Seeker – Autonomous Long Range Rover Navigation for Remote Exploration

Mark Woods, Andrew Shaw, Estelle Tidey, Bach van Pham, Simon Lacroix, Raja Mukherji, Brian Maddison, Gary Cross, Aron Kisdi, Wayne Tubby, et al.

► To cite this version:

Mark Woods, Andrew Shaw, Estelle Tidey, Bach van Pham, Simon Lacroix, et al.. Seeker – Autonomous Long Range Rover Navigation for Remote Exploration. *Journal of Field Robotics*, 2014, 31 (6), pp.861 – 996. 10.1002/rob.21528 . hal-01109203

HAL Id: hal-01109203

<https://hal.science/hal-01109203>

Submitted on 26 Jan 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Seeker – Autonomous Long Range Rover Navigation for Remote Exploration

---

Mark Woods⁽¹⁾, Andrew Shaw⁽¹⁾, Estelle Tidey⁽²⁾, Bach Van Pham⁽³⁾, Lacroix Simon⁽³⁾, Raja Mukherji⁽⁴⁾, Brian Maddison⁽⁵⁾, Gary Cross⁽⁶⁾, Aron Kisdi⁽⁵⁾, Wayne Tubby⁽⁵⁾, Gianfranco Visentin⁽⁷⁾, Guillermo Chong⁽⁸⁾

⁽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, bvpham@laas.fr

⁽³⁾ Univ de Toulouse, LAAS, Toulouse, FR, simon.lacroix@laas.fr:

⁽⁴⁾ MDA Space and Robotics Limited, Didcot, Oxfordshire, OX11 0QX, UK, raja.mukherji@mdacorporation.com:

⁽⁵⁾ Rutherford Appleton Laboratory, Didcot, Oxfordshire, OX11 0QX, brian.maddison@stfc.co.uk:

⁽⁶⁾ BAE Systems, Advanced Technology Centre, Filton, Bristol, BS34 7QW, UK, gary.cross4@baesystems.com:

⁽⁷⁾ ESA/ESTEC, Keplerlaan 1, Postbus 299, 2200 AG Noordwijk, The Netherlands, Gianfranco.visentin@esa.int

⁽⁸⁾ Departamento de Ciencias Geológicas Universidad Católica del Norte, Antofagasta, Chile

Abstract

Under the umbrella of the ESA StarTiger program, a rapid prototyping study called Seeker was initiated. A range of partners from space and non-space sectors were brought together to develop a prototype Mars rover system capable of autonomously exploring several kilometres of highly representative Mars terrain over a 3 day period. This paper reports on our approach and the final field trials that took place in the Atacama Desert, Chile. Long range navigation and the associated remote rover field trials are a new departure for ESA and this activity therefore represents a novel initiative in this area. The 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 Martian day. The paper does not seek to introduce new theoretical techniques or compare various approaches but it offers a unique perspective on their behaviour in a highly representative environment. The final system autonomously navigated 5.05 km in highly representative terrain during one day. 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

Since early 2004, both NASA's Spirit and Opportunity Rovers have logged a remarkable set of distances and have broken new records in terms of autonomous mobile exploration on another planetary surface. In August 2012 Opportunity completed a 21.5 km traverse ("Opportunity Reaches Endeavour Crater", 2011) over a three year period as it traversed from Victoria to Endeavour crater. In that same month, we also witnessed the successful landing of Curiosity, which has begun a long but already fruitful traverse towards Mount Sharp. As of December 2013 the European Space Agency is in the production phases of development for the 2018 ExoMars mission (Van Winnendael et al 2005) which will seek to find evidence of life on the Martian surface. Overall these missions provide a benchmark for state of the art exploration on

Mars. It is an important juncture, therefore, to consider the next steps in terms of improving exploration capability as the space community seeks to enable the next generation of science and exploration goals.

Improving rover traverse distances by at least an order of magnitude represents one significant research challenge. In the wider context, long distance travel alone is not sufficient however as scientists may be reluctant to make the trade-off of fast traverses for lack of scientific analysis en-route. A second challenge is that of science autonomy, whereby rovers would be able to selectively collate information about the surrounding terrain to negate the need for this trade-off. Work towards this challenge has been carried out by some members of the Seeker team (Woods et al 2009), (Paar et al 2012) and provision made for this capability as part of the Seeker software solution although the primary goal was that of long distance travel.

A key feature of the missions outlined above is distance travelled per sol which, in most cases, is in the order of 10, 100 or 200 m per sol, where sol is a Martian day. For future missions such a limit acts as a substantial constraint on remote exploration activities. If we consider Opportunity's recent traverse from Victoria to Endeavour crater, the rover accumulated approximately 21 km over 3 years. Given the original scope of the mission this was a wonderful achievement, but for reference a similar distance was travelled in one day by the Apollo 17 Astronauts on the moon (The Apollo Lunar Roving Vehicle 2011). It is therefore appropriate to seek to extend the reach of our current mobile platforms to allow science teams to cover more ground in less time and ultimately extend the reach of their surface and sub-surface investigations.

There are a number of reasons why state of the art planetary rovers travel slowly such as low power, low processor speeds and limitations in the required vision software. Given recent advances in terrestrial state of the art computer vision systems, it seems reasonable to assume that the software aspects of the problem could be improved, as they have already been shown to be functionally capable of navigating over longer distances on Mars like terrain. A key question however is whether or not such techniques, which rely heavily on their ability to distinguish a well spread feature set, can handle the harsh visual conditions and homogeneity of the desert like terrain observed on many parts of Mars.

This Seeker project specifically set out to establish whether or not vision based techniques could be reliably used over single figure kilometre distances per sol in such terrain. Power issues and computational limitations were beyond the scope of the project, focussing instead solely on the issue of using vision based navigation over long (km) distances to achieve reliable autonomous navigation. In order to fully address the objectives of the project also sought to include high-fidelity field trials by arranging for testing in a suitable Martian analogue site.

Work by other groups such as (Huntsberger, Aghazarian et al., 2002), (Wettergreen et al., 2005), (Fong, Allan et al., 2008) and Barfoot (Furgale & Barfoot, 2010) has demonstrated the significant benefits of testing in such an environment. In short, it provides the diversity of terrain and lighting conditions which a rover can be expected to encounter on Mars, but are difficult or impossible to emulate in small or synthetic test environments or local European natural sites. The Atacama Desert in Chile was selected as the trial site. Section 3 discusses this further.

Figure 1: ESA's ExoMars rover deployed on the surface - artists impression.

1.1 ESA Context and Related Work

ESA mid-term strategy with respect to the exploration of Mars foresees the participation in an International Mars Sample Return (MSR) mission in the mid2020s. ESA has therefore identified that technology which makes it possible for rovers to perform long traverses (several kilometres) in a short time (few tens of sols) is critical for many of the possible MSR mission configurations.

To achieve faster mars rovers, 2 lines of development need to be pursued:

- Improved energy-to-motion ratio: the (solar) energy available to Mars rovers will not increase in future missions. To move for a longer distance each sol, a rover must be able to convert much more of the solar energy it collects to locomotion. Locomotion must be made more energetically efficient.
- Increased autonomy of operation: moving faster means that the distance traversed between 2 telecommunication sessions to ground will be greater; hence the rover will need to navigate autonomously for longer extents and will have to take many more decisions autonomously.

In terms of extra-terrestrial state of the art, the longest single sol traverse record is held by MER Opportunity at 220 m on Sol 410 (Biesiadecki et al 2005). Such drives are a mix of both so-called directed or blind and Auto-Nav modes (Biesiadecki et al 2007). Directed drives are faster but rely on good observability by the operations team which has limits depending on the terrain. Auto-Nav drives are approximately 3 times slower than a directed drive. The use of visual odometry also reduces the traverse speed significantly. Running both Auto-Nav and VO

reduced the MER rover nominal (direct) drive speed from 124 m/h to approximately 6 m/h. NASA's Curiosity Rover traversed approximately 500 meters between sol 21 and sol 56 of its operation. It is ESA's ExoMars rovers will have similar distance travelled characteristics as the MER and Curiosity rovers. (Volpe, 2003) provides a good overview of the functional components of the Curiosity rover's GNC system.

In recent years pioneering work in the field of long distance autonomous navigation has been carried out by (Wettergreen et al. 2008). This work established the value of the Atacama desert for Mars field testing, long-range navigation and science autonomy. In the reported trials, relative localisation was dependent on non-vision based sensors such as wheel encoders, chassis kinematic sensors, yaw gyros and inclinometers. Stereovision was of course used for near-field terrain assessment and obstacle avoidance. In the Seeker case Visual Odometry techniques were the principal source of relative localisation pose estimates.

Visual Odometry (VO) coined by Nister (Nister et al 2004) provides a means of estimating the position and orientation of single or multiple cameras based on image input by detecting and tracking features over time, and has become a widespread means of relative navigation for mobile robots (Konolige, Agrawal et al., 2007)(Maimone, Chang et al., 2007)(Sibley, Mei et al., 2010)(Souvannavong, Lemaréchal et al., 2010)(Lamber, Furgale et al., 2011)(Bakambu, Langley et al., 2012)(Wagner, Wettergreen et al., 2012). This relative localisation technique was complemented by a low frequency absolute localisation process that exploits low resolution Digital Elevation Model (DEM) derived from orbiter data. The DEM refers to a matrix containing only elevation or height data after (Li et al 2009). This allows drift correction over time. Other work (Dupuis et al 2006) has been carried out into long-range navigation development at the Canadian Space Agency (CSA). The path planning element of this work was integrated into Seeker through our MDA partners. A key difference was that the earlier work relied on LIDAR based DEM generation whereas the Seeker system was purely dependent on stereo based DEM construction.

Visual teach and repeat (Furgale & Barfoot, 2010a) and networks of reusable paths (Stenning, et al., 2013) have also been used for long range rover autonomy in analogue environments such as Devon Island, Sudbury impact crater, and the Mistastin impact structure. This work differs from ours in that we consider a single shot scenario i.e. there is no opportunity to teach and then repeat as the rover must navigate through the terrain without any prior in-situ knowledge.

1.2 Objectives and Contribution

1.2.1 Objectives

At the time of commissioning this work, there was uncertainty in the space community about the suitability of using feature based computer vision algorithms for long-range, autonomous navigation in a harsh Mars like environment. It was not clear if such algorithms could be used reliably beyond 100-200 meters. In an attempt to address this uncertainty the Seeker project

sought to determine the feasibility of using such techniques in a highly representative Mars environment by attempting to prototype a rover system which could:

Autonomously navigate 6 km of a high-fidelity Mars analogue site over a 3 day period using so called standard rover sensors i.e. a stereo camera, IMU, wheel odometry and sun sensor.

The following steps were carried out in order to reach this goal:

- A prototype autonomous navigation system was developed which consisted of:
  - Two types of Visual odometry based localisation
  - DEM mapping based on feature-based stereo matching
  - Path Planning
  - 3D surface matching for absolute localisation
  - A high-level decision making executive component and associated operations support
- Integration of the software listed above onto a robust rover platform known as Robovolc (Muscato et al, 2001)
- Conduct test trials in the UK and in the Atacama Desert in Chile
- Execute the above in a short 6 month time frame according to a prescribed “Skunk Works” style methodology

1.2.2 Contribution and Scope

This paper is a field report documenting the prototyping and deployment of a vision based autonomous navigation system in a representative environment. Its contribution can be summarised as follows:

- The paper report focuses primarily on the lessons learned whilst attempting to achieve this difficult challenge from a practical field robotics perspective.
- To the best of our knowledge this is the first non-learning based VO-centric system to autonomously travel km order distances in the highly representative Mars like terrain of the Atacama Desert.
- Prior to this work there was no known demonstration of the feasibility of such techniques in this type of terrain over these distances.
- Such a demonstration is required in order to enable future long-range Mars Rover missions.

- The results of this work have therefore informed both the ongoing ESA ExoMars work and the future Sample Fetch Rover concept.
- We have also demonstrated the integration of absolute localisation and relative localisation techniques in a UK environment.

Although not part of the original scope we have collected a large dataset of highly representative image data which can be used to test other vision based systems – this dataset has been made available to the robotics community through ESA.

Figure 2 compares the proposed Seeker functional and high-level performance scope with that of current space rovers.

Figure 2: Proposed Seeker Functional Scope versus the current Mars Mission State of the Art.

The following sections in this paper outline the Seeker system, the field trials, results and final conclusions.

2. Approach & System Overview

The project methodology, at a high level, was as follows:

- Conceive a system architectural concept
- Assemble component technologies for integration
- Integrate components in a flexible, modular software architecture
- Deploy on a suitable platform(s)

- Conduct outdoor field tests in order to shakedown the system components and architectural concepts
- Extend scale and variety of field trials toward goal conditions
- Select and reconnoitre a suitable Mars analogue site
- Conduct final field trials

Given the short 6 months time-frame the development followed an agile model based around week long iterations and required the introduction of prior technological elements from all partners. Figure 3 below shows the high-level system target architecture, consisting of a set of key components which included an off-board control centre interface and a variety of on-board autonomous components.

The architecture followed a basic three-layer approach with additional, off-board mission planning capabilities. Nominal plans were prepared by the operations team on the ground and uplinked for execution. Plan dispatch was handled by the on-board executive which commanded individual subsystems at a relatively high level of abstraction. For the most part, the navigation sub-system was in control of the vehicle unless it was unable to compute a suitable path or high-level resource decision making was required. For these situations higher-level re-planning and execution control was available.

From the outset it was decided to attempt to achieve the long-range navigation goals with the same sensor suite employed by the current fleet of Mars Rovers, i.e. stereo-cameras, wheel odometry and an IMU. As discussed in the introduction the project scope did not extend to the development or use of highly representative platforms for the demonstrations given that this was not essential to test the software aspects. However it was essential that a robust platform capable of operating in harsh conditions over many months was adopted to support the trials. The Robovolt platform was therefore used as our primary vehicle given that it was initially developed for exploration of harsh volcanic slopes.

Figure 3: System Architecture – High-Level View. DROID and OVO are independent visual odometry (VO) implementations provided by Roke Manor Research and the Mobile Robotics Group at Oxford respectively

Figure 4: Seeker platform –Robovlc provided by BAE Systems. Bumblebee Stereo-pair mounted on the forward facing mast. Generator for extended operations mounted on the rear of the vehicle. D.P.G.S (ground truth testing) and Wi-Fi antenna in the centre.

2.1 Autonomy Software Framework

A key challenge in the Seeker project was the requirement to quickly assemble a complex navigation software system and have it deployed on suitable platforms in a very short time to facilitate quick testing. The system was composed of existing algorithmic implementations for functions such as path planning, mapping and visual odometry, which were brought together from different sources using a mix of co-ordinate reference frames. An autonomy software framework called OVERSEER was used to achieve integration of the various individual functional software functions. The framework defines common interface and abstract wrappers for typical components on an autonomous mobile system. Individual functions such as VO and DEM generation were supplied as libraries and then integrated in a dedicated component which handled co-ordinate transforms, data relevance checks, execution control and component-to-component communication. Although this generated a significant amount of system integration and test work in the early phases of the project, it greatly simplified error detection and analysis in the field where time was often short. Implementations could be swapped out without any software engineering impact on the overall system. This was essential for Seeker as alternate VO libraries were often used in isolation.

The off-board mission planning and control GUI for OVERSEER allowed the users to plan the rover activities at a high level. Initial, target trajectories were specified as a series of well separated, major waypoints on a 3D model of the test terrain. The on-board navigation system was free to decide how to reach each of the target waypoints.

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 were subsequently used by navigation for path generation and obstacle avoidance, and to estimate the absolute localisation of the rover. As a consequence, image capture and distribution provided the core event driven “heartbeat” of the system.

The camera interface library allowed the connection of multiple camera types i.e. USB, Firewire and Ethernet and provided 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 references 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 also the Oxford University OVO (Churchill 2012) algorithm which is being evaluated and integrated for space applications by SCISYS (Shaw et al. 2013 and Woods et al. 2011). 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. (Harris 1993). Since its publication, DROID has been shown to be a robust,

fast and well tested solution for terrestrial robot navigation using vision where GPS and active LIDAR scanning are not available. It has been applied to a range of scenarios; an example is the work carried out on indoor robot exploration under the Systems Engineering for Autonomous Systems Defence Technology Centre (SEAS DTC) (Harris, 2006, 2007; Harris et al, 2008)

The DROID algorithm performs well in Mars-like terrain due to the feature-rich texture of the ground. This is illustrated by Figure 5, 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 5: 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 ensures that the field-of-view contains a large feature-rich area (the ground) and that the cameras are 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 its field of view. A frame rate of 5Hz was selected as being fast enough to achieve good results at the required vehicle speed of 20 m/sec.

- 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 Pan and Tilt Unit (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 6 shows an image captured during calibration.

Figure 6: Camera calibration tool display, showing the vehicle origin and axes and the position of the target.

The OVO algorithm also performs very well in this environment due to the feature richness. To calculate the ego-motion estimate the algorithm assumes that both of the stereo images are captured at exactly the same time using global shutters. During each image capture the following sequence is performed;

1. **Image Rectification:** Each image is rectified to remove distortion,
2. **Feature Extraction:** The FAST corner extractor is used as it produces a high number of candidates at low computational cost, which are ranked using the Harris score. This step is performed at various image scales increasing the robustness to motion blur etc.
3. **Feature Spreading:** To make sure the features are distributed across the image they are processed using a quad-tree which constrains the number in a particular part of the image. This removes the clustering of features produced using Harris score only.
4. **Left-right matching:** Using the sum of absolute differences (SAD) the feature points in the left and right images are matched.
5. **Temporal Matching:** Using image pairs T and $T-1$ a Binary Robust Independent Elementary Features (BRIEF) descriptor is used, which allows fast matching between the image frames.

6. **Sub-Pixel Features:** The BRIEF descriptor provides a integer pixel location which is subsequently refined to sub-pixel accuracy using an efficient second order minimisation.
7. **Motion Estimation:** Using the feature points the 6 DOF motion can be computed using a RANSAC step which removes outliers, then a least squares minimisation is performed, any new outliers are removed and the motion estimation is achieved.

The localisation of the platform in the field was performed by the Localisation Manager component which received inputs from a variety of pose estimation sensors. Each of these pose estimation sensors 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. For example, during the trials, a lot of sandy areas had a hard crust which, when broken revealed soft sand 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 threshold of 15 meters is used to create a coarse DEM (1 meter in resolution) which is used to estimate the global position of the rover, according to a map-based localisation technique that exploits an overall DEM of the terrain generated from orbiter data (also 1 meter in resolution, as can be provided using images from the HIRISE camera of the Mars Reconnaissance Orbiter). 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 7: Left: Rover view of the terrain at a given position during trials in the UK. Right: particles (blue dots) that correspond to likely positions on the global DEM. The red cross is the ground truth, and particles are spread on areas that resemble the viewed terrain – a flat corridor neighbouring a steep slope.

To estimate the global position of the rover with respect to the global orbiter DEM, a discretized particle filter algorithm was used. Instead of randomly throwing the particles inside the global DEM, the particles are initialized by comparing the stitched local DEM around the rover with small patches of the global DEM using a zero-mean normalized correlation criterion

Figure 7 shows an example of a stereo image taken from the current rover position, and the initial position of the particles on the global DEM that are likely to match the current position. 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.

Figure 8: Conventional continuous particle distribution (left) with particles with similar orientation and position in one single cell and the proposed discretized particle distribution (right).

After having initialized the particles, the estimation of the rover global position is improved as the rover moves. For the sake of computational efficiency, instead of employing a continuous particle distribution method, a discretized distribution method is adopted: only a few particles that correspond to a set of possible rover headings are allowed in a single DEM cell (Figure 8 – in our implementation, the orientation is discretized with a step of 3 degrees).

To cope with the particle discretization, a Gaussian model is used to represent the position uncertainty of each particle, with a standard deviation of half a cell size. Once the uncertainty of a particle gets bigger than the resolution of the global DEM (i.e. 1 meter), the particle is decomposed into smaller particles with the same orientation but with different positions. Figure 9 illustrates a situation where a particle is decomposed into 3 particles.

Figure 9: 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 computational complexity low, the particles are only updated after a large distance and only when the local DEM exhibits significant relief.

2.5 Path Planning

The path planning task of our partners MDA, 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 10 illustrates the planned path for the 6 km traverse.

Figure 10: 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. The coarsely-spaced long range way points selected by the scientists are linearly interpolated to produce a sequence of intermediate way points with 2 m spacing. These short range way points are then used as potential local goals for incremental rover traverses. A 2 m planning horizon (determined empirically) allows the system to operate within the local 3D point cloud where the stereo reconstruction is most reliable, and is small enough to force the system to update often to detect and avoid obstacles at a spatial resolution finer than the global map image. Over the 2 m planning horizon, the environment is assumed to be static so that once a safe path is found in the local DEM, the rover only needs to track the generated path in order to reach the next short range goal. By updating often, the path planner effectively implements a real time obstacle path tracking and avoidance function, albeit at a low update rate which is achievable with the Seeker's onboard computational capability.

In order to generate a safe, obstacle-free path, the rover must first classify the terrain into traversable and non-traversable regions. The 3D terrain data was segmented into 15 x 15 cm cells in the horizontal plane. The terrain assessment algorithm (Bakambu, Langley, et al., 2012) is then operated on each cell of data to compute its traversability cost, based on the known ground clearance of the rover, the maximum slope that the rover is capable of traversing, and the roughness of the terrain. Based on these three criteria, the resulting cost map shows cells which are either unknown (due to lack of data), traversable, untraversable, or part of a safety boundary based on the known footprint of the rover. Short range path planning performs a D* search (Stentz, 1995) through the traversable cells to determine the minimum cost path.

Figure 11: 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. Note the rover's current location is at the left end of the short range path.

A higher level of autonomous decision making was included in the short range path planner in order to handle off-nominal cases. For example, if the requested short range goal (interpolated from the global way points) falls within a non-traversable cell, the planner automatically searches for a new goal location. The traversable cell with the smallest distance to the next short range goal is chosen as the new goal, and an attempt is made to find a safe path to this new location. If no path can be found, the planner component pauses the rover by commanding zero linear and angular velocity, and raises a flag so that the higher level autonomy can take corrective action. Possible actions include panning the camera to increase the coverage of the local 3D point cloud and increase the possibility of finding a traversable path, or sending the rover a new long range path to follow, essentially commanding the rover to deviate away from the current obstacles. Once corrective action is taken, path planning and tracking re-engage and the rover continues its traverse.

The rover tracks the short range path by using the pose estimate of the rover (Section 2.3) as feedback to a control law (Astolfi, 1999) which computes the required linear and angular velocities in order to stay on the path. These rate commands are then sent to the rover's motion control system. The control law is tuned to the specific rover platform to reduce tracking errors which could cause the rover to stray into an obstacle.

Once the rover reaches the short range goal, the path planning cycle is triggered again until either the rover reaches the long range goal, the executive requests a rover stop, or a new long range path is sent to the rover. The incremental path planning and tracking cycle can be broken down as follows:


```

while (long range goal not reached){
    load current 3D point cloud
    access point cloud for traversability given rover constraints
    compute short range goal path to nearest non-hazardous intermediate way point
    while (short range goal point not reached){
        Use current pose estimate and path to compute the linear and angular rate
        commands for the rover motion control.
    }
}

```

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 Off-Board OVERSEER Mission Planning

Resource management and high-level decision making for Seeker was provided by the SCISYS Execution and Decision making components. These components controlled the execution of the initial timeline, addressed resource issues, adapted the timeline in response to GNC events 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. Once 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 monitored 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) were 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 UI to check progress of the executing mission plan by overlaying reported pose, live DGPS ground truth, the planned route and absolute localisation matches.

Figure 12: OVERSEER monitoring view showing reported tracks.

2.7 Robovolc Platform and Control

A requirement 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 with 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 and pitch 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 speed range between 0.1 – 1 m/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.

The platform's 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. These enable the monitoring of vehicle telemetry and control of the locomotion of the Robovlc 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. Robovlc 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 2 km range.

3. Field Trials

3.1 UK Trials

After initial integration and test at the RAL test facility, the team moved their UK trials to a nearby mortar sand quarry called Tubney Woods, operated by Hills Quarries. Its primary purpose was to ensure that each component and subsystem met their required functional requirements. Co-ordinate transform harmonisation was a key issue for a project of this nature. Each component was built using heritage technology from various partners. As a consequence a variety of co-ordinate reference frames were used and this required a degree of harmonisation. It is worth considering that in order to complete a full navigation cycle with absolute correction involved up to 8 components processing 6 DoF position information. Each translation was a potential source of error and single point of (complete or gradual) failure. Subsequently the majority of the testing in Tubney Woods was focused on ensuring successful inter-component communication as well as individual functional behaviour utilising the DGPS for highly accurate x,y,z ground truth. The intermediate results of this testing are discussed mainly in the previous subsystem sections and will therefore not be repeated here.

Figure 13: Tubney Quarry during Seeker Mid-term Review (panoramic image).

For the most part the trials were short range – 10's or 100's meters. The maximum traverse was 256 meters in total. Tubney quarry testing also provided an excellent proving ground for UAV operations and integration with the absolute matching method. Furthermore, it allowed the team to test a range of light conditions from soft through to hard and other aspects such as adverse camera shake. A key difference from the final site (and goal for the project) was of course the fact that the quarry terrain was, from a vision perspective feature rich. It also was ideal for the absolute localisation matching method given the abundance of large ridges with unique and sharp elevation deltas.

As trials progressed and the system stabilised, the team moved field trials to a local beach – Weston Beach near Bristol. This provided accessible, long-range terrain which had a more challenging feature density profile than the relatively easy quarry. In addition, the team could quickly deploy to the beach area during the often wet and inclement weather conditions which rendered other inland sites difficult to use. The primary objective for conducting beach trials was to determine if the VO could cope with relatively feature sparse terrain over long distances in the order of several hundred meters. Figure 14 shows the general terrain which exhibits a relatively feature-sparse texture compared to Tubney quarry. The quarry space was quite constrained, allowing only a maximum traverse of 256 m and was visually feature-rich. During the beach trial a maximum distance of 600 m was achieved with VO recording errors in the order of 0.5% of distance travelled. The primary tests were obstacle free however short-range path planning tests were also conducted during this time. This provided evidence that VO could cope with moderately challenging terrain within the bounds of accessible environments within the UK and allowed plans for the remote tests to go ahead.

Figure 14: RoboVolc at Weston Beach during long range VO tests.

Figure 15: RoboVolc at Weston Beach performing obstacle avoidance manoeuvres

3.2 Chile Trials

Globally, there are several sites that can be considered for representative field trials (Preston et al 2012). At the outset of the project several global test areas were considered. Our final selection was based on a variety of factors based on a scoring of terrain suitability with respect to the vision based assessment criteria, logistics support, cost, security, and human factors. The preferred Atacama site scored highly on all criteria 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 Table 1 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.

The research area is located in the Coastal Range of Northern Chile, a main geomorphologic unit of the country. The average altitude can be considered as 2000m a.s.l. The geological setting of this area includes mainly intrusive granitic rocks associated with volcanic and less sedimentary stratified units from Upper Paleozoic up to Lower Cretaceous age. Younger units correspond to alluvial and colluvial rocks and sediments of upper Neogene and Quaternary age. Rhyolitic and andesitic lava flows, breccias and tuffs with minor sandstones and limestones. La Tabla Formation of Carboniferous - Permian age. For further information please see (Marinovic et al., 1995).

Comp	Diff	Terrain Discrimination Features			Lighting			Terrain Hardness			Rock Size			Rock Distribution			Gross Slope Steepness		
		Lo	M	Hi	Soft	M	Hard	Soft	M	Hard	S	M	L	Lo	M	Hi	Lo	M	Hi
VO	E			●	●			●	→	●	●	→	●			●	●	→	●
	M		●			●									●				
	H	●					●							●					
WO	E	NA			NA					●	●	→	●			●	NA		
	M	NA			NA			●			●	→	●			●	NA		
	H	NA			NA										●		NA		
Map ping	E			●	●			NA			NA			NA			NA		
	M		●			●		NA			NA			NA			NA		
	H	●					●	NA			NA			NA			NA		
Nav	E	NA			NA					●	●	→	●			●	NA		
	M	NA			NA			●			●	→	●			●	NA		
	H	NA			NA										●		NA		
Abs	E			●	●			NA			NA			NA			●	→	●
	M		●			●		NA			NA			NA					
	H	●					●	NA			NA			NA					●

Table 1: Terrain Selection Matrix

Figure 16: Pan of the operational area including base camp

The site provided a rich testing ground to support both shakedown and final trials containing:

- High contrast dried out mud-flats to test auto-exposure
- Huge areas of soft, undulating, featureless terrain with no indication of human activity – essential for testing VO performance
- Completely unpredictable slip conditions – apparently firm terrain could often be inconsistent and highly variable even over small scale areas leading to rover wheel slip
- Large boulder fields
- Craters and gullies which provided reference points for absolute localisation

The images below provide some examples of the terrain for context.

Figure 17: Images show the variety of visual conditions available in the Atacama which are ideal for testing feature-based localisation and map-building solutions.

3.2.1 Building a Reference 3D DEM

To support absolute localisation and operations planning UAV flights were conducted to provide 2D orthoimages and DEM's of "Seeker valley" covering a desert area of approximately 2 km x 3 km. These were then downgraded to lower resolution (1 m/pixel) in the z-axis to replicate the quality of orbital data available on a real mission. Initial shakedown activities were confined to an area close to the primary basecamp. The objective was to ensure that the majority of the final long distance trial took place over unmarked and previously unseen terrain. Operators did not physically visit the main trial route area in order to simulate a real rover mission i.e. they relied solely on the 2D and 3D satellite equivalent imagery for planning.

Figure 18: 6 flight orthoimage (out of the total 9 flight over the area) showing base camp (red), flight base locations (orange), safe existing tracks (purple), trial area (green) and RoboVolc (blue)

3.3 DEM Generation using UAV's

A Digital Elevation Model (DEM) was generated prior to final trial traverses to simulate a similar model that would be generated from orbital imagery prior to landing and possibly imagery from the lander during descent. The UAV that carried out the DEM's is a Quest 200. The Quest 200 has been chosen due to many useful features

- Advanced autopilot module from SkyCircuits which enables autonomous flights and data acquisition
- Relatively small size means Civil Aviation Authority UAV license is not required and it is easy to transport and quick to deploy
- Fixed wing design and gimballed camera means imagery can be collected even in high wind conditions
- Thrust propeller, and anti-vibration mount means high quality images and up to 1 cm ground resolution of DEM

The DEM was artificially degraded to 1 m resolution to simulate the current highest achievable data based on orbital imagery on Mars. Nevertheless the high-resolution DEM was used for performance analysis of the system during and after the traverse. Due to slightly higher than usual flights to increase area coverage, the quality of data gathered enabled the team to generate a DEM up to 2 cm in resolution. However the processing power required to generate such a detailed map for the entire trial site and the difficulty in handling such a large DEM file meant that the best DEM generated during the trial was of a resolution of 0.3 meters.

During the study the team identified the automatic generation of the DEM as one of the critical processes, as even on a high-end computer it takes a considerable amount of time (can be up to 12 hours or more). During the Seeker trial, another method using Citrix software was tested. The raw UAV data was transferred back to RAL over the internet and using the DEM generating software (Agisoft's photoscan) running on the RAL server, the study team was able to generate a DEM quicker. However such a process still took many hours since even though the processing time was reduced, this method required other tasks such as the data upload and connecting to Citrix which were time-intensive. This method also needed a reliable internet connection. Post trials alternative methods were investigated to speed up this process. The key methods were identified:

1. Using a dedicated computer with tailored performance: At least a quad core i7 and multi code Nvidia card with CUDA support and preferably of the GeForce series. The clock frequency of the graphics card is also key and should be as high as possible. 16 GB memory is minimum and should be 32 GB if possible. At least 1600 MHz DDR 3

and an at least twice as big UHS-1 SD card or similar speed USB 3.0 flash drive is needed. A fast SSD is preferred for fast access of active data with secondary HDD for backing up already processed data. This PC or laptop can remain at the base and therefore doesn't need to be ruggedized. This would speed up processing time compared to the ruggedized laptops with limited hardware resources. Additionally it can run a light weight operating system to dedicate all resources to the processing.

2. Using an on board data logger connected to the UAV's flight computer to record telemetry and GPS data to speed up image alignment. This can be further improved by using both GPS and IMU data. Additionally differential GPS or other positioning methods can be used to improve the quality of this data.
3. Use of the flight software which is proven to be reliable during the trials to trigger the camera only when the required overlap (no more and no less) is reached compared to previous images. This reduces the number of images taken and cancels out the need for the semi-manual selection of useful images after alignment.
4. Use of automated software to check image quality and correct gamma, brightness, contrast and sharpness.
5. Use software on a dedicated low power computer on-board the UAV to analyse or combine the images on-the-fly. This would mean a faster lower resolution camera and a dedicated DSP attached to a low power computer with a dedicated software package.
6. Use different software capable of accurately modelling the flight path and orientation of the UAV using a combination of the on-board data logger, weather data, GPS data and analytical modelling. Such software could potentially generate the DEM quicker with less input required from the user.

Out of the points mentioned above, (1). can be deployed easily for any future projects using the best available computing technology at a time, (2). has already been implemented and tested since the end of the trials, while (3). and (4). are currently being addressed. (5). is a potential being considered with the availability of small and low powered ARM-processor based computers, but there is some development and optimisation needed in terms of the software such a device could run. Software for (6). has been identified which could be procured for future projects if necessary.

Figure 19: Quest UAV on launch

4. Results

The core objective for the Seeker project was, if possible, to demonstrate the feasibility of vision based, kilometre order navigation in representative terrain. Success would be marked by repeated rover traverses in excess of 1-2 km per day. From a StarTiger project perspective this was the key success criteria used to assess the outcome of the work. Throughout the project many discrete trials were executed in order to work towards this goal. The results presented here were used as a benchmark to measure progress towards that goal and primarily indicate the overall system performance with respect to the distance travelled metric.

For reference - results from the end of the UK preparation period are shown in order to indicate the system readiness prior to commencing field trials in the Atacama. Where possible DGPS or raw GPS was used to present some measure of ground truth in x,y,z and this referenced where appropriate in the results that follow.

The field trials in the Atacama consisted of a several day shakedown period followed by final day trials and ultimately concluded with an attempt to reach all of the project targets in one single day i.e. travel up to 6 km as a set of 3 segments on one single day. Given the relatively short periods of time available for operations at the time of the trials – approximately 08:30 - 17:30 h this was seen as an ambitious target. This section details the results for this final day trial showing distance travelled relative to ground truth where it was available.

In addition to these quantitative results, additional information is provided on the qualitative nature of the test terrain. For vision this is important as it defines the complexity of the localisation, map-building and path planning problem in particular. From a Space perspective the Atacama provides the ultimate challenge for any vision based navigation system given the sheer variety of terrain and lighting conditions tending towards vast swathes of homogenous, featureless terrain.

4.1 UK Trials

As discussed previously, UK trials were carried out in a variety of locations. Before traveling to Chile a final trial was carried out at Weston beach to evaluate system performance in relatively sparse featured flat terrain over several hundred kilometres. To test the system, a large L shaped trajectory was used which consisted of an initial drive of 250 m straight before turning right through 90 degrees for 10 m then a right turn through 90 degrees; traversing back 250 meters parallel to the original straight; then turn left through 90 degrees travel 70 m; turn right 90 degrees travel 25 m; turn right travel 80 m turn 90 degrees; then travel 25 m back to the starting point. The results for this trajectory are shown in the table below and in Figure 20.

Date and test description	Distance travelled VO (m)	Distance travelled GPS (m) (DGPS not available on this set)	Duration (mins)	Observations
May 8 – WestonRun3Oxford	136.4	146.9(error 3.7 m)	9	Navigated a straight line, the first part of the L-trajectory

May 8 – WestonRun4Oxford	719.5	793.9 (error 3.7 m)	58.3	Navigated the whole L-shape trajectory. Result is shown in Figure 20.
May 8 – WestonRun5Oxford	32.8	37.8 (error 5 m)	2.6	Navigated a straight line through a bolder field. Collided with one of the rocks
May 8 – WestonRun6Oxford	54.6	74.9 (error 5 m)	5	Navigated a straight line through the bolder field. Success
May 8 – WestonRun7Oxford	39.4	45.3 (error 4.2 m)	3.8	Navigated a straight line through the bolder field. Success

Table 2: Summary of UK Shakedown Test Results

Figure 20: L-shaped traverse on Weston beach covering 700 m. Red line is the GPS (not DGPS) output and with the green lines showing the fused localisation output from the LocMan component. Note the noise in the GPS measurements. The z error comes from the incorrect calibration of the tilt on the PTU.

Overall the results indicated good OVO and navigation performance with position estimate errors in the order of $<0.5\%$ distance travelled over almost 800 m.

Absolute localisation trials have been performed in the Tubney Woods quarry, using the DEM built from UAV images down-sampled to a 1 meter resolution to mimic a DEM generated from orbital imagery. One can assume that the absolute heading (yaw) angle of the rover is known thanks to the use of a Sun sensor (which was mimicked with GPS ground truth measures during our trials), that yield a heading precision of 3 degrees. Obtaining the absolute heading from the sun sensor allows initial orientation of the DEMs for a more efficient matching. However, the approach can handle no prior heading information, at the cost of more computational time during the initialization step.

To estimate the global position of the rover, we calculate the mean and standard deviation on the basis of the particles' positions. Figure 21 shows the evolution of the error of the estimated position as a function of time. The initial search zone here is the whole area covered by the global DEM, i.e. about one km: this corresponds to a lost-in-space (or "kidnapped robot") situation. In this case, the estimated position of the rover reaches a precision below 2 meters after 9 minutes. Note however that the convergence rate actually does not depend on the time but on the distance travelled, and more importantly on the elevation variance of the local DEM: of course, rather flat areas do not exhibit enough signal to allow discriminative local / global DEM matches – which prevented obtaining significant results in the Chile test site.

Once initialized, the filter update step is only applied when the local DEM has been significantly updated (e.g. a given threshold of the number of mapped square meters is exceeded). The particle's position and Gaussian distribution are continuously updated as the rover moves.

The particle initialization is used once and is by far the most computational expensive step. Initial computation time depends on the size and resolution of the local and global DEMs, on the angular resolution chosen to discretize the heading dimension, and obviously on the size of the area within which the rover is supposed to be located. Lost-in-space situations do not occur in actual missions, and the absolute localisation process is rather to be triggered after several hundreds of meter traverses, when the robot position error due to the drift of visual odometry amounts to several meters. With estimation errors of this order, the time required by the particle updates is at most 0.1 seconds for the worst cases, on a i7 2GHz Linux PC.

Figure 21: Evolution of the error of the positions estimated with the map-based localisation process as function of time. Blue, Red and Green lines represent error versus time for the X, Y and Z axes respectively

4.2 Chile Final Day Results

The following sections outline results from key elements such as localisation and path planning.

Relative localisation subsystem results are presented in three steps, Firstly, we look at raw VO results without any robustness checking or correction. Secondly we discuss the issues that had the biggest impact on the system performance. Thirdly we examine the improved performance that can be obtained using sub-system robustness measures.

During the shakedown period the system travelled in the order of 10 km including several runs in excess of 1 km. In total, the system covered in the order of 15 km over the entire Atacama trials period including debugging and test preparation. The entire final day 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 850 m segment started from base camp at sun-up with an extremely challenging traverse directly into the sun through highly reflective salt flats. These were successfully traversed, but the system appeared to stop analysing the DEMs for the path planning, which was noticed as the platform didn't periodically stop. This was found to be due to a faulty PCMCIA Firewire card that was used to interface the Bumblebee cameras to the laptop. From inspection it appeared that the constant use (10 days) in the hot environment and all the previous testing had taken its toll. The rover was placed back to a location that was calculated as being the point of the initial fail and the system was started again.

The second segment consisted of a 1.2 km autonomous run through the most challenging terrain the system had encountered – namely extensive boulder fields, considerable slopes and extensive soft sand. The system coped well with very difficult environmental conditions before being forced to stop because of thermal stress on the on-board computers and the firewire interface card.

The final segment of 3 km initially headed further south before turning west back across the valley. Along this lower stretch the rover came towards base camp at 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 an extensive boulder field. The test was stopped just after the 5 km mark with sundown approaching.

Figure 22: OVERSEER view of the planned trajectory through the Atacama environment, showing the three primary segments of travel as the rover moved from the start point (top-left) in a clockwise direction towards the end of Segment 3.

During the 5.05 km run, approximately 60 GB data was collected consisting of rectified image data, DEMs, wheel and IMU odometry and absolute locations. In itself this represents a hugely valuable data and image resource, although there is only partial ground truth coverage. 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.

4.2.1 Unfiltered Results

After reviewing the ground truth data a problem was identified with the captured data (the positions were buffered with respect to time) and the points were noisier than expected, (previously ~2cm but in places it was ~2m). This error meant that the position information could be used for ground truth but not correlated by time directly. As the three segments (1, 2 and 3 shown in

Figure 22) were continuous, the DGPS information obtained for the start of the second run could be used to position the end of the first run. Similarly the DGPS data from the start of the third run corresponded to the position at the end of the second run, but there was no data available for the end of the third run.

During the execution of the system on the final day, only the DROID VO was used to provide positional information. Although drift was witnessed, the Seeker architecture allows replay of all the data captured so that offline analysis of both VO systems with respect to the DGPS ground truth. The results presented in the following three figures therefore represent the in-field raw VO results obtained “as is” with no robustness checking/correction fusion applied.

Figure 23 shows the trajectories produced by the system for the first segment (before camera card failure). To align the data sets the same initial pose (aligned to world coordinates using the data obtained from the IMU) was used. The same camera mounting and calibration was also used in both instances. From this figure it's possible to see that both perform very well in the X-Y axis but there is a difference in Z drift which impacts on the overall error estimates in the final analysis. From analysing the image data the difference between the two comes from the miss-calculation of pose estimates when the vehicle traversed over rocks.

Figure 24 shows the same analysis for segment 2 of the traverse, with the starting orientation information obtained by the IMU. From this figure it's possible to see that neither VO is aligned particularly well. This is probably due to the initial orientation estimate from the IMU not being settled as the drift in orientation is from the start and is not gradual.

Figure 25 shows the same analysis for segment 3 of the traverse. Again, the starting orientation information was obtained by the IMU. From this figure it's possible to see that there is more drift in the system here which appears to be caused by the traverse through the boulder field (left side of the graph with movement, also shown in Figure 25). The initial orientation estimate from the IMU appears adequate in this instance.

Figure 23: June 2 segment 1 data processed, VO comparisons with DGPS. Total traverse distance for this segment was 831 m over 113 mins. Note: the issues with the noisy DGPS data are visible in red line and also the issue with the buffering where the line is broken. This graph shows good tracking of the x-y axis by the system and also how quick and large changes in pitch (movement over a rock) can affect the overall pose estimate with the DROID VO.

Figure 24: June 2 segment 2 data processed, review of the trajectories produced by the VOs compared to the DGPS output for the second segment, distance travelled is 1.5 km over 75

mins. Note: the issues with the noisy DGPS data are visible in red line and also the issue with the buffering where the line is broken. This graph shows that there could be a misalignment between the DGPS data and the rover localisation system caused by providing insufficient settling time for the IMU.

Figure 25: June 2 segment 3 data processed, VO comparisons with DGPS. Total traverse distance for this segment was 3400 m over 177 mins. Note: the issues with the noisy DGPS data are visible in red line and also the issue with the buffering where the line is broken.

4.2.2 Analysis

Relative position information was of course essential in terms of the overall system performance. XY estimates were used by path planning to control the platform trajectory. The complete pose estimate, including Z, is of course critical as this drives the accuracy of the DEM stitching which is fundamental for hazard analysis carried out by the path planning.

It's important to note the results presented in the previous section do not incorporate robustness checking or filtering, which has a significant impact on the actual position estimation. A key goal for the long range trials was to discover where the terrain would stress the different approaches to position estimation and then use this information to build in the appropriate system robustness and provide stable results. In a sense this is the key outcome of the work.

Considering raw component performance; Frame-to-frame odometry measurements in the XY plane were generally accurate, however at some points an error in the measured pose change led to a slow drift away from the absolute position measured by the GPS. This generally occurred in the Z direction (altitude), suggesting an error in pitch measurement. In one case the drift was in the XY plane (Figure 20, lower plot) suggesting an error in yaw.

The yaw error illustrates how an error in visual odometry measurement at a single instant can lead to a large drift over 1000s of metres. On investigation it was found that this was due to images not being consistently captured and processed at 5Hz. As described above, this is likely due to the computer overheating.

Errors may also have been caused by the vehicle jolting over rocks, which were seen to cause the PTU (and therefore the camera) to shake. The jolting may have caused the position of the camera to change permanently by a small amount compared to the initial calibration, which would alter the transformation between rover and sensor frames. Jolts may also have caused image degradation (such as motion blur) and reduce the overlap between images.

A possible solution to correct for Z drift would be to provide the estimated vehicle altitude from the Global DEM to the relative localisation component. This could be used to detect where the vehicle is drifting too far from the expected altitude and correct the pitch. Alternatively the IMU accelerometers could be used to provide pitch corrections by measuring the direction of gravitational acceleration when the vehicle is stationary. This is evaluated in the next section.

4.2.3 Impact of Robustness Measures

The graphs that follow show the impact of enabling various localisation fusion and robustness configurations.

Figure 23 shows the results from using just VO and IMU outputs only with the IMU weighted five times less than the VO in the EKF. This graph also shows the output from the wheel odometry where the estimated distance travelled is 820 m compared to the LocMan of 827m. This shows that although the forward distance is nearly correct, the rotational estimate is off. When using the IMU the RMS difference in the final pose estimate compared to the DGPS is 12.9 m for the fused LocMan and 16.6 m for VO alone and wheel odometry is 735 m. In this graph it's also possible to see that the fused output from the LocMan follows the pitch and the roll motion of the DGPS output closer than the output from the VO. This is probably due to the unevenness of the terrain causing constant shaking of the cameras introducing noise.. Figure 24 shows the same run but this time, the weighting of the IMU is the same as that for the VO. From this it's possible to see that there are some errors introduced when the vehicle did a pirouette- this time the LocMan diff was 35 m.

Figure 26: – Results showing how the localisation performed with the two VO's libraries simulating a continual drive through the environment. Here all of the images collected during the entire sequence were loaded into the both VO components and the position with respect to the starting point was measured. This shows how small errors through integration can result in large errors the further through the traverse, showing the need for the fusion with the other sensors.

Figure 27: Replay of segment 1 using Visual Odometry and IMU only. IMU weight is five less than the VO.. This replay only used the OVO library and shows how the drift in roll and pitch

can be minimised using the IMU pose as an input. It also shows how noisy the wheel odometry can be when using a skid steer platform of soft surfaces. Under these circumstances it is better not to include the WO into the pose estimate as it reduces the overall accuracy of the pose.

Figure 28: Replay of segment 1 using Visual Odometry and IMU only. IMU weight is equal to the VO estimates; NOTE the large change in pitch towards the end of the run. Again this image shows that by including the IMU input the overall pose can be improved, until the point where a large error is fused causing the trajectory to then subsequently drift. With further testing and parameter adjustment this could be improved upon.

Fusion alone is not sufficient to counter potential image acquisition and events which may derail the raw VO estimates. When additional techniques are put in place a significant improvement is possible.

A plot of the three sections above superimposed on an aerial image of the test site is shown in Figure 21. The drift in XY position, which begins in the third section, can be seen at the southern end of the path. Figure 22 shows the result of manually resetting DROID at the point at which drift starts to occur (effectively correcting for the missing images shown in Figure 19). This results in an improvement in localisation accuracy. The accuracy is still not as good as in the first two sections, which may have been due to a change in the camera calibration caused by jolting over rocks.

Figure 29: Replay of segment 1 using all the odometry sensors, Visual Odometry, IMU and wheel. IMU weight is five less than the VO and the wheel weight is ten less than the VO; NOTE the large change in yaw towards the end of the run caused by the miss-calculation of the yaw from the wheel slip.

Figure 21 – Uncorrected plot of vehicle position

Figure 22 – Corrected plot of vehicle position

It is clear from the trial results that there are many important considerations for use of visual odometry algorithms for autonomous planetary rovers. These include:

- I. **Calibration.** In order to use visual odometry measurements, it was found to be very important to perform an accurate calibration, both of the camera mounting on the vehicle and the vehicle attitude relative to the ground. An automated calibration scheme was devised for the project, however it requires further detailed testing. It is also crucial that the camera is securely and rigidly mounted on the vehicle, and if a pan or tilt manoeuvre is carried out, that it returns to the original position (or is accurately telemetered so any position change can be measured).
- II. **Frame rate.** Feature tracking between frames requires a reasonable amount of overlap in the scene viewed (greater than $\frac{1}{2}$ the field of view is recommended). The required frame rate will depend on the turn rate of the rover. In the UK trials, a standard laptop was able to capture and process images at the required rate as well as carrying out other processing such as obstacle detection and path planning. However this was not always possible in the Atacama Desert conditions. This could be addressed by actively monitoring the frame rate for mission management – for example, the vehicle might be slowed or even stopped if frame rate begins to drop. If some latency is acceptable, another option is to have an image buffer, and select images for processing according to image overlap.
- III. **Drift.** Failure to process a single frame can result in large drifts over time if the vehicle is executing a turn manoeuvre. It is therefore important to ensure that the required frame rate is met during turns (perhaps by reducing the turn rate or pausing lower-priority processes during this time) – but also highlights the importance of the absolute localisation component in correcting for drift if this occurs.

Overall, both VO systems DROID and OVO performed extremely well in Mars-like conditions, which provided a rich set of features to be tracked. The trials showed that it can provide accurate visual odometry over long distances and times. The experience of the trials highlighted several important system design considerations when using a visual odometry algorithm for localisation purposes, and provided a valuable test data set for future development work

During the third segment of the run, when the vehicle came back across the valley from east to west it encountered a large bolder field. As Robovolc was making its way through it came to a point where it had only rocks in the field of view causing the path planner to signal a “No Path to Goal” event. When a No Path to Goal event is raised the Autonomous Navigation component informs the Executive component it then autonomously pauses the system and commands a new panoramic image acquisition and DEM generation. This panoramic image offers a broader view of the scene around the rover. The wider DEM is then passed to the Autonomous Navigation component which seeks to use the larger map to plan its way out of the current cul de sac. No operator intervention is required whilst responding to a No Path to Goal event. This event was sent to the executive software component. This is an autonomous, software action and is part of the failure management response. The DEM result of this can be seen in Figure 30. Here the blockage can be seen towards the top of the image. Once this DEM

was given to the path planner it was able to find a way out of the “dead end” and it managed to successfully get the rover out.

Figure 30: —: DEM produced at during the trail after a "No path to goal" received. The open area in the middle of the circle section represents the area under the rover.

5. Conclusions, Lessons Learned and On-going Work

In total, the system covered in the order of 15km in fully autonomous driving mode over the entire Atacama trials period including debugging and test preparation. Figure 31 shows the cumulative distances travelled and daily run totals over the shakedown period and final day.

Figure 31: Shakedown Cumulative Distances Travelled.

Figure 32: Light conditions on the final segment of the Seeker traverse terminating at twilight.

During the final day 5.05 km run, approximately 60 GB data was collected consisting of rectified image data, DEMs, wheel and IMU odometry. Using the OVERSEER architecture it is possible to replay the run in “real-time” with the collected data allowing changes in fusion weights, VO configuration, DEM generation parameters etc. so that individual components performance can be evaluated. We are using this capability to continually refine the system performance for future field and potential mission use.

Perhaps the single biggest benefit of the Atacama site was the unparalleled range of terrain lighting conditions with which to test vision based localisation, DEM construction and path planning performance in particular. One of the most challenging aspects of this terrain is the complexity of path planning with limited guidance from an operations team. The complexity and relative variety of the landscape encountered over long distances presents a huge challenge 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 (~60 km over the entire project); ensuring further time for software integration and experimentation. The front and rear axles' articulation provided great agility whilst traversing across the undulating terrain, particularly during the dense boulder fields frequently seen in Seeker Valley.

Autonomy system engineering in this study was challenging given the ambition and scale of the work. However, the final result is a robust software 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, variable slope gradients and 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 Teide in Tenerife, is that a vision based localisation and navigation system can only really be validated and verified in a location such as this.

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. In total we estimate that several 10's of kms of autonomous drives were executed throughout the brief duration of the project.

Ultimately we believe that the Seeker work met its initial goal, which was to determine whether or not vision based localisation algorithms could support long range navigation on the surface of Mars. This work indicates that algorithms such as VO and DEM generation can indeed support long-range navigation provided that each component is part of a larger software subsystem which ensures robustness for situations where the components will produce erroneous results. A key outcome of this work is that the partners involved now have a detailed understanding of the failure cases and how to incorporate the necessary robustness to achieve

the required performance for missions such as ExoMars Rover and Sample Fetch Rover. Since the completion of the project, various components have been evaluated successfully on other space and non-space applications. Flight processor suitability has also been established for key localisation elements incorporating robustness measures derived from the results of these trials. Finally, we recognise the value of the extensive dataset which has been collected during the execution of this work.

Acknowledgements

The authors would like to sincerely acknowledge the support and assistance provided by Hills Quarries, David Wettergreen, David Thompson, Andreas Kaufer, the team at the European Southern Observatory Paranal and Patricio Arias Ortiz.

References

- Astolfi, A. (1999) Exponential stabilization of a wheeled mobile robot, ASME, Journal of Dynamic Systems Measurement and Control, March 1999, Vol: 121, Pages: 121 - 126
- 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
- Furgale, P. & Barfoot T.D., (2010a). Visual Teach and Repeat for Long-Range Rover Autonomy, Journal for Field Robotics, 27(5), 534-560.
- Furgale, P. & Barfoot T.D., (2010b). Stereo Mapping and Localisation for Long-Range Path Following on Rough Terrain, IEEE International Conference on Robotics and Automation, Anchorage, USA
- Biesiadecki, J., Baumgartner, R., Bonitz, R., Cooper, B., Hartman, F., Ledger, C., Maimone, M., Maxwell S., Trebi-Ollenu, A., Tunstel, E., Wright, J. (2005), Mars Exploration Rover Surface Operations: Driving Opportunity at Meridiani Planum, 2005 IEEE Conference on Systems, Man, and Cybernetics (Kona, Hawaii)
- Biesiadecki, J., Ledger, C., Maimone, M., (2007), Tradeoffs Between Directed and Autonomous Driving on the Mars Exploration Rovers, The International Journal of Robotics Research, 26(2), 91-104
- Churchill, W., 2012, Experience Based Navigation: Theory, Practice and Implementation, Ph.D. Thesis, University of Oxford
- Dupuis, E., Bakambu, J., Bakambu, J., I., Bakambu, J., J-L., Gemme, S., Rivest-Caissy, J-P., (2006), Autonomous Long-Range Rover Navigation – Experimental Results, In 9th ESA Workshop on Advanced Space Technologies for Robotics and Automation, ESTEC
- Diamantopoulos, D.; Siozios, K.; Lentaris, G.; Soudris, D.; Rodrigalvarez, M.A., (2012). SPARTAN project: On profiling computer vision algorithms for rover navigation, NASA/ESA Conference on Adaptive Hardware and Systems (AHS), Erlangen Germany

Fong, T.; Allan, M.; Bouyssounouse, X.; Bualat, M. G.; Deans, M.; Edwards, L.; Fluckiger, L.; Keely, L.; Lee, S. Y.; Lees, D.; To, V.; Utz, H. (2008). Robotic site survey at Haughton Crater. In *Proceedings of the 9th Symposium on Artificial Intelligence, Robotics, and Automation in Space (iSAIRAS 2008)*.

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

Harris, C. (2006), “Visual Exploration of Buildings”, 1st SEAS DTC Conference – Edinburgh 2006

Harris, C. (2007), “Strategies for Visual Exploration of Buildings”, 2nd SEAS DTC Conference – Edinburgh 2007.

Harris, C., Evans, R. and Tidey, E. (2008) “Assessment of a visually guided autonomous exploration robot”, Proc. SPIE 7112, Unmanned/Unattended Sensors and Sensor Networks V, 71120Y (October 16, 2008)

Huntsberger, T., Aghazarian, H., Cheng, Y., Baumgartner, E., Tunstel, E., Leger, C., Trebi-Ollennu, A., Schenker, P. (2002). Rover autonomy for long range navigation and science data acquisition on planetary surfaces. In *IEEE International Conference on Robotics and Automation*, Washington DC, USA, May 2002

Konolige, K., Agrawal, M., Sola, J. (2007). Large scale visual odometry for rough terrain. In *Proceedings of the International Symposium on Research in Robotics (ISRR)*, Hiroshima, Japan, November 2007

Lambert, A., Furgale, P., Barfoot, T. D., & Enright, J. (2011). Visual odometry aided by a sun sensor and inclinometer. In *Aerospace Conference, 2011 IEEE* (pp. 1-14). IEEE.

Li, Z., Zhu, Q. & Gold, C., (2009). *Digital Terrain Modeling Principles and Methodology*. Boca Raton, CRC Press.

Maimone, M., Chang, Y., and Matthies, L. (2007). Two years of visual odometry on the Mars exploration rovers. *Journal of Field Robotics*, 3(24):169–286.

Marinovic, N., Smoje, I., Maksaev, V., Hervé, M. & Mpodozis, C., 1995 : Hoja Aguas Blancas. Región de Antofagasta. Carta Geológica de Chile, Escala 1:250.000. N° 70. Servicio Nacional de Geología y Minería. Subdirección de Geología. Chile.

Muscato, G., Guccione, S., Savalli, N., Coltelli, M., Puglisi, G., Briole, B., Faucher, C., Virk, G.S., Azad, A., Semerano, A., Duporquet, V., White, T., Rees, G.S. (Sept 2001) Further progress in the development of the ROBOVOLC system, 4th International Conference on Climbing and Walking Robots (LAWAR 2001), Karlsruhe, Germany

- Nister, D., Naroditsky, O., & Bergen, J., (2004). Visual Odometry, IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Washington D.C., USA.
- Parr G., Woods M., Gimkiewicz C., Labrosse F., Medina A., Tyler L., Barnes D., Fritz G., Kapellos K., (2012), PRoViScout: a planetary scouting rover demonstrator, SPIE Intelligent Robots and Computer Vision XXIX: Algorithms and Techniques, Burlingame CA, USA.
- Preston, L.J., Barber, S.J., Grady, M.M., (2012. CAFE – A New On-line Resource for Planning Scientific Field Investigations in Planetary Analogue Environments), 43rd Lunar and Planetary Science Conference, Woodlands, Texas, USA.
- Shaw, A., Woods, M., and Didot, F., “Robot Visual Odometry for Space Exploration”, Proc, 11th Symposium on Advanced Space Technologies, Robotics and Automation ASTRA 2011, ESTEC, Noordwijk, The Netherlands
- Sibley, G., Mei, C., Reid, I.D., Newman, P.N. (2010). Vast-scale Outdoor Navigation Using Adaptive Relative Bundle Adjustment. In International Journal of Robotic Research. vol. 29, No. 8, 958-980, 2010
- Souvannavong, F., Lemaréchal, C., Rastel, L., Maurette, M. (2010). Vision-based motion estimation for the ExoMars rover. In International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS), Sapporo, Japan, August 2010
- Stenning B E, McManus C, and Barfoot T D. “Planning using a Network of Reusable Paths: A Physical Embodiment of a Rapidly Exploring Random Tree”. Journal of Field Robotics, special issue on “Space Robotics”, 2013. doi:[10.1002/rob.21474](https://doi.org/10.1002/rob.21474)
- Stentz, A. (1995). The focused D* algorithm for real-time replanning. In the 14th Intl. Joint Conference on Artificial Intelligence, Montreal, Canada, August 1995
- Van Winnendael, M., Baglioni, P., Vago, J. (2005). Development of the ESA ExoMars Rover. In Proceedings of 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space, Munich, Germany.
- Volpe, R. (2003, March). Rover functional autonomy development for the mars mobile science laboratory. In *Proceedings of the 2003 IEEE Aerospace Conference* (Vol. 2, pp. 643-652).
- Wagner, M., Wettergreen, D., & Iles, P. (2012), VISUAL ODOMETRY FOR THE LUNAR ANALOGUE ROVER" ARTEMIS, In International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS), Turin, Italy, September 2012

- Wettergreen, D., Wagner, M. D., Jonak , D., Baskaran , V., Deans, M., Heys , S., Pane, D., Smith, T., Teza, J., Thompson, D. T., Tompkins, P., Williams, C., (2008), Long-Distance Autonomous Survey and Mapping in the Robotic Investigation of Life in the Atacama Desert, International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS)
- Wettergreen, D., Cabrol, N., Heys, S., Jonak, D., Pane, D., Smith, M., Teza, J., Tompkins, P., Villa, D., Williams, C., Wagner, M., Waggoner, A., Weinstein, S., Whittaker, W., “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
- Woods, M., Shaw, A., Gily, A., Didot, F., “High-Level Autonomy for Long-Term Exploration Robotics”, Proc, 11th Symposium on Advanced Space Technologies, Robotics and Automation ASTRA 2011, ESTEC, Noordwijk, The Netherlands
- Woods, M., Shaw, A., Barnes, D., Price, D., Long, D., Pullan, D., Autonomous science for an ExoMars Rover-like mission, Journal of Field Robotics, Volume 26(4), 358-390.
- "The Apollo Lunar Roving Vehicle". National Aeronautics and Space Administration. Archived from the original on 15 July 2011. Retrieved 26 August 2011 from nssdc.gsfc.nasa.gov/planetary/lunar/apollo_lrv.html
- “Opportunity Reaches Endeavour Crater “, Retrieved December 2011, from http://marsrover.nasa.gov/mission/status_opportunityAll_2011.html
- Curiosity <http://mars.jpl.nasa.gov/msl/multimedia/interactives/learncuriosity/index-2.html>
Accessed June 2013.