SEEKER-AUTONOMOUS LONG RANGE ROVER NAVIGATION FOR REMOTE EXPLORATION

Mark Woods (1), Andrew Shaw (1), Estelle Tidey (2), Bach Van Pham (3), Unal Artan (4), Brian Maddison (5) Gary Cross (6),

(1) SCISYS, Clothier Rd., Bristol, BS4 5SS, UK, mark.woods@scisys.co.uk:
(1) SCISYS, Clothier Rd., Bristol, BS4 5SS, UK, andy.shaw@scisys.co.uk:
(2) Roke Manor Research, Romsey, SO51 7SA, UK, estelle.tidey@roke.co.uk
(3) CNRS, LAAS, 7 Avenue du colonel Roche, Toulouse, FR
(3) Univ de Toulouse, LAAS, Toulouse, FR, bvpham@laas.fr:
(4) MDA Space and Robotics Limited, Didcot, Oxfordshire, OX11 0QX, UK, unal.artan@mdacorporation.com:
(5) Rutherford Appleton Laboratory, Didcot, Oxfordshire, OX11 OQX, brian.maddison@stfc.co.uk:
(6)) BAE Systems, Advanced Technology Centre, Filton, Bristol, BS34 7QW, UK, gary.cross4@baesystems.com:

ABSTRACT

Under the umbrella of the innovative, ESA StarTiger program, a rapid prototyping study called *Seeker* was initiated which brought together a diverse range of partners from space/non-space to develop a prototype Mars rover system capable of autonomously exploring 6km of highly representative Mars terrain over a 3 day period. This paper reports on our approach and the final, successful field trials which took place in the Atacama Desert. Long range navigation and the associated remote, rover field trials are a new departure for ESA and this study therefore represents a novel initiative in this area. The study's primary focus was to determine if current computer vision and artificial intelligence based software could enable such a capability on Mars given the current limit of around 200m per sol. This work is part of a wider effort to achieve a step change in autonomous capability for future Mars/Lunar exploration rover platforms.

1. Introduction

Currently, Mars rovers are able to traverse in the order of 100m to 200m per sol. This acts as a substantial constraint on remote exploration activities. There are a number of reasons why this is the case such as low power, slow processor speeds and limitations in the required vision software. Given recent advances in terrestrial, state of the art computer vision systems it is believed that the software aspects of the problem could be improved and shown to be functionally capable of navigating over longer distances on Martian terrain. The study did not seek to develop novel power or mechanical solutions in addressing the problem as these aspects are being addressed through other initiatives.

Following the ESA StarTiger methodology, a team of diverse experts were co-located for a period of 6 months at the RAL facility in Oxfordshire in the UK to focus on this difficult technical challenge. Each team member brought existing expertise and software/hardware inputs in areas such as visual odometry, robotics platforms, 3D mapping, short and long range path planning, science autonomy and mission planning. This provided a solid foundation with which to develop a prototype system and work towards the stated goal.

The primary objective for the study was to demonstrate that the system could travel in the order of 2km per day over a three day period i.e. a 6km traverse in total. From system perspective the emphasis for the study was therefore on the autonomous navigation, mapping and localisation system. Given the distances travelled however the study also addressed science autonomy as a lower priority objective. Prior to Seeker trials in the ESA context have been limited to European locations such as El Teide in Tenerife, or dedicated facilities such as the outdoor "Mars yard" at CNES. However it was felt that to truly demonstrate the proposed concept and advance the state of the art it was necessary to carry out tests in the best Martian analogue site available on Earth. Complementary work by other groups such as [1] have demonstrated the many benefits of testing in such environments. In short they provide the diversity of terrain and lighting conditions which a rover can be expected to encounter but are difficult or impossible to emulate in small or synthetic test environments. The Atacama Desert in Chile was selected as the final trial site. Section 3 discusses this further.

2. System Overview

Figure 1 below shows the high-level system concept, consisting of a set of key components which included: OVERSEER Component orientated autonomy framework, Image Management, Visual Odomtery (VO), VSLAM, DEM Generation, Absolute Localisation, Data Fusion, Path Planning, platform Control and Feedback, Science Autonomy, Localisation Management, Timeline Execution Control and High-Level Decision making and off-board operations planning, monitoring and control.

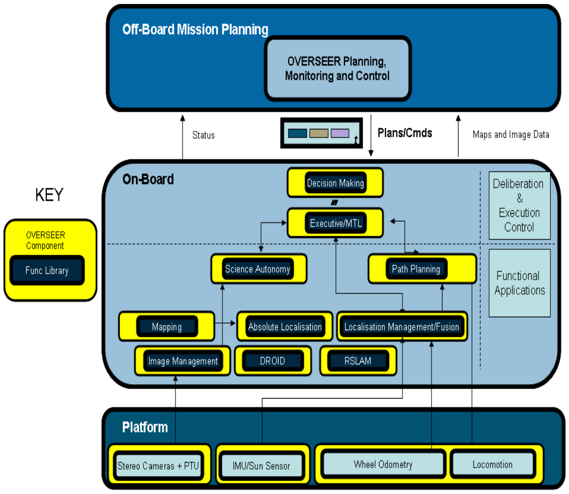

Figure 1: System Overview

2.1. OVERSEER On-Board Framework

The Seeker concept required a means to "drop-in" existing stand-alone autonomy capability for various functions such as mapping localisation and planning/execution and also a solution to the non-trivial problem of integrating position dependent functions which used a variety of reference frames. SCISYS have developed an autonomy framework called OVERSEER [2] see Figure 1 which was used to address this problem. Individual functions such as VO and DEM generation were supplied as libraries and then integrated in a dedicated smart component which took care of co-ordinate transforms, data relevance checks, execution control and component to component communication. This generated a significant amount of system work but it ensured a consistent and robust final system. It also allowed individual function providers to focus on discrete functional requirements whilst off-loading other integration issues to the smart framework development.

2.2. Imaging and Control

Visual Odometry (VO) was the primary localisation means for Seeker. The images were also used by the DEM generation component to produce the 3D terrain models that are subsequently used by navigation for path generation and obstacle avoidance. As a consequence image capture and distribution provided the core event driven "heartbeat" of the system.

The SCISYS camera interface library allows the connection of multiple camera types from USB, Firewire and Ethernet and provides a generic interface to them for image capture, calibration, rectification and saving etc. In this context of Seeker the library was used to obtain images from various colour Point Grey Bumblebee 2 and XB3 cameras, rectify them to remove lens distortion, reduce to half scale and convert to black and white. As multiple components required the image data these were saved in the data storage area and points to which were distributed to the components as a rate of 5Hz.

2.3. Visual Odometry, Localisation Management and Fusion

In the Seeker system, visual odometry was provided by the Roke Manor DROID structure-from-motion algorithm and the Oxford RSLAM [3] algorithm. DROID operates by detecting and tracking visual features in imagery from a geometrically calibrated camera. By analysing the apparent motion of the features (which are assumed to be static), DROID simultaneously determines both the 3D locations of the features and the camera motion, consisting of its change in pose (position and orientation) between successive frames. [4].

The DROID algorithm performs well in Mars-like terrain due to the feature-rich texture of the ground. This is illustrated by Figure 2, which shows plentiful features detected by DROID superimposed on an image captured during the field trials in Chile. The 3D structure of the terrain is shown by colour-coding of the features by range with blue being the most distant.

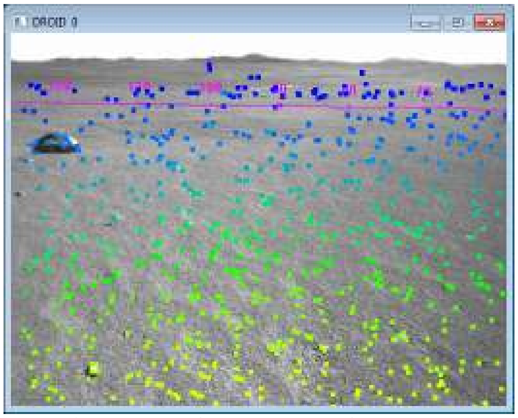

Figure 2: 3D features output by DROID

Important considerations in the use of image processing for odometry include:

Number of cameras: While algorithms such as DROID can operate on imagery from a single camera, it is more straightforward to resolve the speed-scale ambiguity (in which the speed of the camera motion and the scale of the feature point cloud are in error by the same factor) with stereo imagery, as the camera baseline can be used as a yardstick to determine scale.

Camera resolution, field of view and positioning: These affect whether or not it is possible to view and detect features. Tilting the cameras downwards ensured that the field-of-view contained a large feature-rich area (the ground) and that the cameras were less likely to point directly towards the sun, which can cause image saturation. Vehicle speed and frame rate: Visual odometry by feature tracking depends on being able to track features over multiple frames, so the movement of the camera between successive frames must not be greater than the its field of view. A frame rate of 5Hz was selected as being fast enough to achieve good results at the required vehicle speed.

Camera calibration: Visual odometry algorithms output the motion of the camera, not the vehicle, so it is important to accurately measure the pose of the camera relative to the vehicle. A calibration process was devised in which the PTU was moved in a set path so that the camera viewed a target fixed to a known position on the side of the vehicle. The 3D position of the target features relative to the camera was calculated by using DROID to measure the positions of the target features, and the camera pose could then be deduced Figure 3 shows an image captured during calibration.

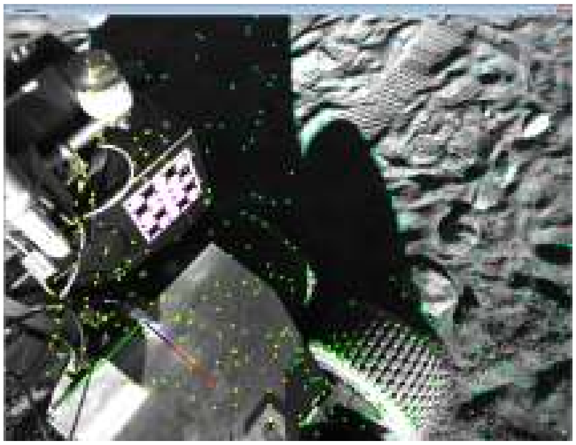

Figure 3: Camera calibration tool display, showing the vehicle origin and axes and the position of the target.

The localisation of the platform in the field was performed by Localisation Manager component which received inputs from a variety of pose estimation sensors. Each of these poses estimation sensor provides the estimation for the change in the pose between camera frame trigger events which are then integrated over time to provide the relative pose from the initial starting location.

Along with the 6DOF delta pose estimates obtained from the VO systems the localisation component obtained estimates from the platform wheel odometry (x, y, yaw) and the IMU (roll, pitch, yaw). A Kalman filter was used to fuse the pose information provided by each of the sensors, where each input could be weighted depending on the axis accuracy of the sensor. The weighting values used for the sensors can also be adjusted depending on the terrain characteristics, ie during the trails a lot of areas the sand had a hard crust which when broken a soft sand was underneath which then introduced noise in the wheel odometry. Using this dynamic approach allows the system to adapt to changing environmental conditions and also system constraints to provide the best vehicle pose estimate within the environment.

2.4. Mapping and Absolute Localisation

In Seeker, only a single Bumblebee stereo camera was used to perceive the environment around the rover. The Digital Elevation Model (DEM) around the robot is captured with a stereo algorithm which provides dense 3D points. As the quality of the 3D points degrades with the distance from the rover, two thresholds are used. The 3D points whose distance to the rover is less than 7 meters are kept to create a fine DEM (5 cm in resolution) which is used for obstacle detection and short path planning. A distance of 15 meters is used to create coarse DEM (1 meter in resolution) which is used to estimate the global position of the map. In order to obtain a large DEM, different single DEMs with the same resolution (e.g. 5 cm or 1 m) are stitched to each other.

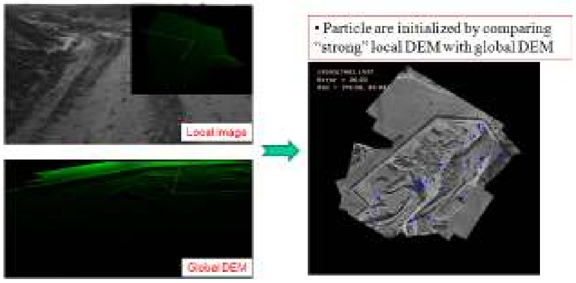

Figure 4: Example of a left stereo image, its local DEM, the global DEM and the initial particle position inside the global map.

To estimate the global position of the rover, a discretized particle filter algorithm is proposed. Instead of throwing randomly the particle inside the global map, the particles are initialized by comparing the stitched local DEM t around the rover with a small patch f of the global DEM using ZNCC:

$$s_{ZNCC} = \frac{1}{n-1} \sum_{x,y} \frac{(f(x,y) - \bar{f})(t(x,y) - \bar{t})}{\sigma_f \sigma_t}$$

Where n *is* the number of point in the local DEM, $t(x,y)$ denotes the elevation of the 3D point at position (x,y), $\bar{t}$ and σ_t are the mean and standard deviation of the local DEM. Figure 4 shows example of a stereo image, its local DEM, the global DEM and the initial position of the particles. Depending on the initial knowledge of the rover's global position and orientation, the search zone could be only a small portion of the global map or the entire global map.

After having initialized the particles, the estimation of the rover global position is improved as the rover moves. Instead of employing a continuous particle distribution method, a discretized distribution method is proposed: only one particle with a given orientation α is allowed in a single DEM cell. But in a DEM cell, there could be multiple particles with different orientations. In our proposed system, the orientation is also discretized with a step of 3 degrees. Alternatively speaking, a particle with a position (i, j) and orientation γ, represented as (i, j, γ), is unique in the proposed system, where (i, j) is the row and column index of the particle in the global map.

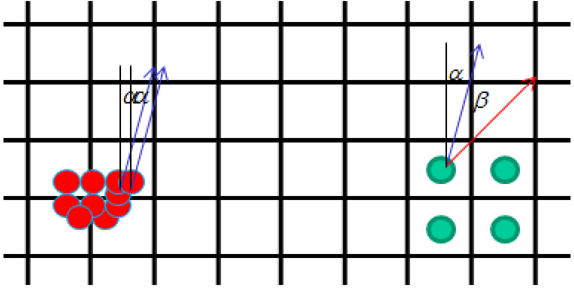

Figure 5: Conventional continuous particle distribution (left) with particles with similar orientation and position in one single cell and the proposed discretized particle distribution (right).

In addition to the discretized distribution method, a Gaussian model is employed to represent the uncertainty of one particle position. Once, the uncertainty of a particle gets bigger than the resolution of the global DEM (i.e. 1 meter), that particle is decomposed into smaller particles with the same orientation but with different position. By then, the uncertainty of each particle is reduced to half of a global DEM cell. Figure 6 illustrates a situation where a particle is decomposed into 3 particles.

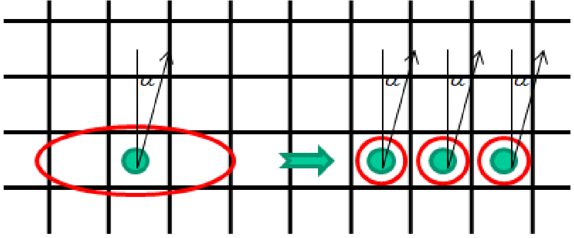

Figure 6: Decomposition of particle uncertainty.

Given a number of particles distributed in the global map, the absolute position of the rover is calculated as the mean value of the position of these particles. Their variance value is used as the uncertainty of the estimation. In order to keep the computation low, the particles are only updated after a large distance and only when the local DEM has strong signature.

2.5. Path Planning

The MDA path planning task is subdivided into two stages, the first being a global or long range path planning and then a second local or short range path planning. The global path planning relies on the scientist to provide long range way points for the rover to traverse while avoiding visible obstacles for a safe path. Figure 7 illustrates the planned path for the 6 km traverse.

Figure 7: Top view of the field trial environment in OVERSEER Operations tool and the selected long range way points for the final 6 km traverse

The global path is then built by connecting the selected long range way points. However, due to global map resolution and limitations due to limited sensor field of views, the local path planning component is required. Using the, coarsely spaced, long range way points selected by the scientist, a finer path is constructed with way point spacing of 2 m. The new long range way points are then used as short range or local goals for the rover to traverse to.

A 2 m value (determined empirically) allows the system to operate within the local 3D point cloud (Section 2.4) and is small enough to force the system to update often to detect and avoid obstacles not distinguishable from the global map image. This is due to no real-time obstacle avoidance in the system (larger computational demand). Over the 2m, the environment was taken to be static such that once a safe path is found the rover could then track the generated short range path to reach the short range goal.

To generate a safe, obstacle free path, the rover must first classify the terrain (terrain assessment) into traversable and non-traversable cells. Cells were constructed by gridding the generated 3D point cloud to a cell size of 15 cm. The terrain assessment [5] took into account; the ground clearance of the rover the maximum slope the rover is capable of traversing and also the roughness of the terrain. Using these three criteria a cost map of traversable and non-traversable cells is able to be constructed. Figure 8 provides one of the generated cost maps overlaid with the long range path and the corresponding short range path during the final 6 km traverse.

A higher level of autonomy was included into the short range path planner, where if the requested short range goal is within a non-traversable cell, an obstacle not seen on the global map image, the system would then efficiently search for a new goal location. The new goal location with the smallest distance to the next goal is then chosen as the best available and the corresponding short range path used for the rover to track. If no path could be found, the component would then raise a flag indicating that no path could be found. This required further action by higher level components and while waiting an action zero valued rover commands are sent to the rover. Once an action is taken such that a short range path could be found then the rover would move. Possible higher level actions is to rather pan the camera to increase the coverage of the local 3D point cloud which could allow the system to find a possible path to the goal or to send the rover a new long range path to follow which would deviate the rover around the current obstacles.

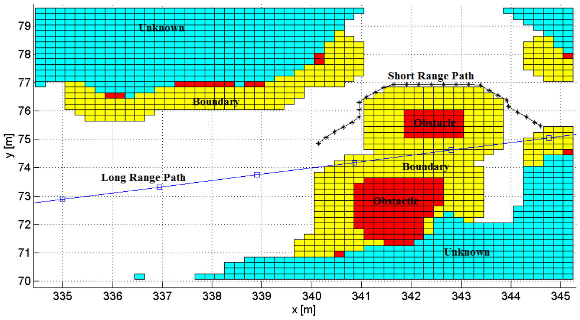

Figure 8: One of the many local cost maps generated during the 6 km traverse. The long range and computed short range paths are overlaid onto the cost map.

The rover stayed on the path by using the pose estimate of the rover (Section 2.3) and the current computed short range path to generate the necessary rover commands (v, w) to keep it on course were then sent to the rover. The commands generated are tuned to the specific rover platform to keep the system from deviating from the generated path, thus reducing the chance the rover strays into an obstacle. Once the rover reaches the short range goal the procedure is triggered again until the rover reaches the long range goal or the executive requests a rover stop or a new long range path is sent to the rover. The short range path planner can be broken down into 5 distinct steps;

Step 1: Load in the current local 3D point cloud
Step 2: Execute terrain assessment
Step 3: Compute a path using D* search algorithm
Step 4: Generate rover commands (v, w)
Step 5: Repeat Step 4 until the short range goal is reached and then repeat Steps 1-3. Repeat Steps 1-5 until the long range goal is reached

During the field campaign, a minimal pause was occasionally observed between short range path updates, providing a more continuous system.

2.6. On-Board High-Level Decision Making and OVSERSEER Off-Board Control

Long-range navigation is clearly impossible without some level of resource management and mid-term tactical decision making. In the Seeker case this support was provided by the SCISYS Execution and Decision making components which controlled the execution of the initial timeline, addressed resource issues, adapted the timeline to GNC failures and implementation of initial plans and priorities generated off-line by the operations team.

Off-line planning, monitoring and control for Seeker were handled via the OVERSEER Operations UI application. This allows operators to view the estimated rover position on a 2D/3D representation of the trial site and plan a route by visually adding waypoints. This can be used as an input to the timeline based planning process where the planning interface is used to assemble an outline plan with appropriate tasks such as navigation, stereo panoramics and science selection. One complete the timeline can be dispatched for execution on the rover.

Execution is handled by the on-board executive which choreographs all aspects of timeline progression by delegating tasks to appropriate components such as path planning and localisation management. Proposed task execution time is onitored within an appropriate window and tasks are dispatched accordingly. The Executive plays a central role in determining the system response in the event of sub-system “exhaustion” e.g. no path to goal. In this case unplanned stereo panoramas can be inserted in the timeline following approval by the on-board decision making system to help the system recover from short-term faults.

Although the issue of computational resources was broadly out of scope for Seeker the system did consider the impact of localisation means versus resources. Although it is anticipated that over the longer term future missions will utilise FPGA technology in order to overcome the computational bottleneck, it seemed prudent to provide some means of intelligently managing the use of computationally expensive algorithms over the course of a route. To this end a smart adaptation system was developed which considered current terrain and localisation performance in order to determine which combination of localisation modalities (e.g. Visual or Wheel Odometry) where required to ensure good progress whilst ensuring an adequate and reliable pose estimate. The component uses estimates of the current terrain type, sinkage and sensor divergence to determine what frequency combination of VO and WO are required for the current route segment.

Although Seeker simulates a real mission in terms of plan upload and end of day telemetry download, the availability of long-range Wi-Fi in the field allowed the team to use the remote monitoring capabilities in OVERSEER to check progress of the executing mission plan by overlaying reported pose, live DGPS ground truth, the planned route and absolute localisation matches.

Figure 9: OVERSEER monitoring view showing reported tracks.

2.7. RoboVolc Platform and Control

The ambition for Seeker was to demonstrate innovative technology in a comparable Martian terrain, consisting of vast boulder fields, soft soil and steep undulations. In order to facilitate a real-world demonstration, BAE Systems' RoboVolc was provided by a team at the Advanced Technology Centre (ATC).

Originally designed for volcanic exploration as part of a collaborative European project of the same name RoboVolc, offered proven all-terrain capabilities with its six wheeled skid steer drive train, in addition to fully articulating front and rear axles designed with 3 degrees of freedom through roll, pitch and linear travel. The drive train consists of six, high torque gear-head motors and custom tyres with the capability of traversing the expected terrain types, maximising mobility. The system control software provides the versatility to operate in the range between 0.1 - 1m/s, catering for the ability to autonomously navigate at speed as well as performing delicate manoeuvres through narrow paths. The option of using an external generator further provides the capacity for continuous operation for seven hours.

Figure 10: RoboVolc Platform in Seeker Valley

The platforms physical architecture and software systems have been designed to allow the rapid integration and validation of 3rd party hardware payloads and software components.

BAE Systems supplied the platform with a rugged safety remote interlock system, remote access, and a number of low level software components that enable the monitoring of vehicle telemetry, and control of the locomotion of the RoboVolc platform from both the payload and remote operators. For Seeker, the platform was modified in two ways, firstly to carry a payload that included essential computer and navigation hardware, and secondly to incorporate a telescopic mast, that provided the high vantage point required for the Stereo Camera and Pan Tilt Unit. RoboVolc was outfitted with a number of state of the art sensors that included high resolution encoders for Wheel Odometry, Point Grey Bumblebee XB3 for Visual Odometry (VO) and an XSens MTI-G for attitude and heading measurement. Wi-Fi infrastructure was also provided, which enabled the OVERSEER Ground Control Station to maintain point-to-point communications over a 2km range.

3. Field Trials Outline

In order to develop the prototype, a series of progressive trials were conceived starting indoors, moving to an outdoor Mars yard, then a red sand quarry site and then long range-beach testing, all in the UK. This was then followed by final trials in Chile. At the outset of the project several global test areas were considered including, Morocco (our first proposed site,) Death Valley in the US and the Atacama. Our final selection was based on a variety of factors based on a scoring of terrain suitability, logistics support, cost, security, and human factors. The preferred Atacama site scored highly on all parameters and we further benefited from special permission from the European Southern Observatory (ESO) organisation to base the team at the Paranal observatory site near Antofagasta. After a reconnaissance and then full regional visit a final site selection was made based on the parameters outlined in the table below in order to allow a range of test conditions for various subsystem elements in one location. The table was used as a template for test route selection. Selected routes should seek to exhibit variations of these parameters to support the basic concept of progressive testing from easy through to hard on the final test days.

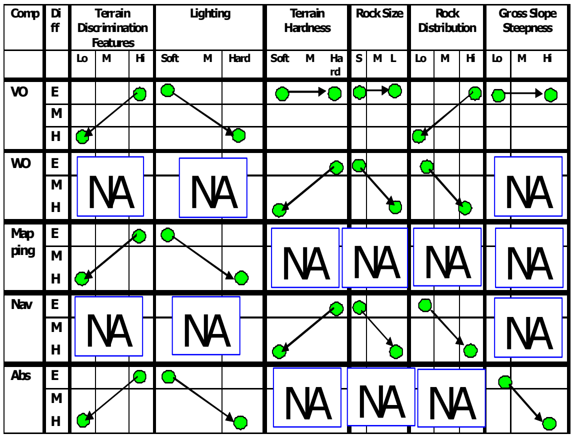

Table 1: Terrain categorisation parameters versus complxity per subsystem

As large areas of the Atacama region are currently only coarsely mapped the site was given the name Seeker Valley. To support absolute localisation and operations planning UAV flights were conducted to provide 2D ortho and DEM's of Seeker valley covering a desert area of approximately 2km x 3km. These were then downgraded to lower resolution (1m/pixel) in the z-axis to replicate the quality of orbital data available on a real mission. Final trial route planning was carried out on a blind basis. Although the team had full access to the valley, the vast majority of the final route (effectively the areas south of the base camp) it was felt that this would both a fully representative test and provide the best learning opportunity for the team which is one of the most significant benefits of such a high-fidelity trial.

4. Results and Conclusions (MW/AS) (0.5-1.0p)

In total, the system covered in the order of 15km over the entire Atacama trials period including debugging and test preparation. As part of the Seeker system a DGPS unit was used to obtain ground truth positional information at 1Hz, timestamp with the other components. This allowed post analysis of the system in terms of positional errors in the location. The objective for the trials was an attempt to traverse 2km per day over three days in order to achieve a 6km cumulative total. Day 1 of the three day evaluation resulted in a km order traverse. This trial moved us into new territory and caused a series of previously undetected issues with the support infrastructure given the distance from our base and operations centre. Day 2 was mostly dedicated to solving these support issues to allow us to perform a more ambitious run (with respect to the original objectives) on Day 3 of 6km.

The final day resulted in a cumulative traverse of 5.05km in the x and y and approximately 7km when z elevation is included. The entire run consisted of three distinct segments made in a loop which took us in quasi-clockwise direction around the Easterly facing base camp position. The first 850m segment started from base camp and sun-up with extremely challenging traverse directly into the sun through highly reflective salt flats which were successfully traversed. The environmental conditions on the final day were much more extreme than previously encountered throughout the trial and resulted in multiple equipment overheat incidents which slowed our progress considerably and a camera interface failure terminated this segment. The second segment consisted of a 1.2km autonomous run through the most challenging terrain the system had encountered – namely extensive boulder fields, considerable slope and extensive soft sand. The system coped admirably with very difficult conditions before being forced to stop because of equipment overheating. The final segment of 3km back towards base camp towards the end of the day in a Northerly direction i.e. directly into the sun and included one operational waypoint failure where the rover was unable to pass through and extensive boulder field. The test was stopped just after the 5km mark with sundown approaching.

During the 5.05km run approximately 60GB data was collected consisting of rectified image data, DEMs, wheel and IMU odometry and absolute locations. In itself this represents a hugely valuable data set resource. The Seeker architecture allows the system to be replayed in real-time with the data that was collected allowing changes in fusion weights, VO configuration, DEM generation parameters etc so that individual components performance can be evaluated.

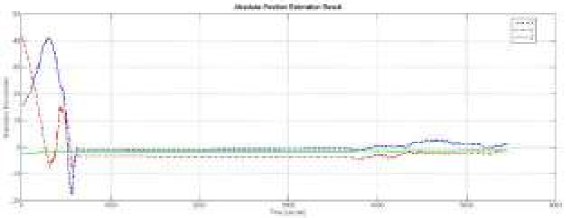

Figure 11: Absolute position estimation result

Figure 11 shows the result of the absolute position estimation algorithm. The horizontal axis shows the timeline and the vertical axis expresses the estimation error in X, Y and Z. The initial search zone has a size of 1x1 km2. As shown in this chart, the correct position of the rover with an error less than 2 meters was found after 9 minutes. It is important to note that the convergence rate does not depend on the time but depends on the distance travelled and more importantly on the elevation variance of the local DEM (e.g. flat area is not good for absolute localization). Besides the estimation in x,y,z the proposed algorithm can also estimate the absolute heading of the rover.

Perhaps the single biggest benefit of the Atacama site was the unparalleled range of terrain lighting conditions with which to test DROID and RSLAM and Localisation Management performance in particular. At the time of writing the results are being assessed/validated by the team and ESA given the relevance to missions such as ExoMars. Initial indications point to an impressive performance overall given the conditions and they also emphasise the need to manage and intelligently fuse all localisation sources to ensure robustness over long range traverses. One of the most challenging aspects of this terrain is the complexity of path planning with limited guidance from and operations team. The sheer complexity of the terrain presents a huge challenges for any autonomous system. Despite this, the system approach and trials demonstrated the benefit of robust, local path planning in this environment.

Throughout the Seeker trials, the RoboVolc platform was both reliable and robust, providing the capability to effortlessly traverse difficult and varied terrain (~60km over the entire project); ensuring further time for software integration and experimentation. The front and rear axle's articulation provided great agility whilst traversing across the undulating terrain, particularly during the dense boulder fields frequently seen in Seeker Valley. In addition, when the fierce heat of the Atacama Desert started to affect the performance of the Seeker payload, RoboVolc and its base safety systems continued to accumulate 2-3 km runs per day over a 3 week period without fault. Given the successful integration and test of the high level localisation and control systems there is now the potential to both increase the capability of the existing system, and also move towards more representative computing, sensor, and rover platforms

Autonomy system engineering in this study was challenging given the ambition and scale of the work however the final result is robust OVERSEER framework which can be re-used at low-cost to support similar field trials in future ESA and other activities. As components were defined at a high-level of abstraction alternate approaches can be evaluated quickly and in context. The purpose of a high fidelity field trial is in a sense to enter the unknown and this certainly proved true for the Atacama field trials. Although the team had carried out extensive and progressive testing in the UK in Mars Yards, quarries and beach locations there is no substitute for the diversity and combination of both adverse terrain and lighting conditions when testing a vision based system. In this regard the Seeker Valley site proved to be extremely challenging and representative. The main challenges included – featureless, saturated terrain, un-predictable slip and variable slope, varying size boulder fields. In our view it is impossible to emulate these conditions in local sites particularly the given the scale and extent of the Atacama range. Our final assessment having also worked in areas such as El Tiede in Tenerife is that a vision based localisation and navigation system can only really be validated and verified in a location such as this or alternatives such as Morocco. The project has achieved a first in the ESA space context with kilometre order autonomous traverses in such challenging terrain. In addition the system framework, data set and components are highly adaptable and could be re-used at low cost to trial similar approaches.

5. Acknowledgements

The authors would like to acknowledge the support of the European Space Agency who funded this work. The authors would also like to express our gratitude to ESO Paranal for their generous and support and assistance during the trial period.

6. REFERENCES

1. D. Wettergreen, N. Cabrol, S. Heys, D. Jonak, D. Pane, M. Smith, J. Teza, P. Tompkins, D. Villa, C. Williams, M. Wagner, A. Waggoner, S. Weinstein, W. Whittaker. "Second Experiments in the Robotic Investigation of Life in the Atacama Desert of Chile", In Proc. 8th Int. .Symp. on Artificial Intelligence, Robotics and Automation in Space, Munich, Germany, 2005

2. M. Woods, A. Shaw, A. Gily, F. Didot, "High-Level Autonomy for Long-Term Exploration Robotics", Proc, 11th Symposium on Advanced Space Technologies, Robotics and Automation ASTRA 2011, ESTEC, Noordwijk, The Netherlands

3. C Mei, G Sibley, M Cummins, P Newman and I D Reid, (2011). RSLAM: A System for Large-Scale Mapping in Constant-Time using Stereo, *Computer Vision,* **94** (2), 198-214,

4. Harris, C. (1993) Geometry from Visual Motion, *Active Vision*, (Eds. Blake & Yuille) MIT Press.

5. Bakambu, J., et al (2012). *Field trial results of planetary rover visual motion estimation in Mars analogue terrain*. Journal of Field Robotics. **29** (3), 413-425