.

5.3 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. Moreover, given that the path execution can be expected to be less than perfect, it is not desirable to find a path which is optimal, whilst closely surrounded by a difficult area. Considering both these facts has led to the choice of a hierarchical approach to the path planning problem.

5.3.1 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 Nanokhod so that the Rover can maneuver com

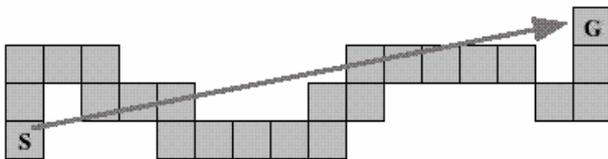


Fig. 13 Path corridor



Fig. 14 Refined path

fortably 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 Fig. 13) in which the rover may safely move. The cells belonging to the corridor are marked in a Restrained Grid which is used in the second stage of the algorithm. The user has the ability to mark additional cells to enlarge the search space.

5.3.2 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 Restrain Grid constructed in the previous stage more complex operators and full available resolution can be used within reasonable time constraints.

The Rover is abstracted as a directed point, representing its position and heading. 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. As such, the set of operator is very customizable. 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 and 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.

5.3.3 Segmentation of the Path

The result of the hierarchical A* algorithm is a high-resolution path (see Fig. 14), represented by an ordered list of Rover poses, bringing the Rover from its start pose to the desired destination. This representation must be converted to a Path Segment List (PSL) which can be executed by the Rover. The PSL is a sequence of straight path segments associated with a Rover action. An iterative procedure is used to approximate the high-resolution path taking into account the traversability along the path.

The high resolution path is not directly usable and needs to be decomposed into a collection of segments and associated way points. This leads to candidates for the path segments. These candidates are validated through simulation, if it is possible to traverse them with single piloting actions (i.e. the rover does not tip over, does not overconsume a resource, etc.).

6 Navigation

The navigation relies on a vision based approach for Rover localization and guidance by the stationary Lander module [16]. The on-board image processing and 3D-localisation system detects significant features of the Rover in the stereo images, determines the current position of the rover w.r.t. to a reference system, and controls the Rover motion to reach the desired target position. To make the feature detection much easier, markers are mounted on top of the Rover's payload cab (PLC, see Fig. 3). Using a sophisticated blob finding algorithm, applied to the difference images (e.g. LED switched off/on), the 2D-coordinates of the markers in the image plane can be calculated. To determine the pose of the Rover, first a stereo reconstruction algorithm generates 3D marker coordinates corresponding to the given pairs of 2D image coordinates, second a matching algorithm is applied, which matches these "measured" 3D markers with the "modeled" ones to get the transformation between the rover and the camera system.

A trajectory control method enforces relative motion corrections to keep the Rover on its desired path: the differences between the measured path position and the desired one will be within a given "on-track margin". Two control modes are applied:

1. Inverse proportional to its distance to the next waypoint, only the Rover's orientation is controlled. Temporary deviations from the reference path will be accepted.
2. In the vicinity of the waypoint, the pure heading control will be expanded by a precise Rover/waypoint matching strategy.

The limiting factor for autonomous navigation is the camera resolution: the more the distance between the cameras and the rover increases, the more the markers' resolution will decrease.

Furthermore, the navigation task is characterized by the following:

- Due to the usage of a tether the action radius of the Nanokhod rover is hard-limited.
- The path control is dominated by discontinuities which abruptly affect the motion behavior in an unpredictable way.
- Due to limited computational power and severe requirements in simplicity and robustness the localization method is based on an active marker concept for vision-based evaluation.

Also, the unknown interaction between the tracks of the rover and the local soil has an unpredictable impact on any command execution, especially for rotations. Although a fixed and undisturbed constellation exists between the planned way points and the Lander, it is not possible to define a precise mathematical model of the path which fits all of the required characteristics in advance. And, finally the formation of the heterogeneous topology, which has to be traversed by the Rover, strongly depends on the direction of motion and cannot be extrapolated based on information from a motion history. Due to those constraints and a mission specific separation of the system into exclusive path planning on ground (OGCS) and autonomous navigation on board (OLCS), the planning operation has to identify extra space for a local refinement of any planned motion.

6.1 Path segment classification

As mentioned in section 5, each path consists of a list of way points. Two consecutive way points mark a path segment. Each path is built up of a finite number of path segments.

Due to the lack of creating a precise path model, the space around each path segment has to be roughly classified using a-priori knowledge of the controlled system (i.e. the characteristics of the Rover and the identified terrain). This firm segmentation marks a region of allowed autonomous activities (see Fig. 15 space between On- and Off-track margin).

On the other hand, within the On-track margin no intervention is necessary and outside of the Off-track margin any autonomous correction is forbidden due to an incalculable risk. By definition each path, which does not violate the Off-track margin, is equivalent to the original planned path.

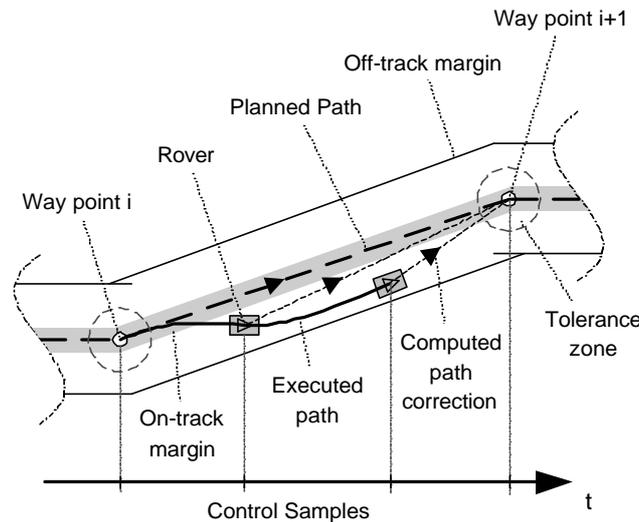


Fig. 15 Path segmentation

6.2 In-situ event classification

In addition to tracing the superior navigation goal, 'In-Situ Classification (ISC)' analyses the Rover behavior in principle. It is mandatory for any local autonomy, which has to deal with unmodeled path execution events, to identify the history as well as the impact of problems w.r.t. the desired navigation goal in order to compensate them. A well-proven method to find discrepancies, is an on-line 'Model-Based Predictor-Corrector (MPC)' approach. While executing the planned path both the progress in motion and the Rover behavior itself will be predicted taking into account the current Rover state and command set. The comparison of that predicted Rover state with the measured one allows fault-detection as well as verification of fault-hypotheses. Additionally, a fault-expansion itself can also be extrapolated accordingly to a fault-propagation law (based on the Rover characteristics). Such a fault tendency is useful in rating the actual fault-severity.

Having preliminary fault estimation the corresponding strategy for recovery is executed in order to overcome the problem. This attempt at correction does not guarantee immediate success. Maybe a set of several iterations will be necessary to locate the real reason of the problem. For that fault analysis ISC takes hold of the complete system state (Lander, Imaging Head, Rover). Additionally, it is capable of reconfiguring system settings for getting a better insight.

In a nutshell: ISC is intended to look on-line for the first-fitted strategy, which enables the Rover to proceed and eventually to attain the navigation target.

6.3 Strategy pool

Typical sources of problems can be described by three failure categories:

- Systematic errors within a device characteristic, which can be modeled as well as directly compensated for by a local control loop (e.g. track synchronism).
- Lack in accuracy or system malfunction due to unmodeled properties, which can or cannot be directly measured (e.g. slippage between Rover and soil, especially while turning).
- Unpredictable events caused by an accidental combination of trouble sources, which cannot be measured (e.g. the Rover being stuck as a result of the Payload Cab getting into contact with hard soil or rock).

For each of these types of problems, the 'Pool of Strategies (**POS**)' comprises at least one recovery method. However, each strategy is highly dependent on both the system characteristic and the system instrumentation. For that reason a detailed strategy description is beyond the scope of this paper.

Just a few words to the principle of the strategy based system intervention:

The main goal in navigation and therefore the basic strategy is the minimization of the distance to a given point. Compliant to that requirement each methodology is applicable, if it supports that goal, keeps given boundary conditions and does not jeopardize the system integrity. The operating costs of a strategy have always to be chosen proportional to the desired quality and system state, and inversely proportional to the remaining control error. In case of a risky ambiguity of results a final decision-maker has to be in the loop (typically the operator at the OGCS).

Having these rules in mind, the proposed set of strategies can be split into two categories:

- On-line supervision of Rover motion tendency (reduced accuracy and low computational load).
- Off-line Rover localization with maximum precision (no motion allowed while taking measurements) and high computational load.
- Reconfiguration of Rover kinematics to improve visibility of the features within the localization process.
- Adaptation of system parameter w.r.t. the progress in the navigation process (e.g. Overshoot avoidance).
- Overcome obstacles which cause local disturbances.

6.4 Principle of the Rover position correction

Whatever strategy has been selected as a hard and fast rule, the Rover will be reoriented first towards the next way point via extra rotational commands and high precision result verification. Thereafter the Rover automatically proceeds its linear trek to the next way point (see Fig. 15: Computed path correction). The reason for that sequential procedure is explained by the high intrinsic sensitivity of Rover motion to angular errors.

The control itself focuses on the dominant degrees of freedom x_{trans} , y_{trans} (i.e. parallel to the surface) and z_{rot} (perpendicular to the surface) defined in surface coordinates. The remaining degrees of freedom are supervised only to detect critical situations (via ISC). Each time a new Rover position is obtained, the running motion command will be superimposed by a new command taking into account the current deviation from the planned path. This control concept guarantees that the Rover motion will be iteratively redirected towards the next way point while traversing a path segment (see Fig. 15). A path segment has been successfully finished as soon as the Rover has reached a predefined tolerance zone at its end.

6.5 Image Processing

The most important point throughout navigation is the precise localization of Rover. As a consequence of the available instrumentation a vision-based 3D localization technique is proposed. Because Imaging Head based applications are typically scheduled for the very beginning of planetary exploration, the reuse of the Imaging Head for navigation is an attractive utilization of free Lander resources later on. Moreover a typical Imaging Head characteristic chosen for scientific usage does also meet the navigation requirements.

In case of vision-based localization there are two possible methods for feature extraction:

- A 'Passive Feature (**PF**)' concept, which allows a most probable recognition of pattern primitives such as edges, corners, circles etc., through a statistical image evaluation.
- An 'Active Marker (**AM**)' concept, which yields a most probable identification of synthetic on-off features in a sequence of camera views.

	PF	AM
Computational load	high	low
Image segmentation	complex	simple
Sensitive to	Illumination	Marker properties
Result quality	fail-safe	accident sensitive

Table 1 Comparison between passive and active marker approach.

As soon as the feature extraction has been completed, an appropriate method for model-based 3D pose estimation has to be executed [17].

Due to the limitations of an operational system for a real mission scenario, the AM concept is preferred because of its lower computational load. The following section describes the AM detection and section 6.6 the proposed Rover localization method in more detail.

6.5.1 Principle of the AM localization method

As a mandatory extension to the basic configuration, a set of four AM's was mounted on top of the PLC (LEDs in Fig. 3) at well-known coordinates. They are used as individually controllable light beacons within the scope of feature detection. The detection method itself is based on a set of two sequential images representing the same scene with unchanged

camera view, but with different AM state (flashing and not flashing). Assuming a strong contrast between the flashing and non flashing marker a 'Difference Image Analysis (**DIA**)' yields the location of the marker in the coordinates of the considered camera head.

Moreover the critical correspondence problem in the field of stereoscopic image evaluation is drastically simplified due to the a-priori known, local identity of a punctual synthetic feature within the view of both cameras. I.e., the feature detection method is simply executed twice, for each camera independently, with the same AM flashing sequence. Later on two corresponding image points in 2D can be converted to an equivalent 3D point via triangulation (back-projection). To do so, the correct knowledge of external and intrinsic parameters of the Imaging Head is mandatory. Both are available the intrinsic due to an extensive calibration at pixel level (done by the camera manufacturer DLR-Berlin), as well as the externals due to an in-situ calibration using the assembled Imaging Head as it is (see section 3).

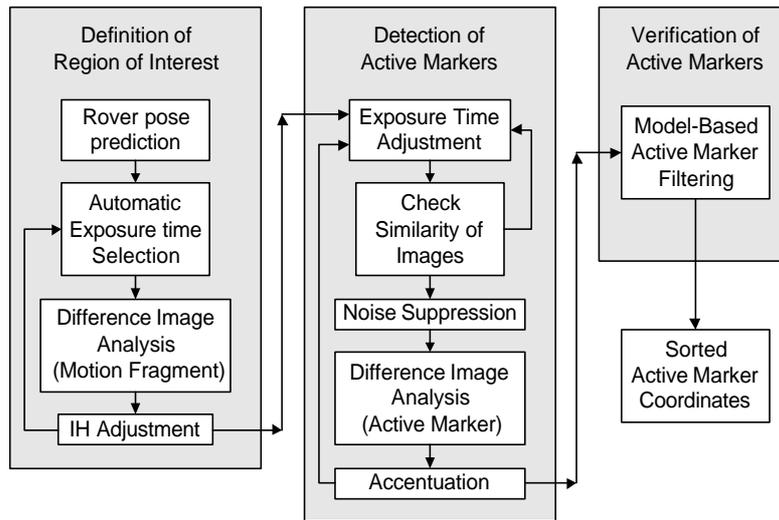


Fig. 16 Control flow of AM detection

The drawback of the AM concept is its accident-sensitivity and the large number of necessary images, taken one by one at different points in time: one reference image and four individual AM images, i.e. at least five images for each camera. For the whole sequence a comparable image characteristic has to be guaranteed.

As shown in Fig. 16 the AM detection is therefore subdivided into three steps:

1. A 'Region of Interest (**ROI**)' will be defined, which reduces the load of handling huge images. Moreover, the most probable Rover location is centered by the Imaging Head via an electrical follow-up (PTU).
2. Potential AMs will be extracted from the selected ROI. Due to the accident-sensitive DIA the quality and similarity of processed images has to be monitored in order to minimize disturbances and accentuate possible AM candidates.
3. The set of potential AMs is subjected to a plausibility check considering general model-based interdependencies.

A robust method to extract interesting regions of an image bases on the 'Motion Fragment Analysis (**MFA**)'. The technique behind is equivalent to the AM concept, but in this case the Rover itself is considered to be a big AM. A DIA of two images with different Rover positions bears noticeable traces of its motion. In such a situation morphological filtering as well as standards in the field of 'Blob Analysis (**BA**)' are optimal tools to define the region of highest probability where the Rover could be. The computational load is quite low and allows an on-line follow-up (e.g. by the PTU). Even big disturbances like unexpected vibrations of the Imaging Head while taking an image does not corrupt the result completely as long as the motion fragment density distribution masses close to the real Rover location.

Differing from the MFA, which is applicable only as a tendency measure, the AM concept has to bring out at best absolute image coordinates of each AM within sub-pixel resolution. Furthermore, the acquirable detection accuracy must not be influenced by the current constellation of the Imaging Head and the Rover. For that reason the basic form of an AM has to be both compact and symmetric. Due to this geometrical constraint pixel-based filter operations (e.g. morphological if applicable) does not distort the interesting focal point of an AM in a wide range.

6.5.2 Limits of the AM localization method

Several additive effects have a negative influence on the visibility of an AM. The order of magnitude is direct proportional to the distance:

- A small inclination angle of the Imaging Head, which is a typical indicator of a huge AM distance, complicates the visibility as well as separability of PLC mounted AMs.
- CCD blurring caused by local overexposure of pixels adulterates the results of the DIA (section 6.5.1).

- Changes concerning the illumination soften the contrast or produce reflections within a sequence of images. For both the discriminatory power of the DIA result degrades.
- The signal-to-noise ratio (SNR) per DIA goes rapidly down and reaches in a distance of approx. 7.5 meters the value 1. I.e., pixel noise produces equivalent results as a real AM. At least at that point the usage of any kind of morphological filtering has completely become meaningless.
- Pixel quantization in combination with a small inclination angle causes big uncertainties during the ongoing localization process especially in case of the depth resolution.

These limits can partly be managed by a work-around, typically a reconfiguration, which results in a better system representation (see section 6.3: Strategy to improve AM visibility via a change in current Rover configuration), or by using model knowledge of the system during data evaluation. However, that is not possible for all of them.

Distance CCD to AM [m]	1.6	7.5	10.0	20.0
Size of AM [pixel]	25.0	1.0	0.6	0.15

Table 2 AM size depending on the distance to CCD

Table 2 indicates the most dominant problem inherent to the AM concept, which cannot be handled. With increasing distance between AM and camera both detected marker size and intensity decreases in a quadratic order. Finally at the claimed maximum distance - under the constraints of the PSPE project - one AM illuminates at most 15% of a CCD pixel, which makes a reliable AM detection quite unrealistic.

For this extreme situation a change to PF based localization is recommended for more reliable results. The advantage of PF over AM is caused by a more universal approach taking into account non-ambiguous silhouettes of the Rover. But the cost is a noticeable increase of the computational load which has to be managed by the Lander computer. PF has also a maximum localization depth, but is certainly beyond the given project requirement.

6.5.3 Automatic follow-up of the stereoscopic range

A very important point in the autonomous navigation concept is the automatic alignment and follow-up capability of the Imaging w.r.t. the Rover movements. At any time the Rover should be focused close to the middle of the stereoscopic camera range. But a simple evaluation of both mono images is not sufficient for that purpose.

Distance [m]	1.6	10.0	20.0
Mono camera range [m]	0.65	4.10	8.20
Stereoscopic coverage [m]	0.44	2.78	5.06

Table 3 Stereoscopic coverage.

In addition to the consideration of a pure stereoscopic aspect, an error model of the effective Imaging Head set-up has to be taken into account. The parameters of this model were typically identified during vision-based terrain exploration at the very beginning. Both predicted and measured Rover locations can be used to align the Imaging Head precisely.

6.6 Vision-Based Pose Estimation

The main problem doing a 'Model-based 3D Pose Estimation (MPE)' is to adjust the position and orientation of an appropriate object model, defined in an 'Object Coordinate System (OCS)', with measured features of a real object in a 'Sensor Coordinate System (SCS)'. Typically a wide-range of equivalent object models can be designed for the same object as long as the selected set of features supports a unique alignment. But the design process is certainly limited to features, which will be optimal in combination with the available measuring equipment. In principle, a model-based approach is a reliable and fail-safe method, especially in case of redundant or ambiguous information. And, it is easy to add different qualities of information to the model in the sense of sensor fusion.

Due to the nature of the problem, our pose estimation approach consists of two steps:

1. We apply an analytical linear approach, which delivers a first guess of the Rover's pose in a numerically efficient way, but not optimal in terms of any error criterion.
2. We use a more accurate nonlinear approach starting with the initial guess of step 1 for an iterative improvement in the sense of least-squares estimation.

The same image acquisition model will be applied to both steps.

6.6.1 Image acquisition model

The Rover, whose pose should be located, is represented by an object model $X_{\langle OCS \rangle}$ which consists of n markers (AMs) expressed in homogeneous coordinates w.r.t. a fix object coordinate system.

$$X_{<OCS>} \triangleq \begin{bmatrix} x_{11} & \dots & x_{1n} \\ x_{21} & \dots & x_{2n} \\ x_{31} & \dots & x_{3n} \\ 1 & \dots & 1 \end{bmatrix}$$

In order to simplify the mathematical representation, the origin of the object's coordinate system is put into the plane, spanned by the markers. And we assume the z-axis of the object's coordinate system perpendicular to that plane. In this special case the 3rd line of $X_{<OCS>}$ becomes to zero, i.e. $x_{3i} = 0, i = 1 \dots n$.

The transformation from the object coordinate system (OCS) to the sensor coordinate system (SCS) is modeled by the 3x4 homogenous matrix $A \triangleq [R, t]$ with the rotational matrix R and the translational vector t . This matrix represents the unknown object pose w.r.t. the SCS.

$$\begin{bmatrix} y_{11} & \dots & y_{1n} \\ y_{21} & \dots & y_{2n} \\ y_{31} & \dots & y_{3n} \end{bmatrix} = \begin{bmatrix} r_{11} & r_{12} & r_{13} & t_1 \\ r_{21} & r_{22} & r_{23} & t_2 \\ r_{31} & r_{32} & r_{33} & t_3 \end{bmatrix} * \begin{bmatrix} x_{11} & \dots & x_{1n} \\ x_{21} & \dots & x_{2n} \\ x_{31} & \dots & x_{3n} \\ 1 & \dots & 1 \end{bmatrix} \quad (1)$$

$$Y_{<SCS>} = A * X_{<OCS>}$$

And finally, the projection of $Y_{<SCS>}$ from the 3D sensor coordinates, expressed in the SCS, into 2D 'Image Plane Coordinates (IPC)' of the left camera ($U_L \in \mathfrak{R}^2$) and the right camera ($U_R \in \mathfrak{R}^2$) is described by the nonlinear equations:

$$\begin{aligned} U_{<IPC>} &= f_L(L_c * Y_{<SCS>}) \\ U_{<IPC>} &= f_R(R_c * Y_{<SCS>}) \end{aligned} \quad (2)$$

$L_c, R_c \in \mathfrak{R}^{3 \times 4}$ are known error matrices, which represent the 'exterior orientations' of the left (L) and right (R) camera w.r.t. the SCS. $f_L(\dots)$ and $f_R(\dots)$ represent the nonlinear 'pin-hole camera' models with the known 'interior orientations', including the lens distortions.

6.6.2 Analytical initial guess

The goal is to find a unique approximation of the unknown object pose A (1). In order to solve that linear equation $Y_{<SCS>}$ will be interpreted as an observation. This observation is computed from measurements of the real markers localized by the image processing module and expressed in the IPC. The applied computation is basically the inversion of (2). The inverse projection maps the stereoscopic IPC into their corresponding 3D representation in the SCS. The result can be considered as an estimate of $Y_{<SCS>}$, which will be used to derive an estimate for A , within one step.

$$A = YX^T (XX^T)^{-1} \quad \text{with } \text{rank}(X) = 4 \quad (3)$$

But in case of coplanar markers, as obtained by the selected marker layout applied to the NR, $\text{rank}(X) = 3$. In this situation the special selection of the object coordinate system, as mentioned before, allows a cut of the 3rd row of $X_{<OCS>}$ as well as the 3rd column of A without loss of information. This elimination results in a problem of reduced dimension, expressed by

$$Y_{<SCS>} = \tilde{A} * \tilde{X}_{<OCS>} \quad (4)$$

$$\text{with } \tilde{A} = \begin{bmatrix} r_{11} & r_{12} & t_1 \\ r_{21} & r_{22} & t_2 \\ r_{31} & r_{32} & t_3 \end{bmatrix} \quad \text{and} \quad \tilde{X}_{<OCS>} = \begin{bmatrix} x_{11} & \dots & x_{1n} \\ x_{21} & \dots & x_{2n} \\ 1 & \dots & 1 \end{bmatrix}$$

which yields the reduced estimate of \tilde{A} by

$$\tilde{A} = Y\tilde{X}^T (\tilde{X}\tilde{X}^T)^{-1} \triangleq [\tilde{R}, t] \quad (5)$$

This solution delivers the first two columns of R and the translational vector t . The 3rd column of R can be easily calculated by the vector product of the first two columns of R .

This analytical method has the advantage that \tilde{A} can be determined very efficiently in terms of computational costs. But it does not respect the following constraints:

- *Feature dependence*: in reality the markers of the chosen object model are subjected to a rigid body constraint. Therefore the markers are not independent among each other. But applying the inverse projection, the corresponding markers are treated as independent.
- *Orthonormality*: the rotational matrix R has to be orthonormal. But the result of (5) does not necessarily guarantee its orthonormality. Therefore the object pose $A = [R, t]$, determined by the analytical approach, will normally not result in an accurate solution for our pose estimation problem.

But this analytical approach is sufficient to get an initial guess for the following iterative improvement.

6.6.3 Iterative improvement

We apply again the object model, as described in 6.6.1. Considering the above mentioned constraints the interdependency of the markers (in 3D-space) will be implicitly taken into account by doing the iterative improvement within IPC, simultaneously. This results in a measurement vector

$$U_{L_c, R_c < IPC >} \hat{=} [u_{1L_c} v_{1L_c} \cdots u_{nL_c} v_{nL_c}, u_{1R_c} v_{1R_c} \cdots u_{nR_c} v_{nR_c}]^T \quad (6)$$

and satisfies the constraint of rigid object geometry.

Additionally, the unknown object pose will not be described by a homogenous matrix of 3x4 independent elements anymore, but by the Euler vector

$$p_{< SCS >} \hat{=} (\mathbf{a} \ \mathbf{b} \ \mathbf{g} \ t_x \ t_y \ t_z)^T. \quad (7)$$

Within this definition the $\mathbf{a} \ \mathbf{b} \ \mathbf{g}$ angles represents the rotation as well as $t_x \ t_y \ t_z$ the translation of the object pose w.r.t. the SCS. Additionally, the angles $\mathbf{a} \ \mathbf{b} \ \mathbf{g}$ describe the rotation matrix R , which fulfils the orthonormality constraint. For example:

$$R_{Euler} \hat{=} R_z(\mathbf{a}) * R_x(\mathbf{b}) * R_z(\mathbf{g}). \quad (8)$$

Due to the usage of (8) and the focus on IPC, the estimation problem (1) changes to a nonlinear one. So (1) and (2) will be rewritten in a more compact form, considering definition (6) and (7):

$$U_{L_c, R_c < IPC >} = f(p_{< SCS >}) \quad (9)$$

whereas $f(\cdots)$ is a known nonlinear projection function from SCS to IPC comprising (1) and (2). For further computation the necessarily linearization of (9) results in an overdetermined system of equations

$$\Delta U_{L_c, R_c < IPC >_k} = J_k^{4nx6} * \Delta p_{< SCS >_k}, k = 1, 2, 3, \dots, \quad (10)$$

where J_k^{4nx6} is the Jacobian matrix at the initial pose estimate. $\Delta U_{L_c, R_c < IPC >_k}$, the so-called residuum, is the difference between the current measurement vector and the simulated one. $\Delta p_{< SCS >_k}$ is the difference between the improved pose estimation and the former one.

Finally, (10) will be solved in the sense of least squares using the pseudoinverse of J_k^{4nx6} :

$$p_{< SCS >_{k+1}} = p_{< SCS >_k} + (J_k^T J_k)^{-1} J_k^T * \Delta U_{L_c, R_c < IPC >_k}, k = 1, 2, 3, \dots \quad (11)$$

In order to obtain the Jacobian matrix (10) - necessary for equation (11) -, the differentials of J^{4nx6} are approximated by the difference quotients, double sided for numerical convenience.

However, before the iterative improvement (11) of the pose estimation can be started, the initial guess \tilde{A} of section 6.6.2 has to be converted into an equivalent pose vector $p_{< SCS >}$ (7). Due to the missing orthonormality of \tilde{A} , an explicit matrix orthonormalization has to be done, which improves the quality of the initial pose vector. The iteration process will be repeated, while the pose corrections $|\Delta p_{< SCS >_{k+1}} - \Delta p_{< SCS >_k}|$ does not fall below a given limit.

The experimental results, gained by the iterative improvement approach, clearly indicate a reduction of the originally remaining pose error. Another possibility, bringing substantial improvement of the pose accuracy, is to strengthen the pose estimation by direct measuring of some of the pose vector elements (e.g. \mathbf{a}, \mathbf{b} if the Rover disposes of inclinometers).

7 Conclusion

An End-to-end Control System has been proposed to allow for highly autonomous operations of a Rover (the Nanokhod) on a planetary surface. The On-Ground Control System allows scientists to select sites for exploration by

the rover in a 3-D model of the terrain surrounding the Lander. The On-Ground Control System then automatically determines the optimal rover path to visit these sites in a safe and efficient manner, and uplinks this information in the form of high-level commands.

The On-Lander Control System is responsible for vision-based supervision and control of the Rover's movements at the remote place in space. The navigation concept is arranged for autonomous execution. The On-Lander Control System subsequently executes the high-level motion commands in an autonomous way and ensures that the Rover reaches its destination as planned. Originally developed for space missions, the system can be adapted to comparable applications on Earth, e.g. for controlling mobile robots in harsh and dangerous environments.

8 Acknowledgements

The described system has been developed under ESA contract no. 13501/99/NL/PA, Payload Support for Planetary Exploration (PSPE), also known as ROBUST (Rover Based Utilisation of Science and Technology). We would like to thank G. Visentin from ESA/ESTEC, Noordwijk, R. Jaumann, T. Behnke, H. Grothues, H. Michaelis and H. Hirsch from the DLR, Institute of Space Sensor Technology and Planetary Exploration, Berlin, as well as R. Dequeker and P. Degezelle from OptiDrive, Belgium, and R. Bertrand, H. von Hoerner and L. Ziegler from von Hoerner & Sulger GmbH, Germany, as well as Dr. R. Rieder for their collaborations to this project.

References

- [1] <http://science.ksc.nasa.gov/mars/Rover/mission.html>
- [2] R. Volpe, T. Estlin, S. Laubach, C. Olson, J. Balam, *Enhanced Mars Rover Navigation Techniques*, in Proc. of the IEEE Int. Conf. on Robotics and Automation, San Francisco, CA, April 2000
- [3] J. Bresina, G. A. Dorais, K. Golden, D. E. Smith, R. Washington, *Autonomous Rovers for Human Exploration of Mars*, in Proc. of the 1st Annual Mars Society Conference, Boulder, CO, August 1998
- [4] R.G. Bonitz, T.T. Nguyen, W.S. Kim, *The Mars Surveyor '01 Rover and Robotic Arm*, in Proc. of the IEEE Aerospace Conference, paper no. 2.0301, March 2000
- [5] B. Fontaine, D. Termont, L. Steinicke, M. Pollefeys, M. Vergauwen; R. Moreas, F. Xu, K. Landzettel, B.M. Steinmetz, B. Brunner, H. Michaelis, T. Behnke, R. Dequeker, P. Degezelle, R. Bertrand, G. Visentin, *Autonomous Operations of a Micro-Rover for Geo-Science on Mars*, 6th ESA Workshop on Advanced Space Technologies for Robotics and Automation, ASTRA 2000, ESTEC, Noordwijk, The Netherlands, December 5-7, 2000
- [6] R. Bertrand, J. Brueckner, M. van Winnendael, M. Novara, *NANOKHOD – A Micro-Rover to explore the Surface of Mercury*, 6th International Symposium on Artificial Intelligence, Robotics and Automation in Space, Montreal, Canada, June 18-22, 2001
- [7] L. Steinicke, B. Fontaine and G. Visentin, *End-to-End Control System for the Nanokhod Planetary Micro-Rover*, i-SAIRAS 6th International Symposium on Artificial Intelligence, Robotics and Automation in Space, Montreal, Canada, June 18-22, 2000
- [8] M. Vergauwen, M. Pollefeys, L. van Gool, R. Moreas, F. Xu, H. van Brussel and G. Visentin, *Calibration, Terrain Reconstruction and Path Planning for a Planetary Micro-Rover*, i-SAIRAS 6th International Symposium on Artificial Intelligence, Robotics and Automation in Space, Montreal, Canada, June 18-22, 2001
- [9] S. Laubach, J.W. Burdick, *An Autonomous Sensor-Based Path-Planner for Planetary MicroRovers*, IEEE Int. Conf. on Robotics and Automation, Detroit, MI, May, 1999
- [10] R. Y. Tsai, *A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf tv cameras and lenses*, IEEE Journal of Robotics and Automation, 3(4): pp. 324-344, 1987
- [11] C. Harris and M. Stephens, *A combined corner and edge detector*, In Fourth Alvey Vision Conference, pp. 147-151, 1988
- [12] M. Fischler, 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.
- [13] S. Maybank, *Theory of reconstruction from image motion*, Springer Verlag, 1992
- [14] M. Pollefeys, R. Koch and L. Van Gool, *A simple and efficient rectification method for general motion*, Proc International Conference on Computer Vision, 99, pp.496-501, Corfu (Greece), 1999
- [15] L. Falkenhagen, *Hierarchical Block-Based Disparity Estimation Considering Neighbourhood Constraints*, Proc. International Workshop on SNHC and 3D Imaging, pp. 115-122, Rhodes, Greece, 1997
- [16] B.-M. Steinmetz, K. Arbter, B. Brunner, K. Landzettel, *Autonomous Vision-Based Navigation of the Nanokhod Rover*, i-SAIRAS 6th International Symposium on Artificial Intelligence, Robotics and Automation in Space, Montreal, Canada, June 18-22, 2001
- [17] K. Arbter, G. Hirzinger, J. Langwald, G.-Q. Wie, P. Wunsch, *Proven Techniques for Robust Visual Servo Control*, in: 'Robust Vision for Vision-Based Control of Motion', ed. M. Vincze and G. D. Hager, chapter 9, pp. 109-125, IEEE Press, New York, 2000