EXOMARS VISLOC – THE INDUSTRIALISED, VISUAL LOCALISATION SYSTEM FOR THE EXOMARS ROVER

Daniel Townson[1], Mark Woods[1], Stuart Carnochan[1]

[1]SCISYS, 23 Clothier Rd., Bristol, BS4 5SS, UK, E-mail: daniel.townson@scisys.co.uk

[1]SCISYS, 23 Clothier Rd., Bristol, BS4 5SS, UK, E-mail: mark.woods@scisys.co.uk

[1]SCISYS, 23 Clothier Rd., Bristol, BS4 5SS, UK, E-mail: stuart.carnochan@scisys.co.uk

1 INTRODUCTION

Maintaining accurate knowledge of the current position of vehicles on the surface of Mars is a considerable problem. The lack of an orbital GPS means that the absolute position of a rover at any instant is very difficult to determine, and with that it is difficult to accurately and safely plan hazard avoidance manoeuvres.

Some on-board methods of determining the evolving POSE of a rover, such as using wheel odometry, orbital tracking, inertial measurement or LIDAR are well known. However there are associated problems:

- Wheels can slip in the Martian soil providing odometry readings which can mislead navigation algorithms.
- In orbit tracking and inertial systems tend not to be accurate enough to use in isolation.
- Complex localization systems, such as LIDARs, tend to be heavy and power hungry.

Visual localization systems, which use simple, lightweight cameras to determine the actual rover motion from images of the terrain, have been shown to be extremely accurate. By measuring movement from the terrain an independent measure of the actual movement can be obtained to a high degree of accuracy.

This paper presents SCISYS's project to develop the industrialised, Visual Localisation flight software system for the ExoMars rover (VisLoc). It describes: the terrestrial history of VisLoc, how it has been adapted and qualified for use on Mars and its potential future directions.

2 EXOMARS VISLOC HISTORY

The core algorithm, known as OVO (Oxford Visual Odometry), was developed at the Oxford Robotics Institute at the University of Oxford [1]. The OVO algorithm was originally targeted at terrestrial based autonomous navigation and was known to be robust and fast, capable of real-time operation on modest hardware.

From this background SCISYS used OVO as a visual localization system on prototype rovers. Trials were performed in the Atacama Desert in Chile. During these trials, the OVO algorithm performed admirably, enabling the system to accurately traverse routes of up to 5Km in a single day. Through the projects described in [2] [3] [4] this system has been evaluated and adapted from its original purpose (navigation systems for terrestrial autonomous vehicles) to be a viable system for the unique challenges associated with extra-terrestrial use. As part of the SEEKER trial [3] it was a key component in achieving a 5km + traverse in a single day in remote and representative trials in the Atacama Desert in Chile. This raised the Technology Readiness Level (TRL) of OVO to a point where it was considered suitable for industrialization by SCISYS UK Ltd. as part of the European Space Agency's ExoMars Rover Mission in an activity called VisLoc under direct contract to Airbus Defence and Space Ltd with Thales Alenia Space Italia as the overall mission prime contractor.

During SCISYS's subsequent development of OVO into VisLoc, SCISYS implemented numerous improvements and optimizations discussed in subsequent sections.

Following an extensive and detailed validation and verification process, VisLoc has passed Qualification Review and is awaiting final acceptance which will complete its TRL 8 level development and categorization. The next step will be operations on Mars.

Figure 1 : SEEKER system at the end of an 5km + drive in the Atacama Desert.

3 THE ALGORITHM

VisLoc obtains its position estimate from a stereo pair of images, captured from twin cameras mounted with 15cm separation low on the rover and angled so that the FoV encompasses the terrain ahead. The algorithm performs by extracting 'features' from the images, which are defined as points of interest that are distinctive enough to be identifiable over both spacial and temporal changes (for example between left and right cameras, between consecutive image pairs, and from the same vantage point after a period of inactivity). The correlation between these feature points from four sources (left-right images over consecutive image pairs) is then used to derive a model of the combined position and orientation of the rover (commonly referred to as a POSE).

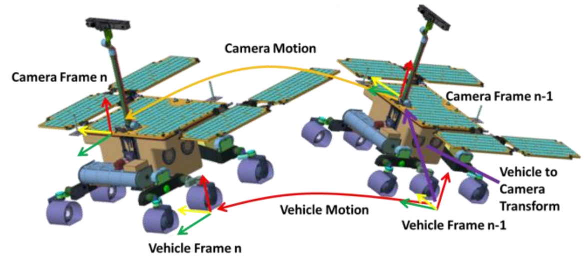

Figure 2 : Relationship between camera motion and vehicle motion

The visual localization algorithm is made of the following basic stages:

1. Image intensity balancing, to correct for exposure differences between the left-right images permit a more accurate match.
2. Corners extraction from both images and determination of feature points.
3. Temporal matching of feature points between previous and current image pairs.
4. Determination of the change in POSE between the consecutive image pairs.
5. Integration of POSE deltas to determine the estimate of the overall POSE in a 'world' reference frame.

For the different stages, various algorithms are used.

The extraction of regions of interest from the image follows a two-stage approach. Following the balancing in image intensity, 'corners' are extracted from both images using the FAST high-speed corner detector [5]. The corner detector, although extremely fast, produces a lot of corners of varying quality so a down-selection mechanism is required to reduce the quantity to a manageable number of good corners. This is done using the Harris score [6], and only the strongest corners, which must be well distributed through the images, are selected. These corners are thereafter known as 'features'.

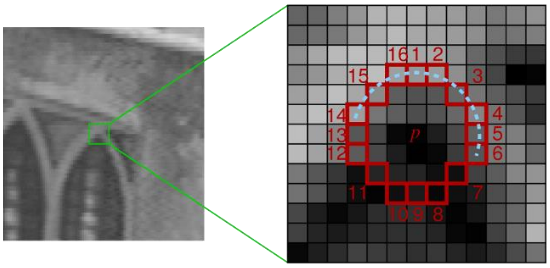

Figure 3 : Example of the FAST corner detection on an image (image courtesy of [5])

These features are matched with those propagated forward from the previous image to find points in the two images that should represent the same location.

In order to determine the POSE change from these matches, a solution must be found to the Perspective-from-three-points (P3P) problem. The solution to this problem is equivalent to finding the length of the tetrahedron connecting CP, A, B and C. (See Figure 4).

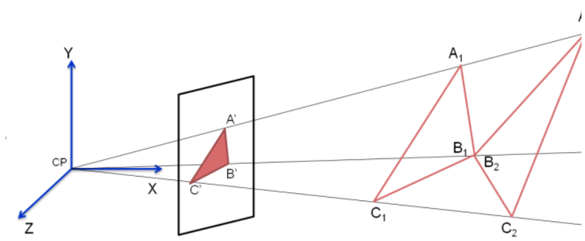

Figure 4 : Illustration of the P3P problem

A RANSAC process [7] is used in order to provide a fast and robust estimate of the motion of the rover between this frame and the subsequent one. The RANSAC process solves the P3P problem for a small random selection of matches then applies this solution to all features from the previous image in order to ascertain how well this solution describes the movement of all the matches. This process iterates until a good solution is found.

The final phase is the integration of a frame-wise POSE change into the rover's new position in the world frame of reference and the propagation of the information of the information in this frame into the next.

4 EXOMARS ROVER REQUIREMENTS

Being part of the on-board software of a rover presents unique challenges for a visual localisation system. In particular the severe limitations on memory and CPU power (by modern standards) present problems for image processing algorithms which attempt to function in real-time, as traditionally these algorithm have required powerful hardware, allowing the rapid processing of images, guaranteeing that the spacing between images will be small.

On a rover the available CPU severely restricts the image throughput. On the MER rovers (with a 20MHz CPU), a Visual Odometry update step took an average of nearly 3 minutes for a single tracking step, with a convergence rate of 95-97%. [8] The requirements for ExoMars on the other hand (with a 96MHz CPU), require an image pair to be processed in under 4.25 seconds, whilst maintaining a position estimate accurate to within 1% of distance travelled (the design goal for the MER system was 10% over 100m [8]).

Because of the nature of the rover hardware VisLoc must operate on images that are 512 by 512 pixels. Further because of the location of the camera on ExoMars and the range of optical depths VisLoc was validated against VisLoc it must be able to produce accurate estimates when significant portions of this image are usable due to shadow, sky and other parts of the rover. Further, the algorithm must not be confused by these artefacts.

The Martian terrain can be extremely bland in some locations. Given the relatively small areas of the image that may be good for VO, VisLoc must be able to extract as much information distributed across the images, picking up good corners where present and poor corners when they are the best available.

A rover can be thought of as in one of a number or operational modes including: traversing, planning or science. VisLoc is obviously critical when traversing, but must be robust to long periods, possibly hours, of inactivity between invocations. It is an operational requirement that VisLoc could not be reset over such a period of inactivity and thus must deal with any environmental changes during that period.

General requirements in space systems value provenance over progression, this means the implementation language mandated for ExoMars is a very mature C compiler and the use of libraries, including the C standard libraries, is prohibited.

5 ADAPTATIONS MADE FOR EXOMARS

To adapt the algorithm to the ExoMars requirements, required in a range of categories such as:

- Parameter tuning optimisations to improve the speed of the algorithm.
- Low level optimisations
- Addition of a 'shadow-rejection' step to prevent false positive solutions arising from the shadow of the rover.
- Extraction of more data from bland surfaces.

The modifications to improve the speed were mostly confined to parameter tuning and low level code optimisations, in order to get as much performance out of the algorithm whilst meeting the accuracy requirements.

Several key performance-affecting aspects of the algorithm were modified in order to reduce the time taken to execute whilst tuning VisLoc to the environment expected on Mars.

The OVO implementation made extensive use of a number of libraries: primarily the standard libraries, Boost and Eigen. These libraries have been highly tuned for performance but do not carry the validation required for flight software and so a number of their functions had to be implemented and validated as flight standard C code. Naïve implementations of such functions would significantly reduce the performance of the algorithm so it was critical that aspects of the algorithm were optimized whilst retaining identical behavior to the original OVO implementation.

The implemented shadow-rejection step prevents corners and features from being extracted from the edge of the rover shadow, see Figure 5.

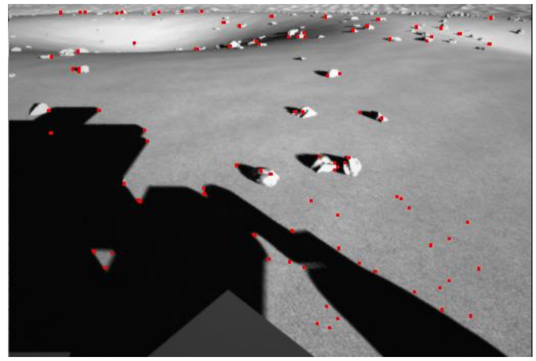

Figure 5: Shadow Rejection

Notice that whilst a large number of features (red dots) have been extracted from this image a significant proportion of them are "pinned" to the rover's shadow. The extraction of features from the edge of the rover shadow is problematic because it moves with the rover, but in a manner inconsistent with the rover's movement, and so presents a competing solution for the change in POSE which does not arise from the movement of the terrain. In the case in Figure 5 the algorithm will be robust to this. However, in edge cases where the algorithm is unable to extract sufficient genuine features this can lead to the false rover shadow features being considered inliers causing the genuine inliers to be discarded. Figure 6 shows an image more likely to be problematic.

There are numerous solutions to this problem. However, most sophisticated solutions significantly increase the run-time of the algorithm; on Mars qualified hardware even a small addition to the processing can be unacceptable. The simplest solution is to simply discard all features that are deemed to be based on shadows, whether from the rover or not. For the general case this appears to be losing information for the benefit of just a small number of edge cases. However, on Mars, it is critical that we handle these edge cases robustly, even at the cost of some general accuracy or performance.

Rover shadows are unlikely to produce useful information for VO. However, in the case of extra-territorial rovers, even terrain based shadows can be misleading. As stated in section 4, VisLoc must be robust to long periods of inactivity during which the environment may change. An example of such a change is that the terrain based shadows may move with the sun, over a period of hours this change may be significant. This generally would only introduce a small error and one that VisLoc would be robust to, but by rejecting terrain based shadows as well as rover based ones even this small error is avoided.

This shadow rejection mechanism allowed the false-positive solutions to be rejected at the first stage of the algorithm, greatly improving the robustness in challenging conditions.

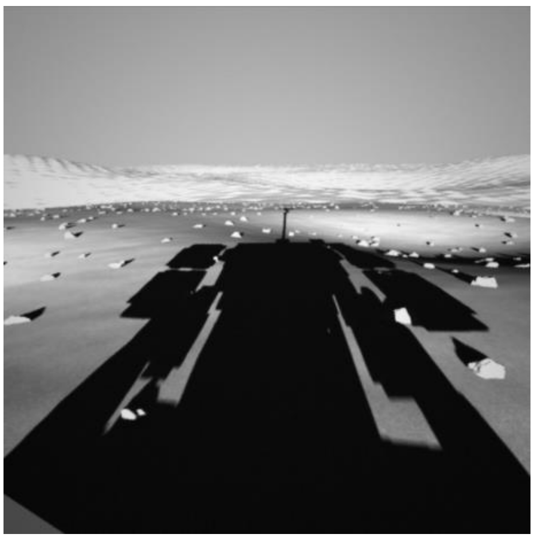

Figure 6: Terrain Detail

VisLoc had to be proven to be accurate under all specified Martian conditions. Figure 6 shows an example of one such a challenging example. We see that the top 1/3rd of the image is sky which, being featureless, is of little use for VO. VisLoc has been adapted not to carefully manage the sky portions of the image to improve performance without compromising accuracy. A second problem with this image is the large rover shadow occluding almost all of the fore and center of the image. Different sections of the image play different roles in terms of sub-POSE estimation. The algorithm has therefore been tailored to be robust to losing true feature density degradation in significant regions of the image

6 VALIDATION FOR EXOMARS

As with any other ExoMars Rover element, the VisLoc has been subject to an ESA and Prime Contractor defined V&V Process commensurate with the levels required for flight software acceptance. The algorithmic validation has been performed chiefly on a series of simulated trajectories which represent the expected conditions that will be faced by the algorithm on the surface of Mars. The test trajectories validated VisLoc under variations of numerous environmental conditions including:

- Motion Blur
- Optical Depth
- Lens Flare
- Feature Poor Terrain
- Shadowing, including from the rover
- Noise
- Impacted Dust

The trajectories also include the rover executing a number of manoeuvres including:

- Point Turns
- Crabbing
- Snaking
- Long Traverses
- Excessive Turning
- Long Periods of inactivity

The use of simulated trajectories provides a completely accurate ground truth which, when compared with the VisLoc output, allows highly accurate analysis of the rover's performance.

Each trajectory was subject to a number of requirements that capture the various operating modes of the rover: long, medium and short traverses, deployment and egress and general inter-frame accuracy. Each trajectory is validated against a computational duration requirement for each frame. These requirements were validated on over 3 million data points. These data points represent either an inter-frame estimate between two consecutive frames or an integrated estimate over a run of frames between lengths of 3 and 70 meters. In all cases VisLoc passes ExoMars' prescribed accuracy and duration requirements with a healthy margin.

Additional verification was performed using trajectories from real hardware on a Martian terrain analogue to show the performance of the algorithm under more realistic conditions. The level of analysis on these was necessarily less given the less accurate knowledge of the true rover POSE, but it still shows the important step towards validation in the real-world.

Further the algorithm was also subjected to a number of robustness tests involving the images being doctored to produce input that could potentially cause system failures. The algorithm was fed images that were blank or whited out, significantly over or under exposed, distorted or rapidly changing in contrast. The aim was not to verify the estimate; the algorithm cannot produce a good estimate in the absence of data, but to ensure that such input could not lead to a software crash.

7 HARDWARE SOLUTIONS

VisLoc's development does not end with ExoMars. VisLoc has been shown to be robust and performant for the ExoMars mission, but future missions will require the rover to be driven faster and use less power whilst remaining as accurate and robust in even more challenging conditions.

In order to meet these goals SCISYS has recently been contracted to port the VisLoc software solution to a FPGA hardware accelerated implementation. Implementing an algorithm on an FPGA effectively addresses these problems:

- Through an FPGAs the algorithm can be significantly improved in terms of execution times. In addition to the increased speed offered from a hardware solution, the algorithm can be deployed so that the processes of the algorithm can be pipelined and parallelized.
- An FPGA solution should also offer a much reduced power consumption both in the way it uses the available hardware and simply through the reduced time the algorithm requires.
- Whilst an FPGA solution may not directly lead to an improved accuracy, the algorithm is intended to produce the same estimate regardless of whether it is running on a purely software solution or is implemented in hardware. However, the decreased execution time would allow the algorithm to be re-parameterised permitting it to consider more data in a given cycle and as a result provide a more accurate estimate or to handle more challenging conditions.

The ability to run the algorithm faster means a reduced time between images samples and, in turn, the rover may travel a greater distance in a shorter time whilst retaining VO accuracy. Alternatively, an increased time between VisLoc processing cycles means the processor is free for longer to engage in other activities.

Power consumption is of critical importance for most extra-terrestrial applications and reducing it is a general goal in such missions. In the case of mars rovers the reduced power also allows the rover to drive longer, particularly earlier in the morning or later in the night, when solar panels are reduced in performance due to the reduced lighting conditions.

Being able to increase the accuracy through re-parameterisation not only allows longer runs through greater accuracy, but would also allow the rover to drive through more challenging conditions such as increased environmental dust.

8 CONCLUSION

The VisLoc project has taken the already robust and proven terrestrial OVO algorithm and implemented it to a flight standard. In doing so it has optimized, tailored and validated it for the Martian environment and use on the ExoMars rover. To this end it is completing the final phase of its Technology Readiness Level (TRL) 8 categorisation and is scheduled to run on ExoMars on the Martian surface in 2020From this position SCISYS are investigating the possibility of accelerating the algorithm using hardware in order that future missions will be able to drive longer and in more challenging conditions whilst using less power and processing resources.

References

[1] Churchill, W., & Newman, P., Practice Makes Perfect? Managing and Leveraging Visual Experiences for Lifelong Navigation, (2012) Proc. IEEE International Conference on Robotics and Automation (ICRA)

[2] Shaw, Woods et al., 2012Robust Visual Odometry for Space Exploration:, in: Proc, 12th Symposium on Advanced Space Technologies, Robotics and Automation ASTRA.

[3] Woods et al., 2014a, 014a. Seeker—Autonomous Long-range Rover Navigation for Remote Exploration. J. Field Robot. 31

[4] Woods et al., 2014b, Demonstrating Autonomous Mars Rover Science Operations in the Atacama Desert. Presented at the 12th International Symposium on Artificial Intelligence, Robotics and Automation in Space – i-SAIRAS

[5]Faster and better: a machine learning approach to corner detection, Edward Rosten, Reid Porter, and Tom Drummond.

[6] C. Harris and M. Stephens (1988). "A combined corner and edge detector" (PDF). Proceedings of the 4th Alvey Vision Conference. pp. 147–151.

[7] Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography, M. Fischler et al., Communications of the ACM, 1981.

[8] JPL. Mars Exploration Rovers. Available at: http://marsrovers.jpl.nasa.gov/home/index.html.