ABSOLUTE LOCALISATION BY MAP MATCHING FOR SAMPLE FETCH ROVER

Michael Dinsdale[(1)]*, Warren Hamilton[(1)], Róbert Marc[(1)], Piotr Weclewski[(1)], Alain Dysli[(1)], Chris Barclay[(1)], Anthonius Daoud-Moraru[(1)], Ben Brayzier[(1)], Todd Cooper[(1)], Max Braun[(1)], Duncan Robert Hamill[(1)], Daniel Aspinal[(1)], Guillaume Du Roy[(2)], Vivien Croes[(1)]

[(1)]Airbus Defence and Space Ltd., Guidance Navigation and Control Dept., Stevenage, SG1 2AS, United Kingdom, Email: GNC.UK@airbus.com
[(2)]Airbus Defence and Space SAS, AOCS/GNC Advanced Studies Dept., Toulouse, 31400, France

ABSTRACT

The challenging autonomy demands of the Sample Fetch Rover mission require an on-board Absolute Localisation algorithm. This paper describes an algorithm that uses cross-correlation of rover-generated maps and reference data stored on-board to localise the rover. This algorithm is presented through its high-level working principles of operations, up to its verification and validation and finally its first preliminary test results obtained from simulations.

1. INTRODUCTION

Mars Sample Return is a joint NASA/ESA campaign consisting of three missions working together to return samples from Mars to Earth by the early 2030s. As part of this campaign, Sample Fetch Rover (SFR) is tasked with driving from the Sample Return Lander 2 (SRL2) landing site to a sample depot where the samples have been placed by NASA's Mars 2020 Perseverance rover, collecting the samples, and transporting them to SRL1 in a timely manner. SRL1 will then place the samples onto the Mars Ascent Vehicle, which will carry them into orbit [1].

SFR is being developed by Airbus Defence and Space Ltd. as prime contractor, and its GNC (Guidance, Navigation and Control) subsystem [2] is heavily based on heritage from the ExoMars Rover Vehicle GNC [3]. This GNC subsystem features a Relative Localisation system based largely on visual odometry. However, the increased length of autonomous drives required by the SFR mission means that its rover position accuracy targets cannot be met with Relative Localisation alone - an on-board Absolute Global Localisation (AGL) algorithm is needed.

AGL is required both during the traverse between SRL and the sample Depot (AGL-Traverse or AGL-T) and during the tube acquisition phase of the mission within the sample depot (AGL-Depot or AGL-D). The required accuracy is very different: approximately 10m for AGL-T, based on the width of the expected traversable regions, and 10cm for AGL-D, based on the accuracy required for sample tube pick-up. This accuracy does not, however, have to be met starting from a "kidnapped robot" state where the rover location within the operational environment is completely unknown. Initial good estimates can be assumed to be provided by ground at the start of a period of autonomous driving, and the AGL algorithm must maintain these without further ground input.

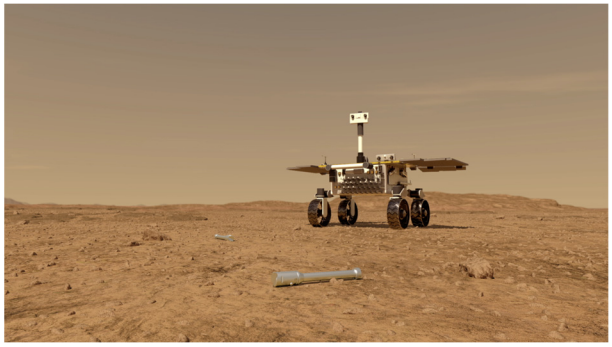

Figure 1 - Artist's concept of the Sample Fetch Rover approaching sample tubes (Credit: ESA [4])

This paper presents the algorithm, which Airbus have developed to implement this functionality and describes results of some of the testing which has been performed so far.

2. PRINCIPLE OF OPERATION

Due to the rover's high-level architecture constraints, the only on-board sensors that the rover can use for AGL are its cameras. Comparing rover images with reference image data stored on-board should allow localisation of the rover position with respect to those reference images. For AGL-T, reference images are available from Mars Return Orbiter's (MRO) HiRISE instrument [5], with ground resolution of ~0.3m. Similarly, for AGL-D, the M2020 rover will generate reference images as it constructs the sample Depot. The resolution of these images is expected to vary from a few mm near the sample tubes, to a few cm over the wider Depot (due to geometrical considerations from the camera's field of view and images' resolution). In both cases, these images will also be used for mission planning, so localising the rover with respect to the images will allow it to execute the planned activities. This is important because it means that there is no need to localise the reference images themselves to 10m/10cm accuracy with respect to a Mars-fixed frame.

In comparing the images, we must deal with the difference in vantage point. Because all three sets of images (SFR, M2020 and HiRISE) consist of stereopairs, a stereo-vision algorithm can be used to construct a point cloud representation of each scene. This may be orthogonally projected onto the horizontal plane, creating a co-aligned Digital Elevation Map (DEM) and Ortho-Rectified Image (ORI), which we refer to generically as "maps". The projection simplifies the comparison of the point cloud data without losing much information because of the rarity of overhangs in the Martian terrain the rover will encounter.

This approach is particularly attractive for SFR because the heritage ExoMars GNC already builds a 4cm resolution DEM for autonomous navigation and hazard detection. Hence, an ORI can be constructed from the same stereo-vision point cloud without significant processing overhead. The benefit of comparing both the DEM and the ORI is that they are complementary: large elevation variations can cause shading and shadowing differences when the light direction changes, degrading the ORI comparison, but they also provide a strong signal in the DEM.

While driving, SFR stops regularly to generate maps. Depending on GNC mode, this occurs every 2.7-5.0m and the maps have a field of view of 110-170 deg and a range of ~8m. In case this is not enough data for the map comparison, the AGL can combine maps from a series of stops before performing the comparison.

Once the combined map is ready, comparison with the reference maps can be used to infer the likely position of the rover. In principle, the comparison could also give information about the rover heading, which is expected to accrue some error during a traverse due to the limited accuracy of Visual Odometry. In practice, we fix the heading when comparing the maps, but use the drift in position over the traverse to learn about the rover heading. This is achieved by passing the result of the map comparison to an estimator, which fulfils several objectives:

1. Infer rover heading
2. Reject outlier map comparison measurements
3. Incorporate data from Relative Localisation
4. Incorporate once-per-sol Sun-Sensing Heading Estimation (SSHE) measurements. These provide an independent source of heading knowledge, which is valid even if the terrain is too featureless for map matching.

The output of this is an estimate of the transform between the Estimated Mars Local Geodetic (EMLG) frame, which is space-stabilised by Relative Localisation (RelLoc), and the frame in which the reference maps (RM) are given. If RelLoc was perfect this transform would be fixed, but it drifts due to RelLoc errors. The EMLG frame is used by SFR for Navigation and Path Planning. Therefore the transformation can be used, for example, to transform a commanded path from the RM frame into the EMLG frame where it can be driven using RelLoc without being affected by the long-term accumulating RelLoc drift.

A high-level overview of the algorithm is given in Fig. 2, below.

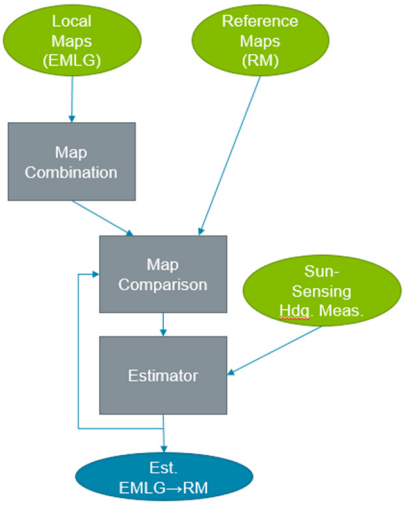

Figure 2 Overview of AGL Algorithm

The following sections describe each component in more detail.

3. MAP COMBINATION ALGORITHM

The rover map data is constructed in the RelLoc-stabilised EMLG frame and accumulated into a combined map buffer centred on the rover position. The ORI data is normalised to make it approximately independent of image exposure and the DEM elevations are expressed in the approximately fixed EMLG frame, so the map data which is being combined should usually be consistent. However RelLoc errors and illumination changes (for example when the combined map stretches over multiple sols) can still occur. To mitigate these when the existing and new data overlap, an additional correction is applied to the existing data to remove the mean discrepancy on the overlap region. The overlapping data itself can be set to a weighted average of existing and new data, but in practice because the new data is usually closer to the rover and hence more accurate we simply use the new data.

4. MAP COMPARISON ALGORITHM

There are many possible approaches to align a pair of images. One group of approaches begins by detecting features in each image, and then aligns those features. In a Mars rover context, matching of ORI features has been used by ground to localise the MER Spirit rover [6]. Another group of approaches computes a pixel-by-pixel metric expressing the difference between the images, and chooses the alignment that minimises this difference. This has the benefit of implementation simplicity. As searching over rotations is not needed for SFR, the amount of computation is acceptable.

Approaches of this type have been used for DEM alignment in several papers. Ref [7] used DEM alignment in the first phase of their localisation process. Their matching metric is a weighted combination of the RMS deviations in elevation and slope. They give an example where the localisation error relative to a manual localisation of MER-A (Spirit) near the Home Plate area is 0.22m, but they note that errors will be larger in relatively flat areas. Ref [8] presents an experimental test of DEM alignment using a normalised cross-correlation metric. Using a 1m resolution reference DEM like for HiRISE, and 15m range rover DEMs assembled within a 30x30m area, they obtained ~2m accuracy after combining measurements with a particle filter. However, their reference 1m/pixel DEM was created by down-sampling a 5cm/pixel DEM created from UAV images. The vertical precision may therefore be superior to that of the HiRISE DEMs. In the SFR context, [9] and [10] propose aligning DEMs with an un-normalised cross-correlation of DEM slopes, for an occasional global pose correction to their SLAM-based localisation. In a field test with their rover on Tenerife they found localisation accuracies better than their global DEM resolution (<0.5m). Returning to Mars rover data, [11] compared 280 MSL DEMs to HiRISE DEMs using an Iterative Closest Point matching approach. They found (for example) that with a 16m radius local DEM and 1m DEM resolution 79.2% of localisations had an error <5m.

ORI alignment based on normalised cross-correlation was examined in [12]. They tested their algorithm in simulations of a lunar rover equipped with a LIDAR sensor. Using 25cm/pixel orbital images (similar to the HiRISE images we would use for AGL-T) they found a 15x15m rover ORI to give better than 1m accuracy at 94% for search windows up to 100m, provided the illumination direction was unchanged from the orbital images. For an equatorial sun position, performance remains roughly level until the solar elevation deviation reaches 40° (*N.B.* SFR will operate at a latitude of 18.4°).

As in several of the mentioned implementations, the SFR AGL uses a zero-mean normalised cross-correlation. This is applied to the gradients (X- and Y-differences) of both the DEM and the ORI. This has the following key advantages:

- It works even if the amount of overlap between the maps varies with the alignment to be tested.
- It is naturally resilient to varying data quality over the maps, in contrast to e.g. a sum-squared-difference criterion, which may move the maps to ensure the overlap is in a low-noise area. This is useful because it is complex to rigorously quantify the error in the maps produced by stereo-vision.
- The zero-mean criterion makes it largely unaffected by overall DEM slope errors (which produce a shift in the gradient)
- The normalisation, although not so important for the DEM, allows for the fact that the absolute ORI magnitude depends on lighting and atmospheric conditions, as well as details of the camera, which could in principle be accounted for but at the cost of some complexity.
- Using the same algorithm for DEM and ORI simplifies the algorithm development, tuning and validation.

As already noted, the AGL does not search over rotations. However, generally there is an arbitrary rotation between the EMLG and RM frames. The combined map is rotated using the latest estimate of this rotation to align it as well as possible with the RM frame before comparison.

5. ESTIMATOR ALGORITHM

The estimator component of the algorithm is a particle filter (see for example [13]), modelling a probability distribution over EMLG→RM frame transformations plus a pair of parameters describing bias in the Visual Odometry increments. Specifically, for a reference point in the EMLG frame near the rover the particle state variables are:

1. The X and Y coordinates of the reference point in the RM frame
2. The heading rotation between the EMLG and RM frames
3. A scale factor error on the RelLoc increments
4. A heading drift rate

The propagation step of the estimator depends only on the EMLG-frame position change since the last update. A tuneable process-noise term whose variance is proportional to the magnitude of the position change is added to each component to model the RelLoc drift uncertainty.

The estimator undergoes a separate update process whenever a DEM or ORI measurement is made. The position of the correlator peak (with quadratic interpolation to obtain the sub-cell part) gives the best- fit alignment of the rotated combined map and reference map. Because the true rotation is unknown, the alignment must be turned into a measurement of the RM-frame coordinates of a specific “measurement point” (MP) in the EMLG frame. The MP is set to the centre of the bounding box of the combined map, which should be positioned more accurately than a point near the edge assuming features are spread across the map. Each particle’s weight is updated based on how close the particle’s transform places the MP to its observed RM-frame location with the best-fit alignment. The weight formula behaves as a Gaussian likelihood for small offsets, but is saturated for large offsets to limit the effect of outliers - this is the only part of the current design that uses the power of the particle filter (compared to e.g. a more compact Extended Kalman Filter design). The parameters of the update weight can be configured to depend on the peak-to-root-mean-square ratio of the

correlator (a kind of signal-to-noise estimate) as this is observed to correlate with measurement quality. The estimator can also be updated with SSHE measurements [2]. This uses a similar saturated update weight depending on the particle's rotations only.

6. VERIFICATION AND VALIDATION

So far we have carried out off-line tests of the AGL algorithm using prototype C code. Flight-representative LEON4 hardware [14] was used for a subset of tests to assess execution time. Rover and reference maps were generated from a mixture of sources including simulation, HiRISE and rover field trials. Details of a subset of these tests and their preliminary results are given in the following sections.

The next phase of testing focuses on incorporating the AGL outputs into the full GNC chain, leading to field trials in Airbus' Mars Yard and planetary analogue environments.

7. EARLY AGL-T COMPARISON RESULTS

For AGL-T the performance depends on the quality of the HiRISE maps, as well as the quality of the rover maps which varies with range from the rover. To get a first idea of the feasibility of AGL-T, it is possible to cross-correlate overlapping HiRISE maps generated from different stereo-pairs. Fortunately, such an overlapping pair exists for part of the Jezero region. The width of the region to correlate (perpendicular to the traverse direction) must not produce rover map errors exceeding the HiRISE ones. Results of this analysis for a 10m x 15m ORI and 40m x 15m DEM are shown as the blue lines on Fig. 3. In this case, the maximum range of the rover map would be 7.5m.

For these map sizes the HiRISE errors should be larger than SFR errors, so one might think the HiRISE-to-HiRISE comparison would be conservative. However MRO's orbit means a particular location is always imaged at a similar time-of-sol, limiting illumination changes. Moreover the ground point-of-view of the rover may cause systematic differences. To account for this we can compare actual rover maps from Curiosity [15] to HiRISE maps of Gale crater. See Fig. 6 for data from one ORI comparison with an SFR-like max range of 8m. A good DEM matching requires an increased range - see for example Fig. 4 (15m range).

To transfer this analysis to the Jezero terrain where SFR will drive, we have used Curiosity data to tune a model which up-scales and adds noise to the HiRISE maps. This relies on the fact that there are overlapping HiRISE maps for part of Curiosity's traverse.

First, Curiosity panoramas with varied terrain were selected. ORI matching of a full panorama at 10m range was used to obtain a "ground truth" alignment. The panorama maps were then sliced in angle and range to create a set of maps with some good and some bad matches. Then the upscaling/noise parameters were tuned so that the application of the model to one of the overlapping HiRISE maps reproduced the observed performance of the Curiosity comparisons (using the other HiRISE map as a reference). See Fig. *5* for example results. Applying this to the Jezero HiRISE comparison process gives the results shown in magenta on Fig. 3.

These results show the ORI matching performs much better than the DEM matching. A substantial subset of the DEM matching results are distributed uniformly across the searched alignments, meaning they convey no information. For a more precise analysis it is necessary to focus on the expected SFR traverse paths, which are smoother and less textured than the Jezero average. ORI matching remains successful with a frequency, which allows the estimates to meet the AGL-T requirements, but DEM matching, provides negligible benefit. Therefore, we expect AGL-T to use only the ORI.

Simulations of the full AGL-T performance (including the estimator) using this HiRISE map up-scaling model along a subset of possible SFR paths show good position accuracy. In addition, the average LEON4 execution time per rover stop is smaller than 5s which is within rover timing budgets.

A limitation of this analysis is that it requires two overlapping HiRISE maps, which are available for only a fraction of SFR's possible operational area. It might seem that one could use the difference of overlapping HiRISE maps to estimate the HiRISE error, and then add this to another map to test the comparison where overlapping maps are not available. However, this turns out to give extremely optimistic results. To see this one can repeat the previous DEM comparison analysis replacing the comparison between DEM2 and DEM1 with a comparison between DEM2 and DEM2+R(DEM1-DEM2) where R is a random shift. In that analysis, ~80% of samples have an error < 1.5m. This occurs because the noise in DEM2 acts as a feature, which AGL-T can use to align the maps.

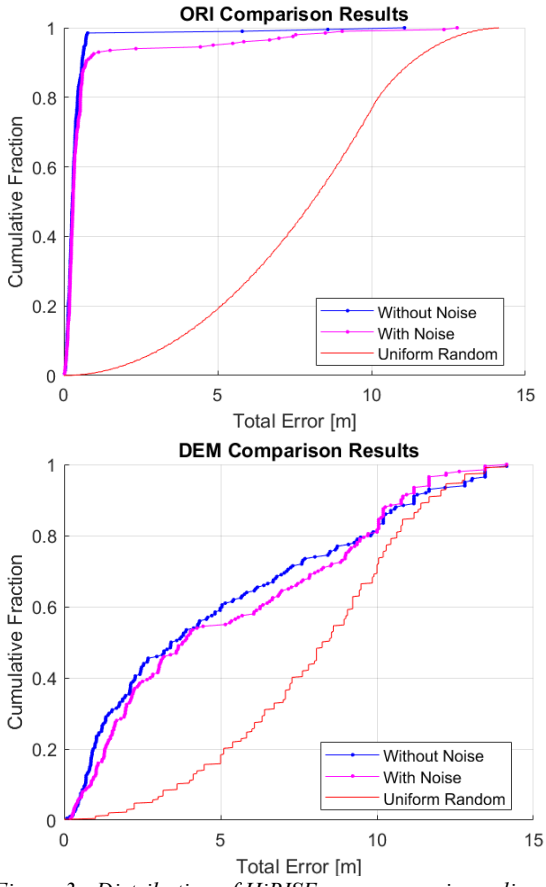

*Figure 3 - Distribution of HiRISE map comparison alignment errors. **Top**: ORI with 15m x 10m. **Bottom**: DEM with 15m x 40m. The red lines arise from random measurements uniformly distributed over the searched window, conveying no information.*

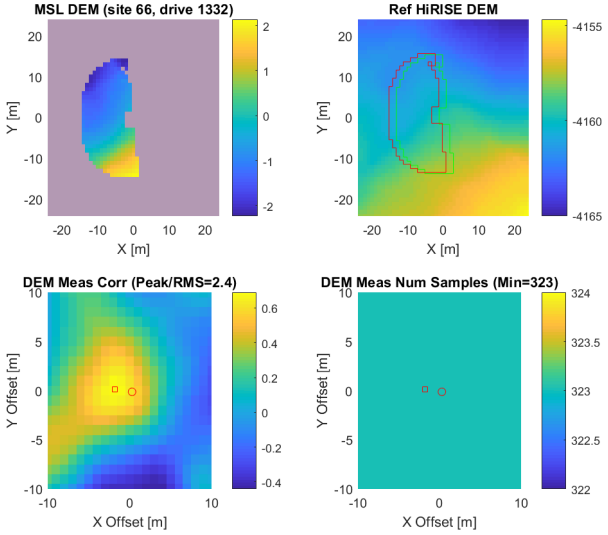

Figure 4 - Example of Curiosity DEM correlations inputs and outputs with a large range of 15m.

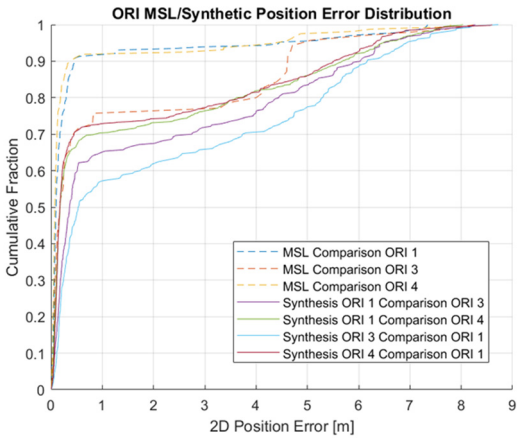

Figure 5 - Example of measurement errors for MSL to HiRISE comparison and synthetic MSL (based on HiRISE) to HiRISE ORI comparison.

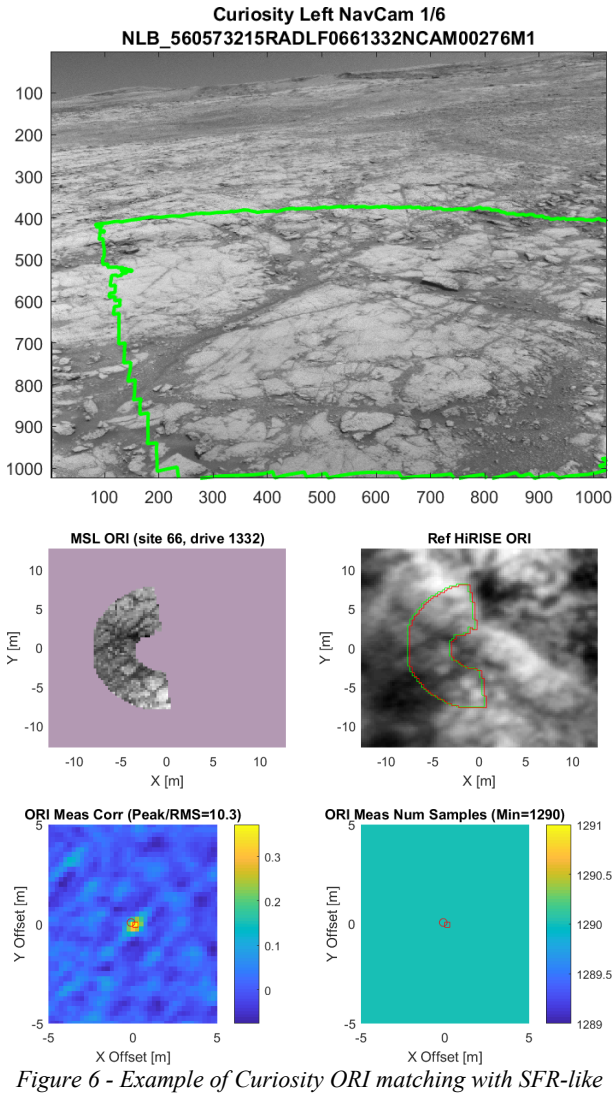

*Figure 6 - Example of Curiosity ORI matching with SFR-like range of 8m. **Top**: example input NavCam image. The green contour shows the part of this image used to build the ORI. **Bottom**: correlation inputs and outputs.*

8. EARLY AGL-D COMPARISON RESULTS

Two data sources have been used to test in-Depot map comparisons, as described in the following subsections.

8.1. Erfoud Trial Results

One set of data comes from the rover trial described in [16]. This consists of a set of NavCam-analogue stereo-pairs with cm-level GPS ground-truth localisations, along with UAV ortho-images of the test sites. A UAV DEM is also available, but the resolution is not sufficient to resolve detail at the 4cm scale.

Comparison of ORI matching and GPS data shows a ~2m offset, varying slightly as the rover moves and turns. Fitting the offset to an 8-parameter model (2D global affine transform + rover-frame translation) reduces the residual error for strongly correlating ORIs to ~1-2cm. The fit is carried out separately for each traverse. Fig. 9 gives examples of the residual error distribution.

Figure 7 - Erfoud test trials site called "Mummy" in Morocco

8.2. SFR Offline Simulator Environment

Simulated maps may be created using the SFR Offline Simulation Environment (OSE). This includes models of the rover hardware (including the Inertial Measurement Unit, cameras, pan-tilt unit and locomotion subsystem) as well as the GNC algorithms. Ideal camera images are rendered using [17] before being degraded to account for hardware-specific effects.

The simulated maps have the advantage of controllability. It is also possible to produce a reference DEM that was missing from the Erfoud tests. To produce reference maps, the rover is driven through the simulated terrain, and the local maps are combined by averaging in a crude representation of the process of M2020 imaging the Depot. Local rover maps are then generated by driving the rover along 3 more paths.

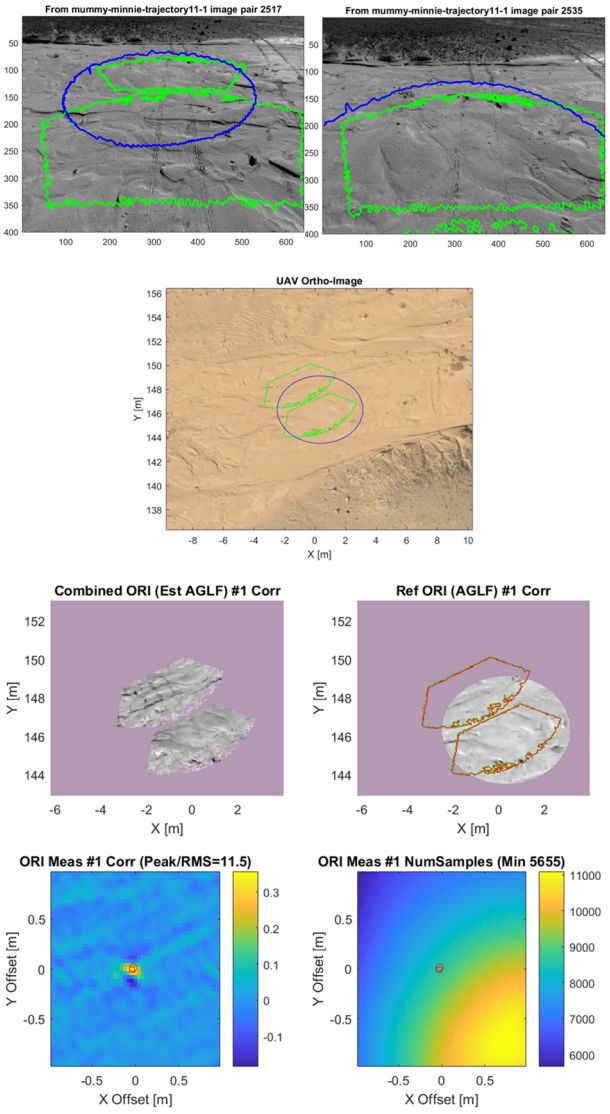

Figure 8 – Example results from Erfoud test trials in Morocco

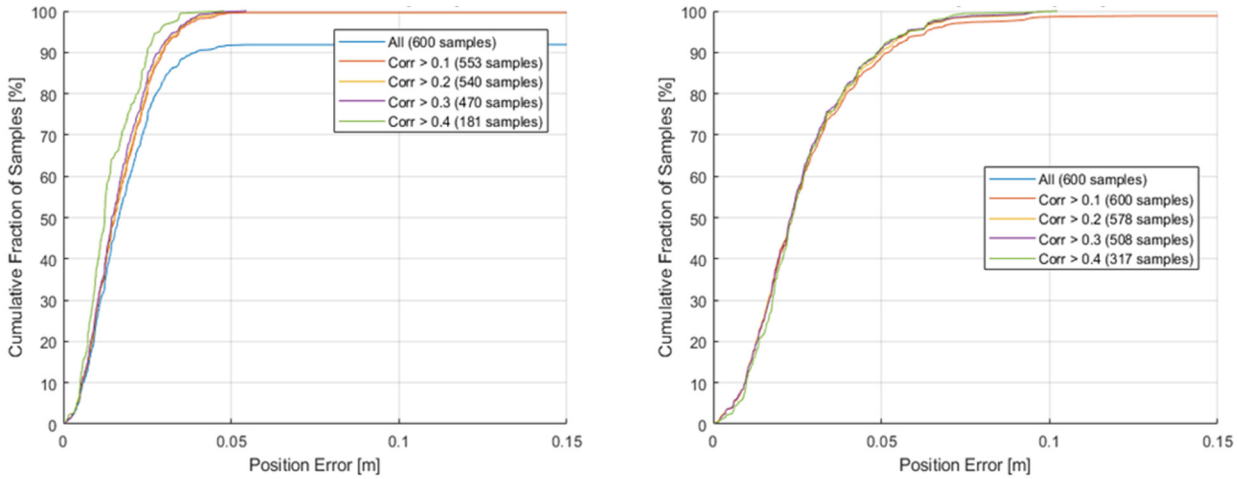

Figure 9 - Residual errors between ORI comparison positions and GPS data for 2 Erfoud traverses

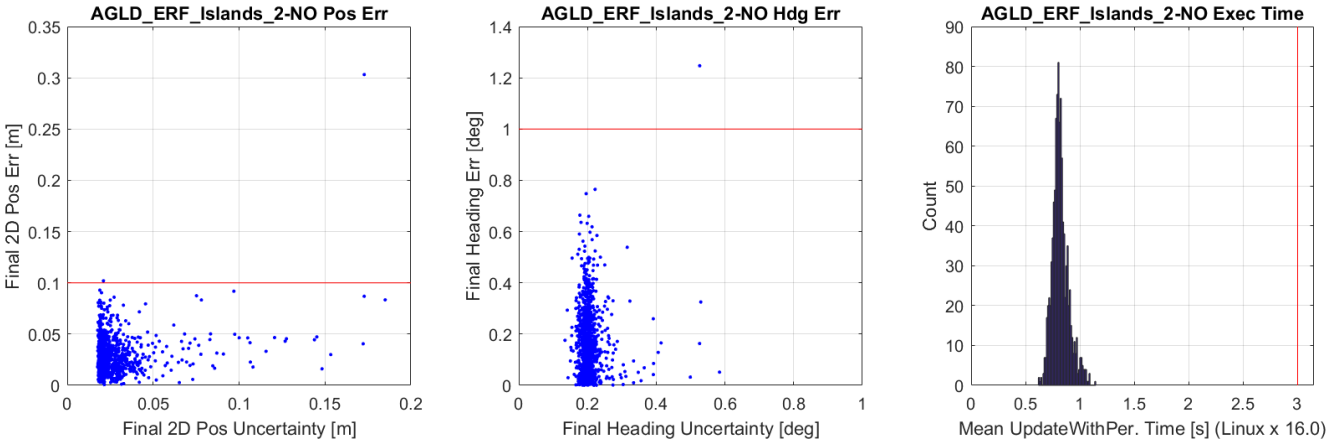

*Figure 10 - Final AGL-D estimation error when reaching the mapped area around a tube. The island radius is 2.8m. **Left**: 2D position error. **Centre**: heading error. **Right**: mean AGL-D algorithm execution time on flight-representative LEON4 HW (does not include local map generation).*

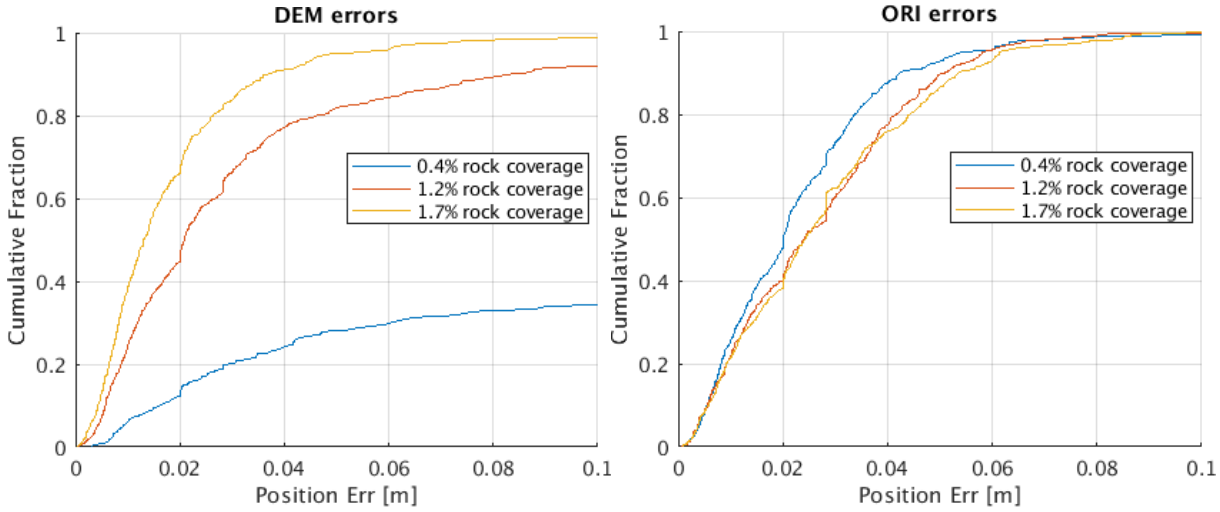

Figure 11 - Example of DEM (left) and ORI (right) matching errors for SFR OSE maps.

Matching errors are shown in Fig. 11 for 3 scenarios with varying rock density/size (though the *maximum* rock diameter is 20cm in both cases because sample tubes will not be placed near large rocks). See Table 1 for details. This gives an example of the complementarity of DEM and ORI matching - larger rocks increase ORI errors when the sun direction differs, but they also allow DEM matching to succeed.

Table 1 - Example rock coverage scenarios (added to flat base DEM) for OSE testing

Depot Rock Scenario	Mean Local Map Fractional Rock Coverage	Mean Local Map Max Rock Radius
1	0.39%	7.1cm
2	1.17%	8.9cm
3	1.71%	9.2cm

9. EARLY AGL-D ESTIMATION RESULTS

The Depot reference maps will be constructed from M2020 rover imagery. Due to bandwidth limitations it may not be possible to map the full area of the Depot. However, the tight 10cm position accuracy requirement only holds in the vicinity of a sample tube. This enables a conservative scenario where an "island" around each sample tube is mapped. The rover drives between samples using RelLoc, and then AGL-D updates with absolute position information once the island is reached.

We have simulated the full algorithm in this scenario, using a model of Visual Odometry performance based on earlier breadboarding. Fig. 10 shows some example results for 1000 20m island approaches, using Erfoud data. The position estimation performance is acceptable and the average run time of the algorithm (including propagation when driving towards the island, and the updates when it is reached) is < 1.5s.

10. CONCLUSIONS

Despite the challenging SFR mission constraints and context, we have developed a single algorithm which, provided with appropriate reference data, appears to meet SFR's absolute localisation needs both in the traverse between SRL and the Depot and within the Depot itself. Adding to the algorithm performance measurement, the average execution time on the LEON4 hardware has been measured as <5s on all scenarios, which is compliant with the rover timing budgets.

ACKNOWLEDGEMENT

The authors would like to thank the European Space Agency for financing the SFR project, and in particular Martin Azkarate and Juan Manuel Delfa Victoria for their support and guidance in the SFR GNC activity.
We would also like to thank Roland Brochard for introducing us to the use of cross-correlation for image matching, and engineers at NASA JPL for their support.

The SFR project received funding from the European Space Agency (ESA) under contract 40000124005/18/NL/PA for its A/B1 and Advanced B2/B2X.

REFERENCES

[1] B. K. Muirhead, A. Nicholas and J. Umland (2020), Mars Sample Return Mission Concept Status, *IEEE Aerospace Conference*, 2020, pp. 1-8
[2] P. Weclewski, R. Marc, B. Brayzier, W. Hamilton, C. Barclay, M. Dinsdale, D. Hamill, A. Daoud-Moraru, A. Dysli, M. Braun, T. Cooper, D. Aspinal, G. Du Roy V. Croes (2022), Sample Fetch Rover Guidance, Navigation and Control Subsystem - An Overview, *ASTRA 2022*
[3] N. Silva, R. Lancaster, J. Clemmet (2015), Exomars Rover Vehicle Mobility Functional Architecture and Key Design Drivers, *ASTRA 2015*
[4] https://www.esa.int/Science_Exploration/Human_and_Robotic_Exploration/Martian_rover_motors_ahead
[5] A. McEwen (2009), Mars Reconnaissance Orbiter High Resolution Imaging Science Experiment, Digital Terrain Model, MRO-M-HIRISE-5-DTM-V1.0, *NASA Planetary Data System*
[6] R. Li, et al. (2011), MER Spirit rover localization: Comparison of ground image– and orbital image–based methods and science applications, *J. Geophys. Res.,* 116, E00F16
[7] JW. Hwangbo, K. Di, and R. Li (2009), Integration of Orbital and Ground Image Networks for the Automation of Rover Localization, *ASPRS 2009 Annual Conference*
[8] B. Van Pham, A. Maligo, and S. Lacroix (2013), Absolute Map-Based Localization for a Planetary Rover, *12th Symposium on Advanced Space Technologies and Automation in Robotics*
[9] M. Azkarate, L. Gerdes, L. Joudrier and CJ. Perez-del-Pulgar (2019), A GNC Architecture for Planetary Rovers with Autonomous Navigation Capabilities, *https://arxiv.org/abs/1911.09975*
[10] D. Geromichalos, M. Azkarate, E. Tsardoulias, L. Gerdes, L. Petrou, CJ. Perez Del Pulgar (2020), SLAM for autonomous planetary rovers with global localization, *Journal of Field Robotics*, **37** (5)
[11] AV. Nefian, L. Edwards, D. Lees, L. Keely, TJ. Parker and M. Malin (2017), Automatic rover localization in orbital maps, *Lunar and Planetary Science* **XLVIII**
[12] A. Sheshadri, KM. Peterson, HL. Jones and WL. "Red" Whittaker (2012), Position Estimation by Registration to Planetary Terrain, *Proceedings of International Conference on Multisensor Fusion and Information Integration (MFI)*
[13] HR. Künsch (2013), Particle Filters. *Bernoulli* **19** (4)
[14] J. Andersson, M. Hjorth, F. Johansson, S. Habinc, (2017), LEON Processor Devices for Space Missions: First 20 Years of LEON in Space, *Proceedings of 6th International Conference on Space Mission Challenges for Information Technology*
[15] M. Justin (2013), MSL Mars Navigation Camera RDR V1.0, NASA Planetary Data System, MSL-M-NAVCAM-5-RDR-V1.0
[16] S. Lacroix, A. De Maio, Q. Labourey, E. Paiva Mendes, P. Narvor, V. Bissonette, C. Bazerque, F. Souvannavong, R Viards and M. Azkarate (2019), The Erfoud dataset: a comprehensive multi-camera and Lidar data collection for planetary exploration, *ASTRA 2019*
[17] IM. Martin et al. (2019), Planetary surface image generation for testing future space missions with PANGU, *2nd RPI Space Imaging Workshop in Saratoga Springs*